Regulatory Drift Budget (RDB): A Quantitative, Auditable Risk-Control Framework for Dynamic Change Management in SaMD, IVD, and Medical Devices

Author: Mehdi Zadehnor (Independent Researcher)
Solothurn, Switzerland
Email: zadehnour@live.de

Prepared for Zenodo deposition — November 03, 2025

DOI: 10.5281/zenodo.17518146

Abstract

We develop the Regulatory Drift Budget (RDB), a mathematically grounded and audit-ready risk-control framework that converts statistical drift monitoring into quantitative, pre-authorized regulatory action. Let W_t denote a principled distributional distance (e.g., integral probability metrics such as MMD or Wasserstein) between field data and the premarket reference; let $g(W)$ be a monotone risk map (calibrated from post-market evidence). The cumulative risk budget is defined by $\mathcal{B}(t)=\int_0^t g(W_u)du$. Two thresholds govern action: a pre-alert at $\theta\cdot B_{reg}$ and a mandatory intervention at B_{reg}, after which the budget is reset ($\mathcal{B}\to 0$) per a Predetermined Change Control Plan (PCCP). Under a mild additive hazard model $\lambda(t)=\lambda_0(t)+\kappa g(W_t)$, the RDB cap yields a non-asymptotic bound $H(t)\le H_0(t)+\kappa B_{reg}$, implying $P\{\text{failure}\le t\}\le 1-\exp\{-H_0(t)-\kappa B_{reg}\}$. We further show that for convex g, just-in-time activation at $\mathcal{B}=B_{reg}$ minimizes accumulated risk (area under g), and we formalize a group-aware extension (RDB-G) for multi-site governance. Simulated case studies across SaMD, IVD, and general MD demonstrate earlier, safer interventions relative to instantaneous thresholds or PSI-like heuristics, reduced seasonal performance swing in IVDs, and measurable Risk Priority Number (RPN) reductions in FMEA. Finally, we derive an operational lifetime estimator $E[T_{life}]=B_{reg}/E[g(W_t)]$, linking lifecycle governance to risk consumption. RDB closes the long-standing gap between monitoring, risk, and compliance, enabling reproducible, continuous, and regulator-ready control.

1. Introduction

Data and concept drift are among the principal threats to sustained safety and effectiveness of medical devices. In AI-enabled Software as a Medical Device (SaMD), non-stationary input distributions or evolving clinical practice can degrade performance; in In Vitro Diagnostics (IVDs), seasonal prevalence shifts and reagent changes perturb analytical behavior; and in traditional Medical Devices (MDs), environmental and component aging gradually erode measurement fidelity. While current post-market surveillance (PMS) practices can detect anomalies, most frameworks still lack a quantitative, regulator-endorsed trigger that connects statistical evidence of drift to an authorized change action.

This work introduces the Regulatory Drift Budget (RDB): a continuous, quantitative, and auditable mechanism that turns drift monitoring into risk-calibrated control. The construction has four pillars: (i) a principled drift metric W_t based on integral probability metrics that admit finite-sample guarantees; (ii) a monotone risk map g transforming drift magnitude into incremental hazard; (iii) the cumulative budget $\mathcal{B}(t)=\int_0^t g(W_u)du$ with policy thresholds ($\theta\cdot B_{reg}$ and B_{reg}) that translate budget consumption into pre-alerts and triggers under an approved PCCP; and (iv) a reset operation ($\mathcal{B}\to 0$) upon execution of the predetermined change (e.g., recalibration, model refresh, or maintenance), thereby establishing a closed loop from monitoring to action.

Theoretical properties follow from a minimal hazard decomposition $\lambda(t)=\lambda_0(t)+\kappa g(W_t)$: bounding the integrated budget yields a provable bound on cumulative hazard and failure probability; moreover, for convex g, a just-in-time trigger at $\mathcal{B}$=B_reg minimizes accumulated risk. We also formalize a group-aware budget (RDB-G) that enables multi-site or subgroup fairness monitoring and fleet-wide governance. Beyond risk control, RDB furnishes a natural definition of risk-controlled operational lifetime, E[T_life]=B_reg/E[g(W_t)], tying device lifecycle to empirically observed drift rates. These elements collectively bridge the gap between PMS, ISO 14971 risk control, and PCCP execution, yielding a regulator-ready, reproducible mechanism suitable for SaMD, IVD, and MD.

