Chapter 2: Why Qualitative Diligence Fails

> *The memo identified all the right risks. It just couldn't tell us why they mattered.*
> *— anonymous IC member, post-mortem review, 2023*

The Change Healthcare diligence work performed during the 2021–2022 review window has not been published. What can be reconstructed from public-record diligence engagements at comparable transactions of the period is the methodology that prevailed: a questionnaire-driven assessment supplemented by external rating-service data, target-CISO interviews, attestation review against SOC 2 and HITRUST controls, and a final memorandum graded on a four-tier severity scale culminating in the recommendation that cyber findings be addressed through "post-close remediation" or "standard rep-and-warranty provisions." This is the qualitative cyber diligence playbook, and it dominates M&A practice through 2026.[1]

The playbook is not a failure of effort or expertise. The professionals running it are competent. The findings they identify are typically real. The remediation recommendations they make are typically reasonable. The methodology fails at a different level: the outputs are not denominated in units the investment committee can use, are not comparable across deals, and are not defensible at the moments the deal team needs them to be defensible.

This chapter documents that failure structurally. Five distinct failure modes are described, each with reference to specific Big Four practice methodologies and commercial cyber-risk-quantification platforms whose published approaches exemplify the failure pattern. The chapter closes by introducing the methodological correction the rest of the book documents — quantitative cyber diligence that produces dollar-denominated, percentile-resolved, comparable, defensible outputs — and forward-referencing the dual-band reporting structure introduced in chapter 10 that operationalizes this correction at the IC level.

2.1 The output denomination failure

The first and most consequential failure of qualitative cyber diligence is that its outputs are not denominated in units the investment committee can use to make deal decisions.

A typical Big Four cyber due diligence deliverable presents findings on a four-tier severity scale: critical, high, medium, low. Deloitte's Cyber Due Diligence service, marketed as part of the firm's broader Mergers, Acquisitions, and Divestiture Services, structures its reporting around "key cyber risk areas" with heat-map presentations and "remediation priority" rankings.[2] EY-Parthenon's cyber transaction services explicitly position their offering as "valuing cyber risk for specific events" and "providing insight into the cybersecurity posture of acquisitions and divestitures."[3] KPMG advertises a "30-to-60-day cyber due diligence window" producing a

[1] PwC, *Global M&A Industry Trends*, 2026.
[2] Deloitte, *Cyber Due Diligence Services*, 2024 marketing material.
[3] EY, *Cyber Transaction Services*, 2024 service description.

"summary report on cyber risks identified" with "recommended remediation actions."[4] PwC's Cyber M&A Diligence service offers "clean-room data analysis" and "post-close cyber integration planning" alongside its diligence reporting.[5]

The qualitative outputs of these engagements typically include:

A finding inventory — typically 30 to 80 specific cyber findings categorized by severity. Each finding describes an observed gap (e.g., "Multi-factor authentication is not enforced for administrative access to the customer-facing platform") and assigns a severity level.

A severity heat map — a two-dimensional grid plotting findings by likelihood and impact, with color coding from green to red. The grid presents findings visually but does not translate them to dollar figures.

A remediation cost estimate — a rough projection of what would be required to address the identified findings, typically expressed as a range (e.g., "$1.5M to $3.2M over 12-18 months") with limited granularity per finding.

A summary recommendation — narrative text recommending a specific deal action, most commonly "address through standard rep-and-warranty provisions" or "include cyber-specific covenants in the post-close integration plan." The recommendation is qualitative and does not specify exact escrow sizes, indemnity caps, or rep-and-warranty insurance sub-limits.

The investment committee receiving this output cannot perform several operations the deal architecture requires. The IC cannot compare the cyber exposure to the target's quality-of-earnings adjustment in equivalent units. The IC cannot size a cyber-specific escrow against a defensible distributional figure. The IC cannot evaluate whether the target's cyber posture warrants a price adjustment versus rep-and-warranty insurance versus walking. The IC must instead either accept the qualitative recommendation at face value, ignore the cyber finding entirely (the modal IC choice when the recommendation is "address through standard reps and warranties"), or commission additional analytical work that produces dollar-denominated figures — typically not feasible within the deal's diligence timeline.

