

Adversarial-Cost Screening

A Validation Framework for Export Control, Sanctions, and Supply-Chain Counterparty Programs

PUBLISHED

13 AUGUST 2026

STATUS

WORKING PAPER – SSRN

SUPERSEDES

VERSIONS 1.0-1.6

VERSION

1.7

AUTHOR

COLLIN B. GEORGE, CISSP

LICENSE

CC BY 4.0

UNCLASSIFIED // OPEN SOURCE

Paper record

ABSTRACT

Counterparty screening programs in export control, sanctions, and defense supply-chain compliance are validated almost exclusively by process metrics — alert volume, clearance rate, list-coverage currency, screening latency. None of these measure whether the program detects the conduct it exists to detect. This paper identifies three structural defects common to these programs and specifies a validation framework that corrects them. The defects are: (1) endogenous label generation, in which outcome data is observable only for counterparties the program already flagged, making measured performance a function of the screening policy rather than of the threat; (2) entity resolution as the unmeasured binding constraint, in which screening accuracy is capped by the ability to recognize that nominally distinct counterparties are one operation, a capability no program reports on; and (3) adversary-controlled attribute dependence, in which screening weight concentrates on attributes the screened party can change at negligible cost while structurally costly attributes go unused. The framework specifies an exogenous sampling requirement, an entity-resolution maturity measure, an attribute-cost tiering scheme, and an evaluation battery that produces defensible estimates of screening effectiveness. It applies the result to reasonable-care determinations under EAR Part 732, to sanctions-compliance testing under OFAC's Framework, and to scienter and materiality analysis under the False Claims Act, including cybersecurity certification under the Civil Cyber-Fraud Initiative, where recent settlements supply worked illustrations of each failure mode.

KEYWORDS

False Claims Act; Civil Cyber-Fraud Initiative; export controls; economic sanctions; counterparty screening; third-party risk; NIST SP 800-171; CMMC; program evaluation; strategic classification; compliance effectiveness; materiality; scienter.

JEL CLASSIFICATION

K23 (Regulation and Business Law); K42 (Illegal Behavior and the Enforcement of Law); D82 (Asymmetric and Private Information); H56 (National Security).

STATUS

Preliminary draft; comments welcome. This version supersedes versions 1.0–1.6. Published concurrently on SSRN.

HOW TO READ THIS PAPER

Claim convention

Assertions are marked where the basis is not self-evident: **[F]** fact traceable to a cited primary instrument (statute, regulation, court record, agency release); **[L]** established proposition in a cited body of scholarship; **[A]** assessment, with confidence stated; **[J]** analytic or design judgment, not a measured result. Unmarked expository text carries no independent evidentiary weight. Named parties appear only as to what the cited public record establishes; settlements resolved allegations without adjudication of liability and, where noted, without admission of wrongdoing.

Disclaimer

This paper is analytic work product. It is not legal advice, creates no attorney–client or consulting relationship, and reflects the author's independent views. It is not an official government product and reflects no agency position. Nothing here should be relied upon without independent verification against primary authority and counsel of record for a specific matter.

1 The Problem With How Screening Programs Are Currently Validated

1.1 Process metrics measure activity, not effectiveness

A representative export-compliance program reports: number of transactions screened, list-update currency, alert volume, false-positive rate, average clearance time, training completion, audit findings closed. Every one of these is an input or throughput measure. The output measure — did the program prevent or detect diversion — appears nowhere, and is generally treated as unmeasurable.

[A, high confidence]

It is not unmeasurable. It is unmeasured because the measurement requires a design element that no program includes.

1.2 Endogenous labels

Screening programs learn from their own history. Counterparties that generated alerts were investigated; investigation outcomes became training data, precedent, and rule refinements. Counterparties that generated no alert were cleared without inquiry and contribute no evidence in either direction.

The consequence is that measured performance is conditional on the screening policy. This is selection on the dependent variable — sampling on the outcome one is trying to explain — and it biases every downstream estimate. [L — Heckman, *Sample Selection Bias as a Specification Error*, 47 *Econometrica* 153 (1979); King, Keohane & Verba, *Designing Social Inquiry* (1994).] A program that has never flagged a particular counterparty profile has no basis to assert that profile is low risk, and its clean record on that profile is not

evidence of anything. Reported hit rates, false-positive rates, and clearance accuracy all inherit this conditioning.