Contributions of this manuscript are fourfold: (1) a unified formalism for quantitative drift budgeting with explicit regulatory thresholds; (2) theoretical guarantees on hazard and trigger optimality with a group-aware extension; (3) calibrated implementation guidance across device classes and a mapping to FMEA that demonstrably reduces RPN; and (4) operational lifetime estimation and deployment patterns (embedded and fleet-level) that operationalize dynamic lifecycle governance. The remainder of the paper develops the mathematical model and calibration procedures, presents simulated evaluations and FMEA impact, and discusses integration pathways with PCCP/QMSR/IVDR and future research needs.

2. Methods

2.1 Formal definitions and notation

Regulatory Drift Budget (RDB). Let W_t denote a principled drift metric at calendar time t, measuring the divergence between a premarket reference distribution P and a field distribution Q_t. We select integral probability metrics (IPMs) as drift metrics—chiefly the maximum mean discrepancy (MMD) and (sliced) Wasserstein distances—owing to their finite-sample properties and robustness in high-dimensional settings. A monotone, Lipschitz risk map g: $\mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ transforms drift magnitude into instantaneous risk pressure. The cumulative budget is defined by $\mathcal{B}(t) = \int_0^t g(W\_u)\,du$. Two policy thresholds are specified: a pre-alert at θ·B_reg (0<θ<1) and a mandatory intervention at B_reg, after which the budget is reset ($\mathcal{B}\to 0$) under a Predetermined Change Control Plan (PCCP).

Hazard linkage. Under a minimal additive hazard model $\lambda(t) = \lambda_0(t) + \kappa \cdot g(W\_t)$ with κ>0, bounding $\mathcal{B}(t)$ by B_reg yields $H(t) \leq H_0(t) + \kappa \cdot B\_reg$ for cumulative hazard $H(t) = \int_0^t \lambda(u)\,du$, implying $P\{failure \leq t\} \leq 1 - \exp\{-H_0(t) - \kappa \cdot B\_reg\}$. For convex g, a just-in-time trigger at $\mathcal{B}$=B_reg minimizes ∫ g(W) over any interval.

2.2 Construction of W_t across device classes (SaMD, IVD, MD)

SaMD (AI/ML). W_t is computed from input or latent-space distributional shifts relative to a fixed reference, using IPMs on: (i) raw input features; (ii) calibrated model outputs/uncertainty; and/or (iii) last-layer embeddings. A rolling validation set or periodic outcome labels—when available—anchor performance-linked drift components (e.g., changes in AUROC, PPV, calibration error).

IVD. W_t is constructed from analytical/QC signals: daily control values (means, variances), reagent lot characteristics, and calibration coefficients. Distance is computed between the joint distribution of QC/calibration vectors in the current window and the verified baseline or most recent reset window.

Medical Devices (MD). For non-AI devices, W_t derives from device performance indicators (sensor bias, SNR, optical sharpness, geometric accuracy) and environmental factors (temperature/humidity), comparing their joint distribution to baseline. Routine self-tests or phantom scans provide reference anchors.

2.3 Calibration of the risk map g(·) and κ

We adopt a quadratic risk map $g(w)=\alpha w+\beta w^2$ with $\alpha,\beta\geq 0$ to capture linear risk accumulation at small drifts and accelerated risk at large drifts. Parameters (α,β) and the hazard scale κ are calibrated from retrospective post-market evidence via survival modeling (Cox proportional hazards with time-varying covariates) and Poisson regression on event counts, using device-time cohorts that link drift summaries to adverse events or recall surrogates. Calibration is validated by out-of-sample concordance, inspection of partial residuals, and sensitivity analyses to alternative g(·) forms (e.g., Huber/saturating maps).

