QRTF (Manuscript v1.3c): A Quality-to-Risk Transfer Format and Monotone Logit Model for QA/QMS Risk Monitoring in Regulated MedTech Contexts

Zenodo Preprint · Manuscript v1.3c · Prepared 2025-12-30

Author: Mehdi Zadehnour (independent work; no employer affiliation or endorsement implied)

Disclaimer: This manuscript describes a methodological framework for quality and regulatory risk monitoring. It is not medical advice, legal advice, or a substitute for an organization's validated QMS processes. Any deployment in a regulated setting requires documented validation, governance, and auditability appropriate to the intended use.

Abstract

Risk management in medical devices requires continuous, evidence-based decisions that integrate design-time risk analysis with post-market and quality system signals. Many organizations rely on static risk matrices or FMEA-derived rankings that do not naturally absorb the evolving health of the Quality Management System (QMS). QRTF™ (Quantitative Risk Trace Fusion) is a tool-agnostic quality-to-risk transfer framework that converts multi-source QMS maturity and operational signals into updated effective failure probabilities (P_eff), risk budgets (ΔB), and a bounded Risk Index, with explicit monotonicity constraints (higher maturity cannot raise risk). QRTF integrates (i) a monotone logistic transfer on the logit scale, (ii) divergence and ordered-ground optimal-transport features between commissioning baselines and operational windows, (iii) quality gating and dual-threshold hysteresis with dwell to prevent alert flapping, and (iv) a budgeted portfolio optimizer that ranks improvement actions by expected risk reduction per cost. We provide the v1.0 data contract (JSON schema and CSV templates), describe compliance-oriented traceability features (release IDs, hashes, signatures, audit events), and demonstrate a 10-day pilot dataset illustrating end-to-end outputs and action selection.

Keywords: Quality risk management; ISO 14971; ISO 13485; QMS maturity; CAPA; post-market surveillance; monotone models; optimal transport; hysteresis; portfolio optimization; auditability; digital traceability.

Zenodo Concept DOI (all versions): https://doi.org/10.5281/zenodo.18097045

CONFIG SNAPSHOT (non-normative drafting metadata)

This section documents drafting constraints and verification status for transparency; it is not a regulatory interpretation and not legal/medical advice.

Manuscript ID	QRTF™ v1.0 Zenodo preprint (final Word export)
Author	Mehdi Zadehnour (independent work; no affiliation implied)
Language	English (EN)
Target venue	Zenodo preprint (journal-style formatting)
Rigor / scope	High-rigor technical manuscript; no fabrication; pilot is illustrative
Web verification	On (targeted verification of load-bearing regulatory and foundational references)
Citations style	IEEE-like numbered references
Figures	Embedded figures (architecture schematic; pilot RiskIndex plot)
Reproducibility pack	Word manuscript + appendix describing included schema/templates and integrity checks
Date	2025-12-30

ASSUMPTION REGISTER (explicit, minimal)

Assumptions are tagged in the text as [ASSUMPTION:A#] where they materially affect interpretation.

ID	Assumption	Used in	Impact	Verification path
A1	Baseline failure probabilities P0(fm) and severities S(fm) come from design-time risk analysis and are suitable as priors over the monitoring window.	Sec. 2.1, 3.3, 7.1	High	Verify against the current risk file/FMEA and change-control history.
A2	Higher QMS maturity (higher RRL_j / better KPI) should not increase computed risk; model	Sec. 3.0–3.1	High	Unit tests + monotonic constraint checks during calibration.

	coefficients enforce monotonicity.			
A3	Aggregation across failure modes (sum/max/weighted sum) reflects the organization's risk policy and independence assumptions.	Sec. 3.3	Med	Document aggregation policy in QMS governance; sensitivity analysis.
A4	Divergence and ordered-ground transport features use meaningful KPI binning and ordering consistent with domain semantics.	Sec. 3.2	Med	Review binning/order with SMEs; compare alternative discretizations.
A5	Pilot10 dataset is illustrative (not performance-validating); thresholds (theta_on/off) and RiskIndex scaling are policy parameters.	Sec. 4, 7	High	Validate on representative field datasets; justify thresholds via risk criteria.
A6	Auditability claims (hashes/signatures) assume secure key management and controlled processes for generating manifests and signatures.	Sec. 5, App. B	Med	Verify key management, access control, and audit trail procedures under QMS/CSV.

OUTLINE (M0) — modules coverage

- M1–M4 (Title/Abstract/Keywords/Introduction): FULL
- M5 (Background & Related Work): SUMMARY-LEVEL (selected foundational references only)
- M6–M8 (Notation/Theory/Methods): FULL
- M9–M10 (Pilot/Results): FULL (illustrative Pilot10 dataset)
- M11 (Validation & Consistency checks): FULL (sanity checks + caveats)
- M12–M13 (Discussion/Conclusion): FULL
- M15 (References): FULL (verified where possible; no invented DOIs)
- M16 (Appendices): FULL (schema summary + reproducibility/integrity)

1. Background and Motivation

Regulated medical device development and maintenance depend on consistent risk-based decision-making across the entire lifecycle. ISO 14971 formalizes a process for identifying hazards, estimating risk, implementing risk controls, and monitoring residual risk during production and post-production. In practice, many organizations keep design-time risk files (e.g., FMEA/FMECA) separate from operational quality signals (CAPA performance, audit findings, supplier issues, training coverage, complaints, and post-market signals). This separation makes it difficult to answer a simple question with evidence: 'Given today's QMS health, what is the effective likelihood of critical failure modes tomorrow?' Key normative anchors include ISO 14971 for risk management and ISO 13485 for QMS requirements [1], [2]. Recent regulatory convergence efforts (e.g., the FDA QMSR final rule) further motivate structured, auditable risk evidence flows [3], [4].