This is not a subtle statistical objection. It is the reason compliance programs are routinely surprised by enforcement actions involving counterparties they screened and cleared many times. [A] The clearance was not a finding; it was an absence of inquiry.

1.3 Entity resolution

Diversion networks, sanctions-evasion structures, and front-company arrangements churn corporate identity far faster than they churn the underlying operating substrate: people, physical premises, freight forwarders, banking relationships, and equipment. [A, high confidence] A screening program that treats each corporate identity as an independent observation discards the only stable signal available.

Every screening program performs entity resolution implicitly. Almost none measure how well. The absence of a resolution-quality metric means the binding constraint on program performance is invisible to program management and to auditors. [A]

1.4 Adversary-controlled attributes

Screening weight concentrates on the attributes that are cheapest for a sophisticated counterparty to change:

Attribute	Cost to change [J]
Corporate name	Days, nominal fee
Registered address	Days
Declared end-use / end-user statement	Zero
Declared commodity, ECCN self-classification	Zero
Beneficial ownership below the 50% threshold	Weeks, restructuring cost
Nominee directors and officers	Days

Against these, the structurally costly attributes are largely unscreened:

Attribute	Cost to change [J]
Operating history and transaction tenure	Years
Physical premises and verified footprint	Months, capital
Relationship tenure with forwarders, banks, brokers	Years
Network position relative to previously actioned entities	Not controllable
Deviation from structurally comparable peer counterparties	Not observable to the subject

A screening attribute's value is not its predictive power. It is its predictive power discounted by the cost to the adversary of falsifying it. [J] This is the strategic-classification problem: a deployed classifier is observed by its subjects, who move the cheap features to defeat it. [L — Hardt, Megiddo, Papadimitriou & Wootters, *Strategic Classification*, ITCS 2016.] High predictive power on a zero-cost attribute is therefore a liability: it produces measured accuracy against unsophisticated counterparties and near-zero accuracy against the counterparties that matter, while generating documentation that the program appeared to be working.

Denied-party list screening is the canonical instance. It is highly accurate against entities that have not changed their names and structurally uninformative against entities that have. [A]

2 Framework

2.1 Component 1 — Exogenous sampling requirement

Specification. The program shall subject a defined fraction of screened counterparties or transactions to enhanced review selected by stratified random sampling, independent of any screening score, alert, or analyst judgment. Stratification is on jurisdiction, commodity or service class, transaction value band, and counterparty tenure. The fraction is set so that the expected number of sampled units per stratum per period supports interval estimation at the program's stated confidence level; 1–3% of transaction volume is an illustrative starting range, validation-pending, not a derived requirement [J].

Precedent. This is the design the IRS uses in its National Research Program, which draws randomly selected returns for detailed examination specifically to estimate the tax gap independent of the enforcement-selection policy. [F — IRS National Research Program.] The transfer to trade and sanctions screening is the contribution; the sampling method is established.

Function. The sample is not an enforcement mechanism. It is the sole source of unbiased base-rate estimates, the denominator against which alert-driven performance can be compared, and the only means of detecting whether illicit activity has displaced into segments the screening rules do not reach.

Assertion enabled. With the sample, the program can state: among counterparties our screening cleared, the estimated rate of undetected disqualifying conduct is X, with confidence interval Y. Without it, the program can state nothing about cleared counterparties at all. That distinction is dispositive in an enforcement or litigation posture. [A]

Cost. The sampling fraction is small and the reviews are ordinary enhanced due diligence. The objection to this component is never cost; it is that it generates a documented estimate of program failure rate. That objection should be recognized as an argument for the component, not against it. The scienter implications are addressed in §3.3, which explains why structured measurement is the defensible posture rather than the dangerous one.