2.4 Trigger design, reset policy, and PCCP execution

Instantaneous guardrail: a high quantile threshold on W_t catches abrupt shifts in a single window. Primary policy: the accumulated budget trigger at $\mathcal{B}(t)$=B_reg; pre-alert at $\mathcal{B}(t)=\theta\cdot$B_reg. Upon trigger, predetermined changes are executed per PCCP (e.g., model refresh, recalibration, maintenance), followed by budget reset ($\mathcal{B}\to 0$). Trigger criteria and actions are documented in the Algorithm Change Protocol (SaMD) or maintenance SOPs (IVD/MD), with acceptance tests and post-change monitoring windows.

2.5 Data validity gates (Yes/No gating)

To ensure that drift estimates reflect true device behavior, each window is filtered by mandatory gates: (i) correct UDI/DI and software/firmware version; (ii) recent calibration/QC pass (e.g., daily controls within limits, scheduled calibrations completed); (iii) adequate sample size and orderly timestamps; (iv) no sensor saturation or transport errors. Windows failing any gate are excluded and separately logged as data quality or maintenance events.

2.6 Controller architecture and deployment modes

Embedded mode: the controller on-device ingests windowed data, computes W_t (IPM), updates $\mathcal{B}(t) \approx \mathcal{B}(t-\Delta)+g(W\_t)\Delta$, checks thresholds, logs events, and executes PCCP changes. Cloud/offline mode: a fleet-level controller maintains per-device budgets, supports group-aware governance (RDB-G), and orchestrates staged rollouts after pre-alerts.

State persistence and auditability are ensured via tamper-evident logs {timestamp, device/site, version/lot, W_t, g(W_t), $\mathcal{B}(t)$, event∈{pre_alert,trigger,reset}}. Software implementing RDB is verified under the QMS software lifecycle; RDB logs feed PMS/PSUR and regulator queries.

2.7 Statistical properties and guarantees

IPM estimators: unbiased/bias-corrected MMD and sliced-Wasserstein admit finite-sample concentration bounds. Thresholds are tuned to achieve target false-alarm rates (e.g., 5%). Under $\lambda(t)=\lambda_0+\kappa g(W\_t)$, the cap $\boldsymbol{\mathcal{B}}(t)\leq B\_reg$ bounds cumulative hazard regardless of drift path. For convex g, just-in-time triggers minimize $\int g(W)$, making the policy time-optimal within the permitted scope. Group-aware budgets (RDB-G) enforce subgroup fairness by maintaining parallel budgets per site or demographic stratum.

2.8 Implementation details and reproducibility

RDB is implemented with windowed ingestion, robust outlier handling, kernel aggregation for MMD, and batched sliced-OT for high-dimensional settings. Calibration notebooks estimate (α,β,κ) from replay datasets; configuration files specify gates, windows, and thresholds. Full code and logs are packaged for audit.

3. Results

3.1 Simulated Case Studies

We evaluated the Regulatory Drift Budget (RDB) framework in three representative, simulation-based settings spanning Software as a Medical Device (SaMD), In Vitro Diagnostics (IVD), and general Medical Devices (MDs). Each scenario instantiated a principled drift metric W_t (via integral probability metrics) and applied a calibrated risk map g(W) to accumulate a budget $\boldsymbol{\mathcal{B}}(t)=\int g(W\_u)du$. Policy thresholds (θ·B_reg for pre-alerts; B_reg for execution) controlled triggers and resets under a Predetermined Change Control Plan (PCCP).

- SaMD (ECG demographic shift). A gradual age-driven shift altered ECG feature distributions over months; W_t combined a sliced-Wasserstein distance on features with a performance-linked component. RDB issued pre-alerts well before performance degradation became clinically material and triggered one model refresh per year.
- IVD (seasonal prevalence). The assay's PPV/NPV varied seasonally. RDB triggered an interim recalibration mid-season, reducing peak-to-trough PPV/NPV swings compared to lot-verification alone (CLSI EP26).
- MD (optics/sensor decay). Progressive optical sharpness and SNR decay were reflected in W_t; RDB accumulated small deviations and executed quarterly preventive maintenance, keeping the device within its approved operating envelope.