Related maturity-model efforts exist for risk management in medical device contexts (e.g., capability/maturity models for RM practices), but they typically assess organizational maturity rather than provide an auditable, quantitative transfer function from QMS signals to updated failure probabilities and decision thresholds. QRTF is positioned as that transfer layer: it can consume maturity assessments (RRL) and operational KPIs and convert them into reproducible, monotone risk updates and action prioritization.

QRTF addresses this gap by defining a standardized quality-to-risk transfer layer. It does not replace a QMS, an FMEA, or regulatory judgment. Instead, it provides an auditable, monotone, and tunable mapping from (a) QMS maturity and operational indicators to (b) updated failure probabilities, risk budgets, and prioritized improvement actions.

Design constraints were guided by three practical requirements observed in QA/QMS work:

- Monotonicity and explainability: improving a quality dimension must never increase computed risk; drivers must be attributable to inputs.
- Stability under noise: alerts should not flap due to transient KPI spikes; gating must include hysteresis and dwell.

- Auditability: the model's inputs, outputs, versions, and configuration must be traceable with hashes/signatures and event logs.

2. QRTF Overview

At a high level, QRTF is composed of four layers: (1) standardized inputs (dimensions and failure modes), (2) feature construction (divergence and ordered-ground transport features), (3) monotone logistic fusion producing P_eff and ΔB, and (4) decision logic that gates risk states and selects improvement portfolios under a budget constraint.

Figure 1 summarizes the architecture.

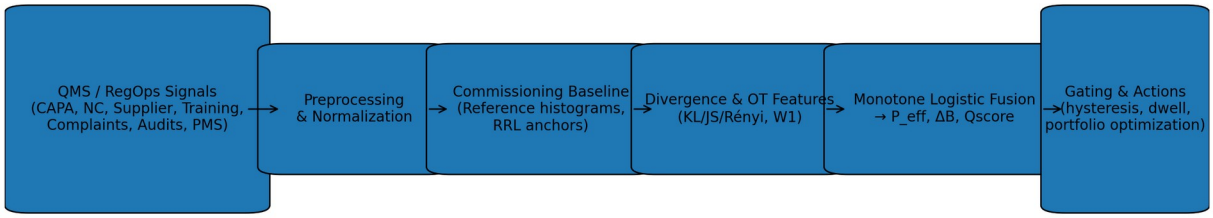

Figure 1. Conceptual architecture of QRTF: quality signals are normalized, compared to commissioning baselines using divergence/OT features, fused via monotone logistic transfer, then gated with hysteresis and connected to action selection.

2.1 Inputs and Outputs

QRTF operates on two primary entities: (a) dimensions describing QMS maturity and operational health, and (b) failure modes describing baseline risk per failure mode.

Signal / Field	Meaning	Typical Range	Notes
P0(fm)	Baseline failure probability per unit time (per failure mode)	0–1	Prior probability before QMS effect; derived from design-time risk analysis or empirical rates.
S(fm)	Severity score per failure mode	1–5 (example)	Scale can be adapted; used to scale risk budget ΔB.
RRL_j	Regulatory Reference Level for dimension j (maturity)	0–5	Discrete maturity anchor; higher is better.
KPI_j(t)	Operational KPI readings for dimension j	varies	Counts/ratios; can be windowed over time.

ψ_j(RRL_j, KPI_j)	Mapped continuous feature for dimension j	≥0	Monotone mapping from RRL (and optionally KPIs) to a continuous risk effect feature.
P_eff(fm)	Effective failure probability after quality-to-risk transfer	0–1	Computed on logit scale; bounded by construction.
ΔB(fm)	Risk budget change for failure mode	real	ΔB = (P_eff − P0)·S (one canonical form).
Risk Index	Aggregated risk score (bounded scale)	0–100+ (example)	Policy-chosen scaling; can exceed threshold to show severity beyond limits.
Qscore / k_qms	Quality/data confidence score for gating	0–1	Used to down-weight or gate processing when data quality is low.
State	Discrete risk state under hysteresis/dwell	GREEN/YELLOW/RED	Dual-threshold hysteresis with dwell prevents flapping.

2.2 Notation and Nomenclature

Table 2 summarizes core symbols and their meaning; symbols are defined at first use in the main text.