2.2 Component 2 — Entity resolution maturity

Specification. The program shall maintain a persistent resolved-entity graph rather than a transaction-keyed counterparty list. Nodes: legal entity, natural person, physical address, freight forwarder, financial-institution relationship, ultimate consignee, end user, registered agent, corporate officer. Edges: shared identifiers and co-occurrence on transactions.

Maturity levels. [J]

Level	Description
0	Exact string matching on counterparty name against lists
1	Fuzzy matching with transliteration and alias handling
2	Identifier-based linkage across records (address, tax ID, officer)
3	Probabilistic record linkage with adjudicated resolution queue; persistent entity identifiers survive corporate re-registration
4	Graph analytics over the resolved network; risk propagation from actioned entities; community detection

Reported metric. Resolution rate — the fraction of screened transactions linked to a persistent entity identifier — and adjudication backlog. Programs below Level 3 should state so explicitly in compliance representations, because their screening accuracy against re-registered entities is materially unknown.

Practice note. OFAC 50% Rule analysis is a graph-traversal problem and is not solvable at maturity levels 0–2. [A] A program conducting 50% Rule analysis without a resolved ownership graph is performing the analysis it can perform, not the analysis the rule requires. The same holds for FOCI determination and for NSPM-33 foreign-talent-affiliation review.

2.3 Component 3 — Attribute cost tiering

Specification. Every screening attribute shall be classified by adversarial cost and the classification documented in the program's screening methodology.

Tier A — negligible cost to falsify. Declared identity, declared end-use, self-classification, stated destination, declared value, self-attested compliance scores. Permissible as interaction terms and as consistency checks against Tier B/C attributes. Aggregate weight capped by policy regardless of measured predictive performance.

Tier B — material cost to falsify. Operating tenure, verified physical footprint, relationship tenure with intermediaries, transaction-pattern consistency with stated business, financial-instrument history, licensing history.

Tier C — not controllable by the subject. Network distance to previously actioned entities, shared-identifier collisions with dissolved entities, deviation from structurally matched peer counterparties, third-party-reported attributes the subject does not supply.

Rationale for capping Tier A weight. [J] A screening program optimized purely on historical accuracy will load weight onto Tier A attributes, because historical enforcement data is dominated by unsophisticated actors who did not falsify them. The resulting program is optimized against the population least likely to cause consequential loss. The cap is a deliberate acceptance of lower measured accuracy in exchange for adversarial robustness, and it must be documented as such or it will read as a defect on audit. This inverts standard feature selection, which chooses on predictive power alone; the departure rests on the strategic-classification premise in §1.4.

Tiering is a judgment, reviewed on a defined cadence. The cost assignments in §1.4 are analytic judgments, not measurements, and change as registry digitization, corporate-formation practice, and verification-service availability change. Review no less than annually.

2.4 Component 4 — Status decay

Specification. Any status that reduces screening intensity — approved end-user, prior license history, trusted-trader enrollment, certified supplier, accepted remediation plan, prior favorable due diligence — shall carry an explicit validity period and behavioral triggers that suspend it prior to expiry.

Triggers. Ownership or officer change; registered-address change; new intermediary relationships; transaction-pattern deviation from the profile on which the status was granted; adverse network events among linked entities; jurisdictional risk change; failure to meet a committed remediation milestone.

Rationale. [A] Status is granted on a point-in-time assessment and applied prospectively to indefinite future transactions. The value of a status to an adversary is exactly proportional to the screening reduction it confers, which makes the most attractive acquisition target the status the program trusts most. Programs that grant durable status without decay have constructed the highest-value target in their own control environment.

Application. This maps onto C-TPAT and FAST in trade, onto approved-end-user and license-exception reliance in export control, onto supplier-certification flowdown under CMMC and DFARS 252.204-7012, onto plan-of-action-and-milestones (POA&M) forbearance in cybersecurity attestation, and onto institutional assurances under NSPM-33. The POA&M instance is worked in §3.4.