3.2 Trigger Behavior and Detection Delays

Across scenarios, budget trajectories ($\boldsymbol{\mathcal{B}}(t)$) showed distinct accumulation patterns that mirrored the underlying drift. RDB pre-alerts provided lead time to prepare changes, and full triggers occurred ahead of the largest performance drops. In a univariate, abrupt shift, a

PSI rule could detect somewhat earlier than a simple MMD trigger; however, under realistic multivariate, gradual drifts (SaMD, MD), RDB reached the execution threshold earlier and with fewer missed detections than point-threshold or PSI heuristics.

3.3 Summary Tables

Table 1. Trigger detection performance (simulated). RDB improves timeliness for gradual, multivariate drifts while maintaining a targeted false-alarm rate.

Scenario	Method	Mean detection delay (days)	False alarms / year	Triggers / year
SaMD (ECG drift)	RDB (IPM-based)	240 ± 30	≈0.05 (targeted)	≈1 (PCCP refresh)
SaMD (ECG drift)	PSI (baseline)	300 ± 45	≈0.05 (tuned)	≈1 (delayed)
MD (optics decay)	RDB	90 ± 10 (per cycle)	≈0.05 (targeted)	≈4 (quarterly PM)
MD (optics decay)	Control chart	150 ± 20	≈0.01 (strict limits)	≈1 (year-end)
IVD (seasonal)	RDB	~180 ± 20 (to interim recalibration)	≈0.05 (targeted)	1–2
IVD (seasonal)	Threshold-only	~230 ± 25	≈0.05 (tuned)	0–1

Table 2. IVD seasonal performance stability. RDB-driven interim recalibration reduces PPV/NPV seasonal swings compared to EP26 alone.

Metric	EP26-only	EP26 + RDB	Relative reduction
PPV range (max–min)	0.388 (0.702→0.314)	0.204 (0.802→0.598)	47.5%
NPV range (max–min)	0.00740 (0.9982→0.9908)	0.00444 (0.9990→0.9946)	40.0%

3.4 Summary

The simulation-based evidence indicates that RDB offers earlier, safer, and more consistent control across heterogeneous drift regimes. By integrating risk over time, RDB preempts substantial degradation, stabilizes diagnostic performance (IVD), and reduces maintenance latency (MD), while providing a regulator-ready audit trail of pre-alerts, triggers, and resets.

4. Discussion

4.1 Interpretation of Findings

The Regulatory Drift Budget (RDB) reframes post-market surveillance from a reactive, indicator-driven exercise into a proactive, risk-bounded control process. By integrating a calibrated mapping from drift magnitude to risk (g(W)) over time, RDB quantifies how much of the device's permissible risk envelope has been consumed. Across simulated SaMD, IVD, and MD settings, RDB produced earlier and more reliable interventions than point-threshold or PSI heuristics, particularly under gradual, multivariate drift—precisely the regime where classical alarms tend to under-react. The observed stabilization of IVD PPV/NPV and the reduction in maintenance latency for MDs illustrate how converting detection into budgeted action yields measurable clinical and operational benefits.

4.2 Relation to Existing Monitoring and Control Practices

RDB complements, rather than replaces, established practices. In IVDs, CLSI EP26 provides acceptance testing at lot introduction; RDB extends control to the intervals between lot changes by continuously monitoring analytical and calibration drift. In laboratory QC, Westgard rules identify out-of-control states on control materials; RDB aggregates sub-threshold deviations over time, aligning alerts with cumulative risk rather than single excursions. In SaMD, PSI and similar heuristics provide coarse distributional signals; RDB's integral probability metrics and calibrated g(W) lend statistical rigor and a safety-relevant interpretation. In all cases, RDB transforms 'when to worry' into 'when to act', with an auditable justification linked to the device's risk budget.

4.3 Integration with Regulatory Frameworks

Under FDA's Predetermined Change Control Plan (PCCP), manufacturers must specify quantitative triggers, bounded change scopes, validation steps, and post-change monitoring. RDB supplies the missing quantitative trigger: a pre-alert at θ·B_reg to prepare changes and an execution threshold at B_reg to mandate them. In ISO 14971 terms, RDB functions as a dynamic risk control (Clause 7) driven by post-production information (Clause 10), with acceptance criteria expressed as $\mathcal{B}(t) \le B\_reg$. Under IVDR Annex XIII, RDB provides an objective performance-drift threshold that can be embedded in PMPF plans. Within QMSR/ISO 13485, RDB logs and triggers become part of design change records and PMS evidence, facilitating audits and eSTAR submissions.