Symbol	Meaning	Unit	Where used
fm	Failure mode index/identifier	—	Sec. 2.1, 3–7
j,k	QMS dimension indices	—	Sec. 2.1, 3
RRL_j	Regulatory Reference Level / maturity anchor for dimension j (0–5)	level	Sec. 2.1, 3.1, 7
KPI_j	Operational KPI window/statistic for dimension j	domain-specific	Sec. 3.1–3.2
P0(fm)	Baseline (commissioning) failure probability for fm	probability per time unit	Sec. 2.1, 3
P_eff(fm)	Effective failure probability under current evidence	probability per time unit	Sec. 3
S(fm)	Severity score for fm	dimensionless (1–5 in Pilot10)	Sec. 2.1, 3.3, 7
α0	Global logit offset	—	Eq. (1)
β_j	Monotone coefficient	—	Eq. (1)

	for dimension j (β_j ≤ 0)		
ψ_j(·)	Monotone mapping from maturity/KPI to feature for dimension j	—	Sec. 3.1
γ_{jk}	Limited interaction term between dimensions j and k	—	Eq. (1)
ΔB(fm)	Risk budget delta for fm: (P_eff − P0)·S	risk points (scaled)	Sec. 3.3
RiskIndex	Aggregated risk index across failure modes/drivers	risk points	Sec. 4, 7
θ_on, θ_off	Hysteresis thresholds for escalation/clear	risk points	Sec. 4, Fig. 2

3. Monotone Logistic Fusion on the Logit Scale

For each failure mode fm, QRTF maps a commissioning baseline failure probability P0(fm) to an operational effective failure probability P_eff(fm) using an additive model on the logit scale. To keep the mapping interpretable and auditable, the quality effect is expressed relative to a commissioning reference state, so that when the QMS is at the reference maturity, the model returns P_eff(fm) = P0(fm) (i.e., no inflation/deflation).

$$\begin{aligned}\operatorname{logit}(P_{\text{eff}}(fm,t)) = \operatorname{logit}(P_0(fm)) + \alpha_0 \\ + \Sigma_j \beta_j \cdot (\psi_j(t) - \psi_j^{\text{ref}}) \\ + \Sigma_{j<k} \gamma_{jk} \cdot (\psi_j(t) - \psi_j^{\text{ref}}) \cdot (\psi_k(t) - \psi_k^{\text{ref}})\end{aligned}$$

where α0 is a global offset (often set to 0 when ψ^ref is defined), β_j are per-dimension coefficients, ψ_j(t) ∈ [0,1] is a monotone (non-decreasing) maturity/health feature for dimension j at time t, and ψ_j^ref is the commissioning (reference) maturity feature for that dimension. The v1.0 spec enforces the constraint β_j ≤ 0 when ψ_j is defined as a health/maturity score (higher ψ means better maturity), so that increasing maturity cannot increase computed risk. If an implementation instead defines ψ as a deficit (higher ψ means worse maturity), then the sign constraint must be inverted (β_j ≥ 0). This manuscript uses ψ as a health score by default.

3.1 Monotone Mappings for Maturity Levels

A minimal mapping normalizes the discrete maturity anchor RRL_j ∈ {0,…,RRL_max} to a unit interval maturity score ψ_j ∈ [0,1] (higher is better):

$$\psi_j(RRL_j) = RRL_j / RRL_{\max}$$

Alternatively, a saturating exponential mapping can represent diminishing returns at high maturity (still bounded to [0,1]):

$$\psi_j(RRL_j) = (\exp(\lambda_j \cdot RRL_j) - 1) / (\exp(\lambda_j \cdot RRL_{max}) - 1), \quad \text{with } \lambda_j > 0$$

Operational KPIs can be introduced either as additional monotone features or by modulating ψ_j via drift features derived from KPI distributions (Section 3.2). In all cases, QRTF requires monotone behavior: improving maturity or KPI health must not increase risk.

3.2 Divergence and Ordered-Ground Optimal Transport Features

QRTF supports comparing commissioning baselines to operational windows using statistical divergence measures and ordered-ground transport distances. In practice, a dimension's KPI readings are summarized into a histogram (or other distributional summary) over a time window; the operational distribution is then compared to a commissioning reference distribution. Divergences (e.g., KL, Jensen–Shannon, Rényi) measure distributional shift, while ordered-ground optimal transport metrics (e.g., Wasserstein-1 / Earth Mover's Distance, W1) quantify how much 'mass' must move along an ordered bin axis to match the baseline. Foundational definitions include Kullback–Leibler divergence [10], Jensen–Shannon divergence [11], Rényi divergence [12], and optimal-transport/Wasserstein formulations [13], [14].

3.3 Risk Budgets and Aggregation

A simple risk-budget form used in the accompanying materials is:

$$\Delta B(fm) = (P_{eff}(fm) - P0(fm)) \cdot S(fm)$$

Total risk budget can be aggregated across failure modes (e.g., sum, max, or weighted sum) depending on the independence assumptions and governance policy. QRTF reports per-failure-mode results (P0, P_eff, severity, ΔB) and can also compute overall b0_total, delta_B_total, and b_total.

Canonical aggregation (v1.0): define a severity-weighted total risk budget at time t as $B_{total}(t) = \Sigma_{fm} w_{fm} \cdot S(fm) \cdot P_{eff}(fm,t)$, and the commissioning baseline as $B0_{total} = \Sigma_{fm} w_{fm} \cdot S(fm) \cdot P0(fm)$, with optional weights w_{fm} (default 1).