2.5 Component 5 — Evaluation battery

Replace process metrics with:

1. **Estimated clearance failure rate.** From the exogenous sample: the rate of disqualifying findings among counterparties the screening cleared. The program's actual error rate, and the only figure that answers the question an enforcement authority will ask.
2. **Lift over random.** Alert-driven detection rate divided by the sample-derived base rate. A lift near 1.0 means the screening rules contribute nothing beyond random selection — a common and undiagnosed condition. [A]
3. **Coverage of unscreened space.** Fraction of transaction volume, by jurisdiction and commodity class, that has received zero non-random enhanced review in the trailing period. Growth is the leading indicator of displacement.
4. **Entity resolution rate and backlog** (§2.2).
5. **Tier A weight share.** Fraction of decisions driven predominantly by Tier A attributes. A robustness measure, reported as a trend.
6. **Time to novel-pattern detection.** Interval between first appearance of a structural pattern in program data and first program-initiated review of it. Measures whether the program can detect anything it was not already configured to detect.

Alert volume, false-positive rate, and clearance latency remain useful for resourcing. They are not effectiveness measures and should not be presented as such.

3 Application to Reasonable Care and to False Claims Act Exposure

3.1 EAR reasonable care and the red-flag framework

EAR Supplement No. 3 to Part 732 sets out red flags and an affirmative duty to inquire, and warns firms against structuring themselves so that red-flag information is cut off before it reaches the officials responsible for evaluating it — conduct the guidance terms “self-blinding.” [F — 15 C.F.R. pt. 732, supp. no. 3.] The regulation thus already takes the position that structured non-inquiry is not a safe harbor — a point that recurs with force in §3.3 and §3.4.

In practice, adequacy of an inquiry system is assessed retrospectively against a known bad outcome, which invites hindsight bias in both directions: programs that missed a diversion are found deficient, programs that were never tested are presumed adequate. [A] The framework supplies a prospective

standard. A program that maintains an exogenous sample, reports a clearance failure rate with a confidence interval, operates at a stated entity-resolution maturity level, and documents its attribute cost tiering can demonstrate the adequacy of its inquiry system independent of any particular outcome. A program reporting only alert volume and clearance rate cannot; its representations of reasonable care rest on process compliance rather than on evidence of effectiveness.

3.2 Sanctions compliance

OFAC's Framework for OFAC Compliance Commitments (May 2, 2019) identifies five essential components of a sanctions compliance program — management commitment, risk assessment, internal controls, testing and audit, and training — without specifying testing methodology. [F — OFAC, *A Framework for OFAC Compliance Commitments* (2019).] The exogenous sample is the natural instrument for the testing-and-audit component: it is the only design that produces an estimate of undetected exposure rather than an assessment of control execution. Programs that test only alert handling are testing the branch of the process where errors are already visible. [A]

3.3 False Claims Act — scienter and materiality

Where a contractor certifies compliance with export-control, supply-chain-security, or cybersecurity requirements as a condition of payment or eligibility, the certification's factual basis is discoverable, and the screening methodology becomes evidence.

The governing standards. *United States ex rel. Schutte v. SuperValu Inc.*, 598 U.S. 739 (2023), holds that the FCA's scienter element turns on the defendant's subjective knowledge and beliefs, not on what an objectively reasonable person would have believed; the statutory term “knowingly” reaches actual knowledge, deliberate ignorance, and reckless disregard. [F] *Universal Health Services, Inc. v. United States ex rel. Escobar*, 579 U.S. 176 (2016), holds that a misrepresentation must be material to the government's payment decision, that materiality is a “demanding” and “rigorous” standard, and that government payment in full despite actual knowledge of noncompliance is strong evidence the requirement is not material. [F]

These are distinct elements, and the *SuperValu* remand illustrates the distinction — not a materiality holding. On remand, *Schutte v. SuperValu*, No. 3:11-cv-03290 (C.D. Ill.), was tried to a jury over three weeks in February 2025, with the verdict returned March 5, 2025. The court's instructions required relators to prove both that the defendant knew its claims were false and that the false claims caused damages to the government. The special verdict form records that the jury found knowing submission of false claims but found no damage causation, producing a defense verdict; the district court subsequently denied relators' motion for a new trial. [F] The case therefore failed on causation of damages, a separate element, not on materiality, and not on scienter — which relators established.