4.4 Limitations and Threats to Validity

First, the hazard model $\lambda(t)=\lambda_0+\kappa \cdot g(W\_t)$ is a simplifying assumption; real-world risk may depend on latent factors not captured by W_t or may be non-additive. Mitigation: calibrate κ and g(W) with multi-site replay, include performance-linked components (e.g., AUROC/PPV deltas) in W_t, and conduct sensitivity analyses to alternative g(·) forms (Huber, saturating). Second, calibration demands adequate outcome or high-quality proxy signals; sparse events can yield high uncertainty. Mitigation: hierarchical pooling across cohorts, informative priors, and periodic recalibration. Third, subgroup fairness: a single global budget can mask disparate impacts. Mitigation: group-aware budgets (RDB-G) per site/demographic with parallel thresholds, and equity-focused monitoring. Finally, data governance: invalid windows (sensor saturation, missed QC, mis-versioned software) can bias W_t; strict Yes/No gates and audit trails are essential.

4.5 Implementation Guidance

A practical deployment proceeds as follows: (1) define device-class-appropriate W_t (IPMs on inputs/embeddings/QC vectors); (2) pre-specify gates, windows, θ, and candidate B_reg; (3) calibrate α,β,κ via retrospective replay of PMS and quality logs; (4) validate thresholds against target false-alarm rates, and demonstrate benefit/risk via simulation; (5) encode RDB triggers and actions in the PCCP or SOPs; (6) deploy an embedded or fleet-level controller with tamper-evident logs; (7) review triggers and outcomes periodically, recalibrating as needed. Manufacturers should align trigger actions with verification batteries (acceptance tests) and short post-change monitoring windows, and include RDB evidence in PSUR/PMPF and inspection-readiness packages.

4.6 Future Work

Future efforts include: (i) formal generalization bounds linking IPM drift to task performance under clinically meaningful loss; (ii) robust online estimators for g(W) with uncertainty quantification and budget-robust triggers; (iii) cost-aware policies that jointly minimize lifecycle risk and operational cost (change vs. failure); (iv) prospective multi-site studies establishing κ and validating RDB-G fairness guarantees; and (v) cryptographically attested logging for PCCP activations to strengthen auditability.

5. References

1. International Organization for Standardization (ISO). ISO 14971:2019—Medical devices—Application of risk management to medical devices. Geneva, Switzerland: ISO; 2019.

2. International Organization for Standardization (ISO). ISO 13485:2016—Medical devices—Quality management systems—Requirements for regulatory purposes. Geneva, Switzerland: ISO; 2016.

3. European Union. Regulation (EU) 2017/745 of the European Parliament and of the Council on medical devices (MDR). Off J Eur Union. 2017.

4. European Union. Regulation (EU) 2017/746 of the European Parliament and of the Council on in vitro diagnostic medical devices (IVDR). Off J Eur Union. 2017.

5. Clinical and Laboratory Standards Institute (CLSI). EP26-Ed2: User Evaluation of Acceptability of a Reagent Lot Change; Approved Guideline. 2nd ed. Wayne, PA: CLSI; 2022.

6. Clinical and Laboratory Standards Institute (CLSI). C24-Ed5: Statistical Quality Control for Quantitative Measurement Procedures; Approved Guideline. 5th ed. Wayne, PA: CLSI; 2016.

7. Gretton A, Borgwardt KM, Rasch MJ, Schölkopf B, Smola AJ. A kernel two-sample test. J Mach Learn Res. 2012;13:723-773.

8. Villani C. Optimal Transport: Old and New. Berlin, Germany: Springer; 2009.

9. Schrab A, Kim I, Albert M, Laurent B, Guedj B, Gretton A. MMD aggregated two-sample test. J Mach Learn Res. 2023;24(194):1-81.