This is not a failure of the cyber diligence team. The professionals running the engagement have produced exactly what their methodology specifies. The methodology, however, is misaligned with the IC's decision needs. A financial diligence team that produced its quality-of-earnings analysis as a four-tier severity heat map would be replaced; a cyber diligence team that does so is the industry default.

2.2 The comparison failure

The second failure mode is that qualitative cyber diligence outputs do not compare across deals.

---

[4]KPMG, *Cyber M&A Services*, 2024.
[5]PwC, *Cybersecurity in Mergers and Acquisitions*, 2025.

A PE platform evaluating ten potential acquisitions in a calendar year, or evaluating one acquisition against the sponsor's existing portfolio of portcos, requires that cyber findings be expressible in commensurable units. Severity ratings are not commensurable. A "high severity" finding at Target A and a "high severity" finding at Target B may correspond to materially different dollar-denominated exposure depending on the target's data holdings, vendor concentration, regulatory environment, and operational scale. The qualitative methodology produces no mechanism for normalizing across these dimensions.

Commercial cyber-risk-quantification platforms have attempted to address this gap. C-Risk, a French CRQ platform, markets FAIR-based quantitative analysis aimed specifically at M&A diligence.[6] Kovrr offers cyber risk quantification with M&A-applied use cases.[7] ThreatNG markets external attack-surface quantification with deal-applied tooling.[8] CyCognito provides external surface intelligence with risk scoring.[9] Each of these platforms produces dollar-denominated outputs of some form.

The platforms address the denomination failure but not the comparison failure. Each platform's methodology is proprietary, internally consistent within its own engagements but not specified in detail externally, and not calibrated against a published reference standard. A C-Risk output is comparable to other C-Risk outputs; a Kovrr output is comparable to other Kovrr outputs; the two are not comparable to each other. A PE platform using mixed CRQ vendors across its diligence engagements faces a comparison problem the vendors do not solve.

The framework documented in this book addresses the comparison failure directly. The five-pillar decomposition produces commensurable per-pillar outputs across deals. The CCOD aggregation produces a single dollar-denominated figure expressed as percentage of enterprise value. The dual-band reporting structure produces three figures (central, contained, escalated) with explicit anchor band selection. Practices applying the framework against different targets produce comparable outputs, calibrated against the framework's published reference parameters and updated through the engagement-evidence learning loop documented in chapter 17.

2.3 The defensibility failure

The third failure mode is that qualitative cyber diligence outputs are not defensible at the IC, in front of regulators, or in post-close validation.

A defensible analytical output has four properties: explicit methodology, calibrated parameters, documented assumptions, and reconcilable outputs. Each property fails for typical qualitative cyber diligence.

Methodology is rarely documented at the level of specificity required for an IC to question it. A finding rated "high severity" in a Big Four memo represents the engagement team's professional judgment, but the criteria for that judgment vary across engagement teams within the same firm

[6]C-Risk, *FAIR for M&A Cyber Diligence*, 2024.
[7]Kovrr, *Cyber Risk Quantification for M&A*, 2025.
[8]ThreatNG, *External Attack Surface Quantification*, 2025.
[9]CyCognito, *Risk Intelligence for M&A*, 2024.

and across firms. Two competent diligence teams can produce divergent severity ratings for the same finding without either team being wrong. The output is not reproducible.

Calibrated parameters are absent. Severity ratings have no calibrated probability or magnitude. The remediation cost ranges are estimates without uncertainty bands. The summary recommendations are narrative judgments without documented confidence levels. An IC member asking "what would change this from medium to high?" cannot get a parameter-level answer; the question is methodologically incoherent within the qualitative methodology.

Documented assumptions are typically incomplete. The diligence engagement makes assumptions about the target's threat environment, the regulatory landscape, the post-close integration trajectory, and the acquirer's residual risk tolerance. These assumptions are typically not documented in the engagement deliverable; they are baked into the engagement team's expert judgment.

Reconcilable outputs are not produced. The qualitative methodology does not produce outputs that can be reconciled against post-close realized cyber outcomes, because the outputs are not denominated in units that map to realized outcomes. A "high severity" rating cannot be reconciled against a $14M post-close breach cost, because the rating is not a prediction of the cost. The methodology does not learn.