Two cautions on how much that verdict carries. It does not establish a rule that the absence of damages defeats FCA liability: the statute imposes civil penalties per false claim independent of damages, 31 U.S.C. §3729(a); the verdict operated as it did because the verdict form structured damages causation as an

element of the liability determination, and judgment was entered on that form. Commentary has treated the framing as case-specific rather than doctrinal. [A] The proposition the remand supports is narrow and sufficient: subjective knowledge can be fully established and the defendant can still prevail on a downstream element. [A]

The strategic problem, stated precisely. A screening program that measures its own clearance failure rate creates a record capable of establishing the subjective-knowledge element under *Schutte*: the program knew, to a quantified degree, that some cleared counterparties carried disqualifying conduct. A program that never measures cannot have that specific knowledge attributed to it from the measurement — but it remains exposed to the deliberate-ignorance and reckless-disregard prongs, which *Schutte*'s definition of “knowingly” expressly reaches, and which a documented decision not to examine a known blind spot tends to support. EAR Supplement No. 3's “self-blinding” language (§3.1) is the same principle in the export-control regime.

Why this argues for measurement, not against it. The naive reading — do not measure, so as not to manufacture knowledge — is wrong for three independent reasons, and this section is not to be read as counseling non-measurement.

First, scienter is not the only element, and *Escobar* materiality is a partly independent defense that rewards disclosure. A contractor that measures, discloses its stated failure rate, and can show the government continued to pay and to contract with that knowledge builds *Escobar* materiality evidence in its own favor — on *Escobar*'s continued-payment reasoning, not on *SuperValu*, which decided a different element. Suppressing measurement forfeits that defense to preserve a scienter posture the deliberate-ignorance prong already undermines.

Second, the knowledge that creates exposure is knowledge that the program was ineffective, not knowledge that some residual risk exists. A program that measures a failure rate and acts on it — tightening screening where the sample shows leakage — converts the measurement from an admission into evidence of diligence. The adverse posture attaches to measuring and doing nothing.

Third, the alternative to measurement is not safety. It is a compliance certification with no evidentiary basis, which collapses the moment the underlying conduct surfaces and offers no materiality defense when it does.

For expert application. The framework supports opinion testimony on whether a program's screening design was capable of detecting the conduct at issue — distinct from whether it did detect it, and a question on which process-metric documentation is silent.

3.4 Cyber certification and the Civil Cyber-Fraud Initiative

Cybersecurity certification is where the FCA docket is currently most active, and where the framework's materiality argument is strongest, because the certified object is a program, not an outcome.

The mechanism. The DOJ established the Civil Cyber-Fraud Initiative in October 2021 to use the FCA against cybersecurity fraud by contractors and grant recipients. [F — DOJ, Civil Cyber-Fraud Initiative (2021).] Under DFARS 252.204-7012 and NIST SP 800-171 — applicable to DoD contracts since 2017 and continuing under the CMMC program — contractors self-attest their control implementation through a summary assessment score reported in SPRS. [F — DOJ, Georgia Tech Research Corp. settlement release (Sept. 30, 2025).] The score represents the state of a security program. Enforcement has proceeded, and settlements have resolved allegations, with no breach, exfiltration, or attack alleged: the asserted false statement is the attestation, not the occurrence of a compromise. [A] (These matters resolved by settlement; allegations were not adjudicated.)

The cases map directly onto the framework's failure modes.

LOGZONE, Inc. (N.D. Ala.) (FCA settlement, June 18, 2026, \$507,144, of which \$253,572 restitution; resolved by settlement agreement without a filed complaint or relator): in October 2021, LOGZONE submitted a perfect self-assessment score of 110 to SPRS; in February 2024, a Defense Contract Management Agency DIBCAC assessment of its NIST SP 800-171 implementation returned a score of –170, at the low end of the –203-to-110 range. [F — DOJ settlement agreement and release.] A self-attested 110 against a government-assessed –170 is a 280-point gap on a single Tier A attribute. This is the cleanest public-record instance of the §2.3 failure mode precisely because of the corrective score's provenance: it came from a government assessment rather than the contractor's own consultant — the scored object was a self-declaration with no verified footprint behind it, and independent government assessment inverted it. No breach was alleged; the covered conduct was submitting claims while knowing the DFARS 252.204-7012 requirement was unmet. [A/J]