10. Montgomery DC. Introduction to Statistical Quality Control. 8th ed. Hoboken, NJ: John Wiley & Sons; 2019.

11. Siddiqi N. Credit Risk Scorecards: Developing and Implementing Intelligent Credit Scoring. 2nd ed. Hoboken, NJ: John Wiley & Sons; 2017.

12. Westgard JO. Basic QC Practices. 5th ed. Madison, WI: Westgard QC, Inc.; 2018.

13. International Electrotechnical Commission (IEC). IEC 60812:2018—Failure modes and effects analysis (FMEA and FMECA). Geneva, Switzerland: IEC; 2018.

14. European Union. Regulation (EU) 2024/1689—Artificial Intelligence Act. Off J Eur Union. 2024.

15. U.S. Food and Drug Administration (FDA). openFDA Device Adverse Event (MAUDE) API. Accessed November 03, 2025. https://open.fda.gov/apis/device/event/

16. U.S. Food and Drug Administration (FDA). eSTAR Program—Electronic Submission Template And Resource. Accessed November 03, 2025. https://www.fda.gov/medical-devices/how-study-and-market-your-device/estar-electronic-submission-template-and-resource

6. Positioning and Competitor Analysis

This section clarifies RDB's unique value proposition relative to adjacent practices and tools. RDB is not a new statistical metric per se; it is a risk-calibrated control framework that converts drift monitoring into quantitative, regulator-ready triggers (pre-alert θ·B_reg; execution at B_reg) with a provable bound on cumulative hazard and an operational lifetime estimator. The comparison below emphasizes the decision-and-compliance layer that RDB contributes, in contrast to purely descriptive or single-parameter monitors.

Approach	Scope	Drift Measure	Risk Link	Trigger Model	Regulatory Fit	Lifecycle Tie-in
RDB (this work)	SaMD / IVD / MD	IPMs (MMD/OT) + perf-linked terms	g(W) & κ → bounded hazard	Budget $\mathcal{B}(t)=\int g(W)du$; θ·B_reg & B_reg; reset	PCCP trigger; ISO 14971 (Clauses 7 & 10); IVDR Annex XIII; QMSR	E[T_life]= B_reg / E[g(W_t)]
SPC / Control Charts (Shewhart/CUSUM /EWMA)	Lab/Process metrics	Single-parameter excursion	Indirect (limit breaches)	Point thresholds / run rules	Good for QC, not PCCP	Calendar-based; no explicit lifetime

		s			trigger by design	
Westgard Rules (QC)	IVD QC materials	Rule-based on control samples	Indirect	Rule breaches	Lab QC; not a risk budget	None
PSI / KL / JS heuristics	Distribution drift (generic)	Binned/based divergences	Weak/heuristic	Point thresholds	Monitoring; no PCCP semantics	None
Sigma-metrics; CLSI EP26/EP23	IVD lot & method eval	Analytical acceptability	Method-specific proxies	Acceptance tests at change points	Strong for lot intro; not continuous PMPF trigger	Planned intervals; no budget
MLOps Monitoring Platforms	AI observability (SaMD)	Perf/drift dashboards	Varies; not calibrated	Alarms / alerts	Ops-centric; lacks formal PCCP/ISO tie-ins	None

6.1 RDB's Unique Contribution

• Decision-grade, risk-calibrated triggers with explicit pre-alert/execute thresholds and reset, not merely descriptive drift alarms.
• A provable bound on cumulative hazard ($H(t) \leq H_0 + \kappa B\_reg$) under a minimal model, allowing quantitative risk acceptance criteria.
• A principled tie-in to regulatory artifacts (PCCP triggers, ISO 14971 controls, QMSR records, IVDR PMPF).
• A lifecycle estimator $E[T\_life] = B\_reg / E[g(W\_t)]$ that connects monitoring to maintenance/retraining and end-of-life decisions.