Define an allowed budget increase $B_{budget} = \kappa \cdot B0_{total}$ with $\kappa > 0$ as a governance/policy parameter. Then define the (non-negative) Risk Index as:

$$RiskIndex(t) = 100 \cdot \max(0, B_{total}(t) - B0_{total}) / B_{budget}.$$

This definition yields RiskIndex = 0 at baseline (no risk inflation vs commissioning), RiskIndex = 100 when the risk budget exceeds baseline by $\kappa \cdot B0_{total}$, and allows straightforward thresholding (e.g., theta_on/theta_off in Section 4). Implementations may optionally cap RiskIndex (e.g., at 200) to avoid unbounded scales during extreme regimes; if capped, the cap must be reported and treated as a loss of resolution rather than as a stable measurement.

4. Control Layer: Quality Gating and Hysteresis

To avoid alert flapping under noisy inputs, QRTF uses dual-threshold hysteresis with dwell time. Two thresholds are defined: theta_on (trigger) and theta_off (clear), with theta_on > theta_off. A dwell time requires the threshold condition to persist for a minimum duration (e.g., 2 days) before changing state.

- Escalation: if Risk Index ≥ theta_on continuously for dwell time → enter or maintain a higher state.
- De-escalation: once in a higher state, Risk Index must fall below theta_off continuously for dwell time to clear.

The pilot materials use theta_on = 65, theta_off = 50, and dwell = 2 days, producing a stable RED state when Risk Index is persistently high.

4.1 Data Quality Score (Qscore / k_qms)

QRTF reports a quality/data confidence indicator (Qscore, sometimes denoted k_qms) in the 0–1 range. When data quality is low, QRTF can reduce sensitivity or trigger conservative fallback behavior. This is intended as an engineering control against spurious alerts due to missingness, delayed updates, or known measurement issues.

5. Data Contract and Traceability

QRTF v1.0 includes a versioned JSON schema defining RiskRequest, RiskResponse, and AuditEvent messages, plus CSV templates for dimensions, failure modes, improvement options, and test vectors. Core request fields include schema_version, release_id, org_id, product_id, dimensions_payload, failure_modes, and options. Core response fields include k_qms, per-dimension partial contributions, per-failure-mode results, and overall risk budgets.

To support auditability in regulated environments, the schema includes release_id, timestamps, and a signature field. The accompanying spec emphasizes hashing and signing of inputs/outputs and maintaining event logs for traceability (e.g., for Part 11-like expectations on electronic records).

6. Budgeted Portfolio Optimization for Risk Reduction

Beyond risk estimation, QRTF supports selecting a subset of improvement actions under a budget constraint to maximize expected risk reduction. Each option i is described by (dimension(s), expected improvement ΔRRL or Δψ, cost, and mapping to affected failure modes). A canonical objective is:

$$\text{maximize} \ \ \Sigma_i \, \Delta B_i \ - \ \lambda \cdot \text{Var}(\Delta B)$$
$$\text{subject to} \ \ \Sigma_i \, \text{cost}_i \leq \text{Budget} \quad \text{(and prerequisites/constraints)}$$

In the simplest case, this reduces to a 0-1 knapsack problem; in richer cases, it becomes an integer program. The output is a ranked list of actions with estimated risk-index reduction and ROI (risk points reduced per currency unit).

7. Pilot Demonstration: 10-Day Scenario (Supplementary Pilot10 Package)

Pilot10 is an end-to-end pipeline demonstration (including an illustrative capped regime). The semi-synthetic evaluation section provides variable-regime dynamics used to quantify state transitions, flapping behavior, and sensitivity.

The supplementary Pilot10 package provides a compact end-to-end example: four QMS dimensions (CAPA, SUPPLIER, TRAINING, COMPLAINTS), three failure modes with baseline probabilities and severities, daily operational signals, and computed outputs including Risk Index, state, driver contributions, and recommended actions.

7.1 Pilot Inputs

Initial maturity anchors (RRL 0–5) at 2025-11-01:

Dimension	RRL (0–5)	Comment
CAPA	2	Corrective and Preventive Action
SUPPLIER	1	Supplier quality
TRAINING	3	Training/competency
COMPLAINTS	2	Complaint handling

Baseline failure modes used in the pilot:

Failure mode ID	P0 (per day)	Severity (1–5)	Notes
FM1_overinfusion	0.00010	5	Illustrative placeholder
FM2_alarm_failure	0.00005	4	Illustrative placeholder
FM3_occlusion_FN	0.00002	3	Illustrative placeholder

7.2 Pilot Outputs and Interpretation

The pilot's one-page report summarizes: Max RiskIndex = 200.0, Days RED = 10, Days YELLOW = 0, Days GREEN = 0, Top Driver = CAPA, theta_on = 65, theta_off = 50, and dwell = 2 days. The Risk Index remains at 200 across all 10 days (Figure 2). Driver contributions and Qscore vary by day, and the action selector recommends a CAPA backlog sprint as the top intervention.