Quantitative cyber diligence produces defensible outputs across all four properties. The methodology is documented in the methodology specification[10] at the level of specific formulas, calibration parameters, and procedural steps. Calibration parameters are specified explicitly with documented sources. Assumptions are documented in the engagement deliverable's methodology disclosure. Outputs are denominated in dollars and produce post-close validation engagements that reconcile predicted figures against realized outcomes, feeding the engagement-evidence learning loop.

2.4 The temporal failure

The fourth failure mode is that qualitative cyber diligence does not operate at deal speed.

A typical M&A transaction in current practice runs on calendar pressures the cyber diligence engagement must accommodate. Letter of intent to close runs 60 to 120 days for mid-market deals, 90 to 180 days for large deals, longer for megadeals subject to regulatory review. Within that window, cyber diligence is one of approximately fifteen functional workstreams (legal, financial, commercial, operational, technical, IT, HR, environmental, regulatory, etc.), competing for target-team availability and deal-team attention. The cyber engagement's window is typically 30 to 60 days, and it must produce a deliverable in time for IC review at minimum two weeks before close.

The qualitative methodology accommodates this timeline by limiting depth. The 30-to-60-day window is consumed primarily by questionnaire issuance and review (10 to 15 days), target-CISO interviews (3 to 5 days), external scan ordering and analysis (5 to 10 days), and report drafting

[10]Yolonda Smith, *QCD Methodology Specification*, companion methodology paper, 2026.

(10 to 20 days). The methodology does not produce parameter-level analysis because the timeline does not support parameter-level analysis. The output is a memorandum at the level the timeline permits.

The temporal pressure interacts badly with the IC's needs. The IC asks the deal team to produce a defensible cyber number under the same temporal pressure that produced the qualitative memo. The deal team cannot produce the figure because the methodology underlying the engagement does not support figure production. The deal team's options at IC are to commission additional analytical work (creating timeline pressure that may delay close), to accept the qualitative output without challenge (the modal outcome), or to escalate the cyber finding to a deal-specific structural overhaul (rare, typically only when the qualitative finding is severe enough to threaten deal-team consensus).

The framework operates at deal speed. Tier 1 engagements (chapter 12) deliver in 10 to 14 days. Tier 2 engagements deliver in 25 to 35 days. Both produce dollar-denominated, percentile-resolved CCOD figures and dual-band IC anchors with associated structural recommendations. The temporal advantage comes not from compressed analytical depth but from a methodology designed for deal-speed application: pre-built reference-class dictionaries, calibrated parameters, automated tooling for external scan and vendor analysis, and a structured deliverable format that compresses memorandum drafting time.

2.5 The structural-recommendation failure

The fifth failure mode is that qualitative cyber diligence produces structural recommendations not aligned with the realized loss distribution.

When cyber findings produce a "high severity" rating in a qualitative memo, the typical recommended structural response is "address through standard rep-and-warranty provisions with a cyber-specific basket and sub-limit." The recommendation is structurally identical across deals, regardless of whether the underlying loss distribution is concentrated (most realizations near the central case) or tailed (most realizations modest with rare large outliers). The same recommendation is produced for a target whose dominant cyber exposure is operational disruption (where typical realization fraction is 1.8% of theoretical maximum) and a target whose dominant exposure is long-tail litigation (where typical realization fraction is 11.8% of theoretical maximum).

The dual-band reporting structure introduced in chapter 10 makes this distinction operational. Targets with no tail-driver flag firing produce a contained-band IC anchor; targets with one or more tail-driver flags firing produce an escalated-band IC anchor. The two bands map to materially different recommended deal architectures: contained-band recommendations cluster around price adjustment, modest escrow, and reps-and-warranties cyber sub-limits; escalated-band recommendations cluster around larger escrows, indemnity carve-outs, and pre-close covenants requiring specific remediation work. The qualitative methodology's single-recommendation pattern collapses both into the same structural posture, over-prescribing for contained-band targets and under-prescribing for escalated-band targets.