United States ex rel. Craig v. Georgia Tech Research Corp., No. 1:22-cv-02698-JPB (N.D. Ga.) (qui tam filed July 2022; United States intervened February 2024; complaint-in-intervention filed August 2024; settled September 30, 2025, \$875,000): the alleged falsity was a December 2020 summary score of 98 premised on a “fictitious” or “virtual” environment corresponding to no actual system that processed covered defense information. [F] In the framework's terms this is a pure Tier A failure — a self-attested declared value with no verified footprint behind it — and it is exactly the failure §2.3 predicts when the scored object is a self-declaration the subject controls. [A/J]

Pennsylvania State University (settled October 22, 2024, \$1.25 million, without admission of wrongdoing): the enforcement theory did not require an inflated score. The covered conduct, per the government, was failing to implement contractually required controls and failing to adequately develop and implement the plans of action committed to correct the deficiencies the institution had itself identified, across fifteen DoD and NASA contracts, alongside noncompliant use of an external cloud service provider. [F — DOJ release and settlement agreement.] The POA&M strand is a status-decay failure (§2.4) in its purest form: a POA&M is a granted forbearance that should decay on non-performance, and the enforcement theory is that the granted status outlived the conduct justifying it. [A/J]

United States ex rel. Berich v. MORSECORP, Inc., No. 23-cv-10130-GAO (D. Mass.) (settled March 26, 2025, \$4.6 million; \$851,000 to the relator): MORSE submitted a January 2021 SPRS score of 104 — near

the top of the –203-to-110 range — then, after a July 2022 third-party gap analysis returned an actual score of –142 reflecting roughly 22% of NIST SP 800-171 controls implemented, did not correct the reported score until June 2023, three months after a federal subpoena. [F] A self-attested 104 against an actual –142 is a 246-point gap on a single Tier A attribute, and it is a second stark public-record illustration of the §2.3 failure mode: the scored object was a self-declaration the subject controlled, and no verified footprint sat behind it. The matter also fixes the temporal point in §2.4 — status is assessed as of the attestation, and the delayed correction, not any later remediation, was the exposure. Reported as the first FCA settlement premised on failure to promptly update a self-assessment score after a third-party assessment returned a lower one, it makes the §2.4 status-decay mapping predictive rather than retrospective: the enforcement regime has only recently begun charging this failure mode. The remediation trajectory is itself status decay in one line — MORSE moved from the stale 104 to a third-party score of 57 (June 2023), 82 (October 2023), and 110 (May 2024) — none of which cured the attestation that preceded it. [F/A]

The §3.3 materiality strategy is already being litigated — and remains untested. Georgia Tech did not immediately settle. Its motion to dismiss argued materiality directly: that DoD did not treat cybersecurity compliance as the essence of the bargain, that DoD never asked GTRC to verify its assessment score, and that DoD continued paying on the contracts after learning of the alleged noncompliance. [F] That is precisely the continued-payment argument §3.3's first reason recommends, run in live litigation. DOJ opposed, the matter went to mediation, and it settled with no ruling reported on the motion — so the strategy is demonstrated but unadjudicated. Its ceiling is also bounded: continued government payment is probative of immateriality but not dispositive. At the pleading stage in *United States ex rel. Cimino v. IBM Corp.*, 3 F.4th 412 (D.C. Cir. 2021), the court declined to resolve the point — holding it plausible that the government continued paying for reasons that did not render the alleged fraud immaterial, while noting continued payment could bear on immateriality later, at summary judgment. [F] The measurement-and-disclosure posture §3.3 recommends builds the evidence for this defense; it does not guarantee it.

The materiality argument, in this domain. Where the government pays for a security program rather than a security outcome, *Escobar* materiality turns on whether the program the contractor operated was the program the government paid for. A self-attested SPRS score, unaccompanied by any measurement of whether the attested controls actually detect anything, is a Tier A artifact: it documents a program without evidencing one. The evaluation battery in §2.5 is what distinguishes a documented program from a functioning one, and that distinction is where both materiality (was the deficiency material to the award or payment decision) and the expert question (was the program capable of detecting the conduct at issue) are decided.