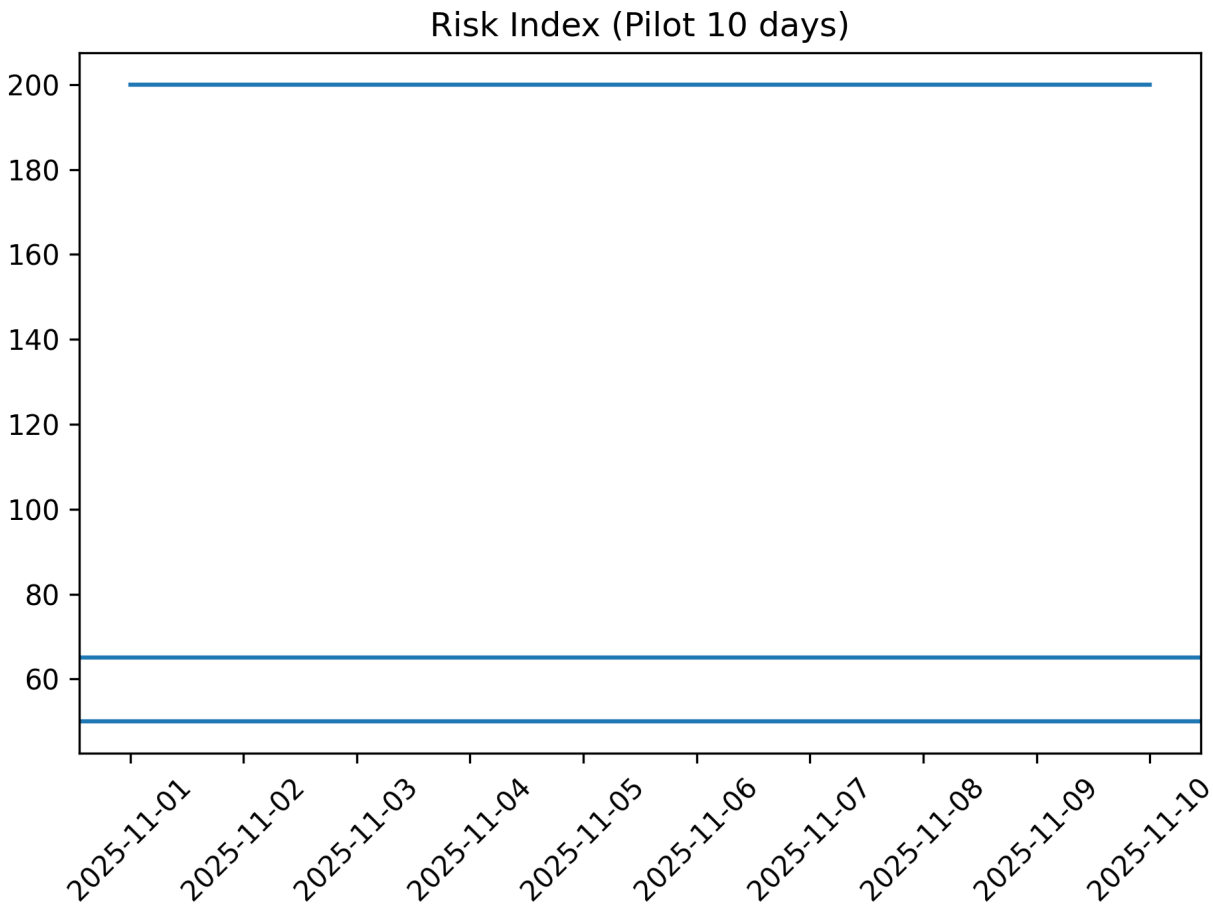

Figure 2. Pilot10 Risk Index over 10 days (constant 200 in this illustrative dataset) with theta_on = 65 and theta_off = 50 shown as horizontal lines.

Selected Pilot10 daily outputs (first 5 days shown):

date	RiskIndex	State	Q	top_driver	multiplier
2025-11-01	200.0	RED	0.86	CAPA	2.16
2025-11-02	200.0	RED	0.84	TRAINING	2.08
2025-11-03	200.0	RED	0.96	CAPA	2.34
2025-11-04	200.0	RED	0.72	CAPA	2.45
2025-11-05	200.0	RED	0.89	SUPPLIER	3.33

Note: The Risk Index is constant at 200 in this pilot dataset, indicating a capped/illustrative high-risk regime. In real deployments, Risk Index should be validated against field outcomes and expected to vary meaningfully with signal changes.

7.3 Recommended Actions (Pilot10)

The top recommended actions and their estimated reductions and costs are:

7.4 Semi-synthetic evaluation (reproducible, non-clinical) [SYNTHETIC]

Purpose: provide a reproducible sanity-check that (i) the RiskIndex scale behaves as intended, (ii) hysteresis+dwell reduces alert flapping, and (iii) the monotonicity constraint (improving maturity cannot increase risk) holds in randomized perturbation tests. This evaluation uses a semi-synthetic generator (random-walk RRL with occasional regressions) and does not claim clinical or field-calibrated performance.

Setup (fixed for reproducibility): 60 days; dimensions = {CAPA, SUPPLIER, TRAINING, COMPLAINTS}; RRL_max = 5; linear ψ_j = RRL_j/RRL_max; commissioning references ψ^ref correspond to RRL_ref = {3, 3, 3.5, 3}. Model coefficients enforce monotonicity (β_j ≤ 0). RiskIndex uses κ = 0.5 and is capped at 200 for reporting. Hysteresis parameters: theta_on = 65, theta_off = 50, dwell = 2 days.

Summary (linear ψ): RiskIndex min=29.7, max=103.1, mean=55.6; Days GREEN=38, YELLOW=2, RED=20; State transitions=6 (flap rate=0.102/day).

Monotonicity check: 0 violations in 800 randomized perturbation checks (increase one RRL by +0.5, recompute RiskIndex; RiskIndex never increases).