The Schwab/TD Ameritrade transaction (chapter 10) illustrates the contained-band over-prescription pattern. The retrospective application of QCD produces a contained-band anchor sized for the conservative-realization scenario; realized cost came in below even the contained band. The qualitative methodology applied at the time produced structural recommendations consistent with a more conservative posture than the realized outcome warranted. The Marriott/Starwood transaction (chapter 5) illustrates the escalated-band under-prescription pattern. The qualitative methodology in 2016 produced standard reps-and-warranties recommendations; the realized cost — long-tail litigation, regulatory investigation, multi-year settlement tail — required materially larger structural protections than the rep-and-warranties framework provided. Marriott absorbed the difference.

The framework's dual-band structure addresses this failure by selecting structural recommendations that match the realized loss distribution. The chapter 10 treatment walks the four tail-driver flags against the sixteen-case retrospective, demonstrating how flag firings produce escalated-band recommendations that match the realized outcomes of the documented case set.

2.6 What quantitative cyber diligence produces

Quantitative cyber diligence — the framework documented in this book — produces outputs that address each of the five failure modes structurally:

Dollar-denominated outputs at the per-pillar and aggregate level, expressed as percentage of enterprise value, with explicit confidence intervals and percentile resolution.

Comparable outputs across deals, with calibration version stamping that supports cross-deal comparison and portfolio-level rollup.

Defensible outputs with explicit methodology, calibrated parameters, documented assumptions, and post-close validation reconciliation. The framework's calibration evolves through engagement-evidence learning rather than remaining frozen.

Deal-speed outputs through Tier 1 and Tier 2 engagement modes that deliver in 10–14 days and 25–35 days respectively, with calibrated reference-class dictionaries and pre-built tooling that compresses analytical time without compressing analytical depth.

Structural recommendations matched to realized loss distributions through dual-band reporting with tail-driver flag taxonomy. Contained-band anchors produce recommendations consistent with bulk-realization scenarios; escalated-band anchors produce recommendations consistent with tail-realization scenarios.

The methodology is not a replacement for the qualitative diligence team's expertise; the engagement teams running QCD bring the same domain knowledge, professional judgment, and target-cooperation skills that qualitative engagements require. What differs is the methodology's analytical apparatus and reporting structure. The professional judgment enters the framework as

parameter elicitation, scenario evaluation, and tail-driver flag assessment — rendered into commensurable units that produce IC-defensible outputs.

2.7 Forward reference

The next part of the book documents the framework's five pillars in detail. Each pillar chapter cold-opens with a retrospective case demonstrating the pillar's signature failure mode under qualitative diligence: chapter 3 (Pillar 1) opens on UnitedHealth/Change Healthcare and the absent-MFA Citrix portal that external assessment would have surfaced; chapter 4 (Pillar 2) opens on PowerSchool and the SaaS-essential-infrastructure concentration pattern; chapter 5 (Pillar 3) opens on Marriott/Starwood and the long-tail litigation exposure; chapter 6 (Pillar 4) opens on Yahoo/Verizon and the maturity gap that drove a $350M repricing; chapter 7 (Pillar 5a) opens on MGM/Cosmopolitan and the inverse-integration finding; chapter 8 (Pillar 5b) opens on Brookfield/CDK and the operating-disruption sub-component.

Each chapter establishes the formal mathematics, the calibrated parameters, the engagement workflow, and the named pitfalls. Each chapter's worked Apex micro-example demonstrates the pillar's calibrated computation against the running synthetic target. The five pillars converge in chapter 9's CCOD aggregation, chapter 10's dual-band reporting and IC conversation, and chapter 11's translation from CCOD anchor to valuation adjustment and recommended deal architecture.

The framework's value proposition is not analytical novelty — most of the underlying machinery (FAIR-style decomposition, Monte Carlo aggregation, calibrated probability assessment) exists in the broader cyber-risk-quantification literature. The framework's value is methodological coherence at the deal-time application: a pre-calibrated, deal-speed, IC-defensible, post-validation-reconciled methodology specifically targeted at the M&A cyber diligence problem. The qualitative methodology fails at this problem because the qualitative methodology is general-purpose; the framework is purpose-built for the deal-time application and produces outputs the IC can use.