The exposure is widening, on the figures that actually bear on cyber. The FY2025 record of roughly \$6.8 billion in total FCA recoveries is about 83% healthcare and does not measure cyber exposure. The cyber-specific figures do: DOJ's FY2025 FCA Fact Sheet reports over \$52 million recovered across nine cybersecurity fraud settlements, with civil cyber settlements having more than tripled over the past two years, against a record 1,297 qui tam filings overall — the highest ever, against the prior record of 980 — meaning most of the pipeline is whistleblower-initiated and does not depend on agency-initiated

enforcement. [F — DOJ, FY2025 FCA Fact Sheet (Jan. 16, 2026).] The Administrative False Claims Act enacted in December 2024 adds an independent agency mechanism for smaller cyber lapses below DOJ's suit threshold. [F]

The CMMC program was designed to move the certified object from self-attestation toward third-party assessment and thereby narrow the Tier A gap the LOGZONE, Georgia Tech, and MORSECORP matters exploited. On July 13, 2026, DoD suspended the Phase 2 third-party certification requirement, and subsequent implementation milestones, pending a 60-day program review by a newly established CMMC Reform Task Force; Phase 1 self-assessment, SPRS score reporting, executive affirmation, and DFARS 252.204-7012 obligations remain in force. [F — USD(A&S) Memorandum 26-P-1023 (July 13, 2026).] For the interim, the certified object has reverted to a self-attested SPRS score backed by an executive affirmation — the precise Tier A configuration the settled matters arose from — with DIBCAC assessment and FCA enforcement as the operative external checks. The suspension does not narrow the exposure this section describes; it extends the regime that produced it. [A]

4 Validation Case: High-Volume Land-Border Cargo Targeting

Reserved. The framework was derived in part from analysis of a high-volume commercial-cargo screening environment in which all three defects are present and the data conditions are unusually observable: transaction volume is very large, the screening authority controls both the referral policy and the outcome measurement, status-based facilitation programs are mature and well documented, and an independent terminal outcome measure exists. That case will be developed as a separate companion piece against public-record facts only, with no named parties and no operational referral logic. It is illustration, not the paper's subject.

LIMITATIONS

What this paper does not establish

1. The exogenous sample estimates the failure rate against conduct that enhanced review can detect. It does not bound exposure to conduct that no available review method detects, and must not be represented as doing so.
2. Entity-resolution maturity above Level 2 requires data many commercial programs cannot lawfully obtain in every jurisdiction, particularly beneficial-ownership and officer data. The maturity scale should be reported per jurisdiction, not as a single figure.
3. Attribute cost tiering is a judgment, not a measurement (§2.3), and must be reviewed on a defined cadence.

4. Tier C attributes derived from network position carry a base-rate problem: association with an actioned entity is weak evidence in dense commercial networks. Propagation should be distance-decayed and validated against the exogenous sample.
5. The framework assesses screening design capability. It does not establish that any particular program did or did not detect any particular transaction.
6. The FCA analysis in §§3.3–3.4 describes a strategic structure under current authority. It is not a prediction of outcome in any matter and is not a substitute for matter-specific advice of counsel. The named FCA matters are cited only as to what their public records establish; settlements resolved allegations without adjudication of liability.

CONTRIBUTION

Author's note on contribution

None of the underlying methods are novel in their home disciplines: selection on the dependent variable is elementary program evaluation; randomized audit sampling to recover unbiased base rates is the IRS National Research Program design; feature manipulation under a deployed classifier is the strategic-classification literature; entity resolution is standard AML infrastructure. The contribution is the translation of these into export-control, sanctions, and supply-chain compliance validation, where none has been absorbed, together with two specific claims — that screening-attribute weight should be capped by adversarial falsification cost even where predictive power is high, and that the FCA scienter/materiality structure makes structured measurement the defensible posture rather than the dangerous one. Reviewers in the methods disciplines will find the imported machinery familiar; the target audience will not. The paper is written for the second audience and cites the first honestly.