Figure 3. [PLACEHOLDER] Semi-synthetic RiskIndex(t) over 60 days with hysteresis thresholds (theta_on/theta_off) and resulting state.

Table 5. Semi-synthetic daily excerpt (first 10 days):

date	RiskIndex	State	top_driver
2025-11-01	85.7	GREEN	SUPPLIER
2025-11-02	103.1	RED	SUPPLIER
2025-11-03	99.1	RED	SUPPLIER
2025-11-04	94.3	RED	SUPPLIER
2025-11-05	95.1	RED	SUPPLIER
2025-11-06	91.3	RED	COMPLAINTS
2025-11-07	93.9	RED	COMPLAINTS
2025-11-08	90.2	RED	COMPLAINTS
2025-11-09	78.2	RED	COMPLAINTS
2025-11-10	71.1	RED	COMPLAINTS

7.5 Sensitivity and robustness (semi-synthetic) [SYNTHETIC]

To quantify sensitivity to modeling and policy choices, we rerun the same semi-synthetic RRL series under alternative ψ mappings while keeping β monotone (β_j ≤ 0), κ = 0.5, and the same hysteresis thresholds (theta_on=65, theta_off=50, dwell=2). The goal is not to claim real-world performance, but to demonstrate how design choices change scale and alert dynamics.

Table 6. Sensitivity summary (same RRL series; different ψ mappings):

ψ mapping	RiskIndex min	RiskIndex max	Mean	Days RED	Transitions

Linear ψ = RRL/RRL_max	29.7	103.1	55.6	20	6
Exponential saturating ψ (λ=0.7)	0.0	48.3	17.2	0	0

Interpretation: the ψ mapping changes the effective scale of (ψ − ψ^ref) and therefore changes the RiskIndex distribution and alert frequency. For deployment, ψ mappings, κ, and thresholds must be calibrated against outcomes and governance risk appetite; otherwise the same underlying signals can appear either "quiet" or "alarming" under different (yet monotone) mappings.

Option ID	Dimension	Est. RiskIndex reduction (points)	Cost (CHF)
OPT_CAPA_BACKLOG_SPRINT	CAPA	11.6	80,000
OPT_TRAINING_REFRESH	TRAINING	6.2	60,000
OPT_COMPLAINT_TRIAGE	COMPLAINTS	4.5	50,000
OPT_SUPPLIER_TARGET_AUDIT	SUPPLIER	9.3	120,000

8. Discussion

Figure (placeholder): Governance stack alignment diagram (RRL-QMS/RRL-P → QRTF → RDB → controlled change execution) with traceability via release_id and signed audit events.

QRTF → RDB triggers: RDB defines outcome-focused, audit-ready triggers for dynamic change management under a PCCP. In an integrated workflow, QRTF can supply a risk-state/trajectory (e.g., RiskIndex(t), driver contributions, and quality confidence) as an input signal for RDB's budgeting and trigger logic, while RDB governs when change execution is permitted/required. The separation is deliberate: QRTF provides risk quantification; RDB provides regulator-ready trigger governance.

RRL-P → optional domain adaptation: RRL-P is the pharma analogue and can be used to define maturity anchors for GMP/ICH Q10 contexts. QRTF itself remains tool-agnostic; the main adaptation is the definition of dimensions and the mapping ψ_j(·) appropriate to the domain. No pharma claims are made in this manuscript.

RRL-QMS → QRTF inputs (RRL_j): RRL-QMS provides a domain-specific maturity ladder mapped to regulatory expectations. When available, an organization's RRL-QMS level(s) can populate QRTF's per-dimension maturity anchors RRL_j directly. This makes QRTF's monotone constraint (β_j ≤ 0) semantically consistent: higher readiness cannot increase computed risk.

QRTF is intentionally scoped as a quality-to-risk transfer layer. In the broader governance stack, it can be aligned with three adjacent public frameworks published by the same author: RRL-QMS

(a staged readiness ladder for medical-device QMS), RRL-P (a pharma-specific readiness ladder), and RDB (Regulatory Drift Budget, a drift-trigger framework for PCCP-style change governance). This section aligns terminology and clarifies boundaries; it does not disclose private implementation details.

Positioning and Alignment with RRL-QMS, RRL-P, and RDB (Public Interfaces)

QRTF's value proposition is not a single formula; it is the combination of monotone risk fusion, distribution-shift features, stability controls, and auditable data contracts. In regulated environments, interpretability and governance are often as critical as predictive accuracy. Monotonicity constraints help reduce counter-intuitive behavior and simplify explanation in audits.

Metric gaming and aggregation artifacts are real risks: optimize against hard-to-game outcomes (e.g., recalls/serious events) and interpret aggregates cautiously to avoid Simpson-type effects. These points are governance considerations and do not substitute for calibrated validation.

Limitations: the Core Package v1.0 materials are sufficient to reproduce the data exchange and pilot outputs, but the approach must be validated on representative datasets before being used for operational decisions. The pilot dataset is illustrative and does not demonstrate calibrated performance.

9. Conclusion

QRTF™ (Core Package v1.0) provides a structured and auditable bridge between QMS maturity and risk monitoring. By defining a standardized data contract, monotone logistic fusion on the logit scale, divergence/transport features for drift detection, and governance-ready gating and portfolio optimization, QRTF supports risk-based quality management in a way that is compatible with regulated expectations. Future work should focus on calibrated validation, causal and Bayesian extensions, and deployment governance (traceability, access control, and change management).