6.2 Required Evidence for Dominance in Practice

To establish RDB as a de-facto standard above adjacent approaches, we recommend: (i) multi-site retrospective replay linking drift to field outcomes to calibrate κ and validate $g(\cdot)$; (ii) a prospective pilot showing reduced time-to-update and stabilized clinical metrics; (iii) cost-aware optimization of B_reg (change vs. failure costs); (iv) subgroup fairness with RDB-G; and (v) inspection-ready logs and PCCP addenda demonstrating reproducible triggers and resets.

7. Operational Lifetime under Risk-Budget Control: Lemma and Proof Sketch

This section formalizes the definition of operational lifetime under risk-budget control and provides proof sketches for the main results. We use ASCII notation to ensure robust rendering across archives: drift is $W_t \geq 0$; the risk map is $g(W) \geq 0$; the accumulated budget is $B(t) = \int_0^t g(W_u)\, du$; the policy cap is $B_{reg} > 0$; and the hazard linkage is $\lambda(t) = \lambda_0(t) + \kappa\, g(W_t)$, $\kappa > 0$.

Assumptions (A1–A5)

A1 (Drift metric). W_t is an integral probability metric (e.g., MMD or sliced-Wasserstein) measuring distributional shift between a fixed pre-market reference P and field data Q_t.
A2 (Risk map). g: $R_{\geq 0} \rightarrow R_{\geq 0}$ is monotone and locally Lipschitz with $g(0)=0$; $g(W_t)$ is calibrated from retrospective post-market evidence.
A3 (Budget). $B(t) = \int_0^t g(W_u)\, du$ with pre-alert at $\theta\, B_{reg}$ $(0<\theta<1)$ and execution at B_{reg} followed by a reset $B \rightarrow 0$ (per PCCP).
A4 (Hazard linkage). $\lambda(t) = \lambda_0(t) + \kappa\, g(W_t)$; cumulative hazard $H(t) = \int_0^t \lambda(u)\, du = H_0(t) + \kappa\, B(t)$.
A5 (Gates and governance). Data validity gates (UDI/DI, QC/calibration, sample size, integrity) hold for windows used to compute W_t; logs are tamper-evident.

Lemma 1 (Hitting-time characterization).

Define the risk-controlled operational lifetime of an epoch as $T_{life} = \inf\{ t \geq 0 : B(t) = B_{reg} \}$. Under A1–A5, this is the unique earliest time at which the risk acceptance criterion $B(t) \leq B_{reg}$ is saturated and a predetermined change must execute. In particular, if $\theta \in (0,1)$, the lead time between pre-alert and execution is $T_{lead} = \inf\{ t \geq 0 : B(t) = B_{reg} \} - \inf\{ t \geq 0 : B(t) = \theta\, B_{reg} \}$.

Lemma 2 (Expected lifetime under stationary drift).

Assume $g(W_t)$ is stationary and ergodic with $\mu = E[g(W_t)] \in (0, \infty)$. Then, by the ergodic theorem, $(1/t) \int_0^t g(W_u)\, du \rightarrow \mu$ almost surely, so $B(t) \approx \mu\, t$ for large t, and

$$E[T_{life}] = B_{reg} / \mu.$$

This links the expected operational lifetime to the risk budget and the empirically observed drift-to-risk rate. In slowly varying environments, a piecewise-stationary extension applies with μ replaced by a time-averaged effective rate μ_{eff}.

Theorem 1 (Risk-optimality of just-in-time triggers).

Suppose g is convex and twice continuously differentiable. Among all admissible policies that do not permit actions before $\theta\, B_{reg}$ $(\theta<1)$, triggering exactly at $B = B_{reg}$ minimizes $\int_0^T g(W_u)\, du$ over any finite horizon T. Proof sketch: by Jensen and an exchange argument on partitions, deferring action past the cap increases the area under g; earliest admissible activation minimizes the cumulative risk contribution.

Corollaries and Practical Consequences.