SOURCES

References — primary authority

STATUTES

False Claims Act, 31 U.S.C. §§3729–3733.

Administrative False Claims Act of 2024 (amending the Program Fraud Civil Remedies Act, 31 U.S.C. §3801 et seq.).

CASES AND SETTLED MATTERS

Universal Health Services, Inc. v. United States ex rel. Escobar, 579 U.S. 176 (2016).

United States ex rel. Schutte v. SuperValu Inc., 598 U.S. 739 (2023).

United States ex rel. Schutte v. SuperValu Inc., No. 3:11-cv-03290 (C.D. Ill.) (jury verdict Mar. 5, 2025; new-trial motion denied 2025).

United States ex rel. Craig v. Georgia Tech Research Corp., No. 1:22-cv-02698-JPB (N.D. Ga.) (settled Sept. 30, 2025).

United States ex rel. Berich v. MORSECORP, Inc., No. 23-cv-10130-GAO (D. Mass.) (settled Mar. 26, 2025).

LOGZONE, Inc., False Claims Act settlement (no complaint filed; settlement agreement and DOJ release, June 18, 2026; USAO N.D. Ala.).

United States ex rel. Cimino v. IBM Corp., 3 F.4th 412 (D.C. Cir. 2021).

REGULATIONS AND AGENCY INSTRUMENTS

BIS “Know Your Customer” Guidance and Red Flags, 15 C.F.R. pt. 732, supp. no. 3.

OFAC, *A Framework for OFAC Compliance Commitments* (May 2, 2019).

NIST SP 800-171, *Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations*.

DFARS 252.204-7012, *Safeguarding Covered Defense Information and Cyber Incident Reporting*.

Cybersecurity Maturity Model Certification (CMMC) Program, 32 C.F.R. pt. 170.

USD(A&S) Memorandum 26-P-1023, suspending CMMC Phase 2 implementation pending program review (July 13, 2026).

National Security Presidential Memorandum 33 (NSPM-33) (Jan. 14, 2021).

U.S. Dep't of Justice, Civil Cyber-Fraud Initiative (announced Oct. 6, 2021); DOJ settlement releases: Georgia Tech Research Corp. (\$875,000, Sept. 30, 2025); Pennsylvania State University (\$1.25M, Oct. 22, 2024); MORSECORP Inc. (\$4.6M, Mar. 26, 2025); LOGZONE, Inc. (\$507,144, June 18, 2026).

U.S. Dep't of Justice, *Fact Sheet: False Claims Act Settlements and Judgments, Fiscal Year 2025* (Jan. 16, 2026) (total recoveries ~\$6.8B; 1,297 qui tam filings; healthcare ~83%; over \$52M across nine cybersecurity settlements).

IRS National Research Program (tax-gap measurement via randomized examination).

SCHOLARSHIP

J. Heckman, *Sample Selection Bias as a Specification Error*, 47 *Econometrica* 153 (1979).

G. King, R. Keohane & S. Verba, *Designing Social Inquiry* (1994).

M. Hardt, N. Megiddo, C. Papadimitriou & M. Wootters, *Strategic Classification*, ITCS 2016.

Secondary practitioner analyses were used only to locate the primary instruments above and are not relied upon as authority. Citations, regulatory sections, and settlement figures in this paper were verified against primary or authoritative agency sources.

SUGGESTED CITATION

George, Collin B. *Adversarial-Cost Screening: A Validation Framework for Export Control, Sanctions, and Supply-Chain Counterparty Programs*. Working Paper SCREEN-01, version 1.7. Sanctir LLC, 13 August 2026.

AUTHOR

About the author

Collin B. George, CISSP, is the principal of Sanctir LLC, an independent research and advisory practice working on CMMC and NIST SP 800-171, export controls, sanctions, and defense industrial base risk.

Sanctir is a solo practice. This paper was subjected to adversarial self-review rather than external peer review, and is not affiliated with any government agency, academic institution, or defense contractor.

ORCID 0009-0007-8162-6839 • Full background • Contact