Appendix A. v1.0 Interface Schema Summary

The accompanying qrtf_schema_v1_0.json defines the following main message types:

- RiskRequest: schema_version, release_id, org_id, product_id, dimensions_payload[], failure_modes[], options, timestamp, signature.
- RiskResponse: schema_version, release_id, k_qms, dimensions[], results[], b0_total, delta_B_total, b_total, timestamp, signature.
- AuditEvent: schema_version, release_id, event_type, event_payload, timestamp, signature.

DimensionPayload items include dimension_id, rrl (0–5), score, kpis[], timestamp. FMResult includes fm_id, p0, p_eff, severity, delta_B, and a confidence_interval.

Appendix B. Reproducibility and Integrity

The QRTF v1.0 package includes a manifest.json and SHA-256 hashes for bundle integrity. For regulated usage, storing each release with its hash, inputs, configuration, and outputs enables end-to-end reproducibility and supports audit trails.

Reproducibility & Integrity (Summary)

Integrity note: the submission bundle includes SHA-256 checksums so any reader can verify file integrity and detect modifications.

To reproduce the provided pilot outputs: (i) use the included schema/templates to assemble a RiskRequest with RRL_j, P0, S, thresholds (θ_on/θ_off/dwell), (ii) run the monotone logit transfer to obtain P_eff and ΔB, and (iii) apply hysteresis gating for the discrete state.

This manuscript is supported by a versioned data contract and reproducibility artifacts included in the Zenodo deposit. The Pilot10 package is illustrative (pipeline demonstration) while the public-data benchmark pack specifies an outcome-aligned validation path using authoritative public sources.

References

1. ISO 14971:2019. Medical devices — Application of risk management to medical devices. International Organization for Standardization (ISO), 2019.
2. ISO 13485:2016. Medical devices — Quality management systems — Requirements for regulatory purposes. International Organization for Standardization (ISO), 2016.
3. J. Burton, F. Mc Caffery, and I. Richardson, “A Risk Management Capability Model for use in Medical Device Companies,” Proceedings of WoSQ’06 (Workshop on Software Quality), Shanghai, China, May 21, 2006. Available (PDF): https://www.researchgate.net/profile/Ita-Richardson/publication/228337572_A_risk_management_capability_model_for_use_in_medical_device_companies/links/53fef2040cf21edafd154524/A-risk_management_capability_model_for_use_in_medical_device_companies.pdf
4. Office of Rail and Road (UK), “Risk Management Maturity Model (RM3) 2019,” ORR document, 2019. Available: https://www.orr.gov.uk/sites/default/files/2020-09/risk-management-maturity-model-rm3-2019.pdf
5. U.S. Food and Drug Administration (FDA), “Medical Devices; Quality System Regulation Amendments,” Final Rule, Federal Register, 89 FR 7496 (Feb. 2, 2024). Available: https://www.federalregister.gov/documents/2024/02/02/2024-01709/medical-devices-quality-system-regulation-amendments (effective Feb. 2, 2026).
6. U.S. Food and Drug Administration (FDA), “Medical Devices; Quality System Regulation Amendments; Correction,” Federal Register (Oct. 15, 2024). Available:

https://www.federalregister.gov/documents/2024/10/15/2024-23701/medical-devices-quality-system-regulation-amendments-correction (effective Feb. 2, 2026).
7. U.S. Food and Drug Administration (FDA), “Medical Devices; Quality Management System Regulation Technical Amendments,” Federal Register (Dec. 4, 2025). Available: https://www.federalregister.gov/documents/2025/12/04/2025-21955/medical-devices-quality-management-system-regulation-technical-amendments
8. U.S. Food and Drug Administration (FDA), “Quality Management System Regulation (QMSR): Final Rule Amending the Quality System Regulation – Frequently Asked Questions,” FDA web page (last updated Aug. 27, 2025). Available: https://www.fda.gov/medical-devices/quality-system-qs-regulationmedical-device-current-good-manufacturing-practices-cgmp/quality-management-system-regulation-final-rule-amending-quality-system-regulation-frequently-asked
9. U.S. Food and Drug Administration (FDA), “Quality System (QS) Regulation/Medical Device Current Good Manufacturing Practices (CGMP),” FDA web page (Jan. 31, 2024). Available: https://www.fda.gov/medical-devices/postmarket-requirements-devices/quality-system-qs-regulationmedical-device-current-good-manufacturing-practices-cgmp
10. International Council for Harmonisation (ICH), “ICH Harmonised Guideline: Quality Risk Management Q9(R1),” Final version adopted Jan. 18, 2023 (minor correction Jan. 15, 2025). Available: https://database.ich.org/sites/default/files/ICH_Q9%28R1%29_Guideline_Step4_2025_0115.pdf
11. U.S. Food and Drug Administration (FDA), “Electronic Records; Electronic Signatures,” Final Rule, Federal Register, vol. 62, no. 54, Mar. 20, 1997 (Doc. 97-6833). Available: https://www.govinfo.gov/content/pkg/FR-1997-03-20/pdf/97-6833.pdf
12. C. A. E. Goodhart, “Problems of Monetary Management: The UK Experience,” in Papers in Monetary Economics, Reserve Bank of Australia, Sydney, 1975.
13. M. Strathern, “‘Improving ratings’: audit in the British University system,” European Review, vol. 5, no. 3, pp. 305–321, Jul. 1997. doi:10.1002/(SICI)1234-981X(199707)5:3<305::AID-EURO184>3.0.CO;2-4.
14. S. Kullback and R. A. Leibler, “On Information and Sufficiency,” The Annals of Mathematical Statistics, vol. 22, no. 1, pp. 79–86, 1951. doi:10.1214/aoms/1177729694.
15. J. Lin, “Divergence Measures Based on the Shannon Entropy,” IEEE Transactions on Information Theory, vol. 37, no. 1, pp. 145–151, Jan. 1991. doi:10.1109/18.61115.
16. A. Rényi, “On Measures of Entropy and Information,” in Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, vol. 1, pp. 547–561, 1961. Available: https://projecteuclid.org/euclid.bsmsp/1200512181
17. C. Villani, Optimal Transport: Old and New. Springer, 2008. ISBN: 978-3540710493.
18. Y. Rubner, C. Tomasi, and L. J. Guibas, “The Earth Mover’s Distance as a Metric for Image Retrieval,” International Journal of Computer Vision, vol. 40, no. 2, pp. 99–121, 2000. doi:10.1023/A:1026543900054.
19. R. P. Adams and D. J. C. MacKay, “Bayesian Online Changepoint Detection,” arXiv:0710.3742, 2007. Available: https://arxiv.org/abs/0710.3742

20. M. Zadehnour, “QRTF™ v1.0 specification, schema, and Pilot10 dataset,” supplementary materials accompanying this Zenodo deposit, 2025. [REF: To be verified after DOI assignment].

10. Public Datasets for Validation (Planned)

This section identifies authoritative public post-market and regulatory signal sources that can be used to validate QRTF™ beyond the illustrative Pilot10 package. This is a validation plan; no new empirical results are claimed here.

10.1 FDA openFDA MAUDE device adverse event reports

The openFDA device adverse event API provides access to FDA MAUDE reports (Medical Device Reports), covering publicly releasable records from roughly 1992 to the present. These reports can be aggregated to outcome intensity series (e.g., monthly counts by product code / device key) for backtesting RiskIndex trajectories.

10.2 FDA openFDA device recalls and FDA recall portal

The openFDA device recall endpoint provides structured recall records; the FDA recall portal provides recall classifications since November 2002. Recall class and reason can serve as severity proxies and outcome anchors.

10.3 Inspection observation summaries (Form 483)

FDA publishes fiscal-year spreadsheets summarizing the areas of regulation cited on system-generated FDA 483s. While not comprehensive, these data are useful as macro-level governance signals and for contextualizing shifts in regulatory observations.

10.4 UK MHRA field safety notices (FSNs)

UK MHRA publishes device safety communications, including field safety notices. These can be normalized into a monthly FSN/FSCA intensity series for selected device categories.

10.5 Swissmedic FSCA database (CH/LI)

Swissmedic publishes field safety corrective actions and associated field safety notices for devices on the Swiss/Liechtenstein market. These can support cross-jurisdiction robustness checks.

10.6 Health Canada recalls and medical device incidents

Health Canada provides recall and safety alert feeds (CSV/JSON) and a medical device incident listing with a downloadable extract. These can serve as additional outcome sources.

10.7 Germany BfArM FSN/FSCA postings

BfArM provides public postings of manufacturer measures including FSNs/recalls, enabling a complementary EU-member-state lens while EUDAMED is not fully public.

10.8 Normalization and evaluation approach

All sources should be mapped into a common OutcomeSignal schema (date, jurisdiction, device_key, signal_type, severity_proxy, traceability URL). Evaluation should focus on calibration, lead time, false alerts, flapping index, and monotonicity violations. See the accompanying “Public Data Benchmark Pack v1.0” supplement for acquisition recipes and metric definitions.

References for Section 10

openFDA device/event: https://open.fda.gov/apis/device/event/

openFDA device/recall: https://open.fda.gov/apis/device/recall/

FDA recall portal: https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfres/res.cfm

FDA inspection observations: https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/inspection-references/inspection-observations

MHRA alerts: https://www.gov.uk/drug-device-alerts

Swissmedic FSCA: https://fsca.swissmedic.ch/

Health Canada recalls: https://recalls-rappels.canada.ca/en ; MDI extract: https://hpr-rps.hres.ca/mdi_landing.php

BfArM manufacturer measures: https://www.bfarm.de/DE/Medizinprodukte/Aufgaben/Risikobewertung-und-Forschung/Massnahmen-von-Herstellern/_node.html

Additional related works (public):

Zadehnour, M. (2025). Regulatory Drift Budget (RDB): A Quantitative, Auditable Risk-Control Framework for Dynamic Change Management in SaMD, IVD, and Medical Devices. Zenodo. doi:10.5281/zenodo.17518146

Zadehnour, M. (2025). Regulatory Readiness Level–Pharma (RRL-P): A Nine-Level Quality-Maturity Framework and Predictive Model for First-Pass Drug Approval. Zenodo. doi:10.5281/zenodo.15486646

Zadehnour, M. (2025). RRL-QMS Overview (slide deck). Zenodo. doi:10.5281/zenodo.15468504