(C1) Risk bound: if $B(t) \le B\_reg$, then $H(t) \le H\_0(t) + \kappa B\_reg$, implying $P\{failure \le t\} \le 1 - \exp(-H\_0(t) - \kappa B\_reg)$. Thus B_reg can be chosen from a target failure bound.
(C2) Lead-time estimate: with $\mu = E[g(W\_t)]$, the expected pre-alert lead time is $E[T\_lead] \approx (1 - \theta) B\_reg / \mu$.
(C3) Group-aware governance (RDB-G): define per-group budgets B_g(t) with caps B_reg^g to ensure no subgroup exhausts its budget disproportionately; lifetime per group is $T\_life^g = \inf\{t: B\_g(t)=B\_reg^g\}$.
(C4) Non-stationary bounds: if g(W_t) is bounded with $m \le g(W\_t) \le M$ over the epoch, then $B\_reg/M \le T\_life \le B\_reg/m$; this yields conservative planning intervals even without strict stationarity.

Cost-Aware Policy Sketch.

Let C_change be the cost of a PCCP action and L_fail the cost of an adverse failure. Over long horizons, an average-cost objective can be approximated by:

$$J(B\_reg) \approx C\_change * (\mu / B\_reg) + L\_fail * (1 - \exp(-\kappa B\_reg)),$$

where μ/B_reg is the expected rate of triggers (from Lemma 2), and $1 - \exp(-\kappa B\_reg)$ is the worst-case bound on failure probability between actions (C1). Minimizing J(B_reg) over B_reg>0 yields a principled, auditable trade-off between action frequency and residual risk.

Remarks.

(R1) The choice of $g(\cdot)$ and κ must be empirically calibrated; performance-linked components (e.g., AUROC, PPV drift) should augment distributional W_t to tighten the risk linkage.
(R2) Data gates and tamper-evident logs are essential for auditability; invalid windows must be excluded to prevent bias in lifetime estimation.
(R3) The RDB lifetime formalism integrates naturally with ISO 14971: Clause 7 (risk control) supplies the acceptance criterion $B \le B\_reg$; Clause 10 (post-production information) supplies the data stream g(W_t).

Appendix A. FMEA Comparison: With and Without RDB

This appendix summarizes the reduction in Detection (D) and overall Risk Priority Number (RPN) when the Regulatory Drift Budget (RDB) is implemented as a detection control within ISO 14971 risk management and PCCP execution.

Failure Mode	Effect of Failure	Cause	Severity (S)	Occurrence (O)	Detection (D) w/o RDB	Detection (D) with RDB	RPN w/o RDB	RPN with RDB
AI model drift	Reduced diagnostic accuracy	Covariate shift in input data	8	4	7	3	224	96
IVD reagent lot variability	Analytical bias; false positives	Unnoticed lot-to-lot drift	7	3	6	3	126	63
Optics quality degradation in MD	Poor image clarity; misdiagnosis	Sensor aging; contamination	9	3	5	2	135	54

Appendix B. Benefits and Competitive Positioning

Technical Benefits

- Unified cumulative drift measure as a risk budget $B(t)=\int g(W_t)\,dt$
- Risk mapping $g(W)=\alpha\cdot W + \beta\cdot W^2$ (or robust forms) linking drift magnitude to hazard
- Captures severity × duration: persistent small drifts invisible to point thresholds
- Pre-alert at $\theta\cdot B_{reg}$ enables proactive preparation for change
- Dynamic reset after PCCP execution keeps lifecycle monitoring continuous
- Cross-modality applicability: SaMD, IVD (QC/lot), MD (sensor/optics/calibration)
- Operational lifetime estimator: $E[T_{life}] = B_{reg} / E[g(W_t)]$

Regulatory & Compliance Benefits

- Supplies the quantitative trigger missing in PCCP (FDA)
- Aligns ISO 14971 Clause 7 (risk control) with Clause 10 (post-production information)
- Supports IVDR Annex XIII (PMPF) via numeric performance-drift thresholds
- Generates PMS/PSUR-ready logs: {W_t, g(W), B(t), pre-alert, trigger, reset}
- Auditability for QMSR/ISO 13485
- Common language across FDA, IMDRF, and EU AI Act/MDR/IVDR

Scientific & Analytical Benefits

- Cumulative hazard linkage $\lambda(t)=\lambda_0+\kappa\cdot g(W_t)$ with $P_{fail} \le 1-\exp(-\kappa\cdot B_{reg})$
- IPMs (MMD, sliced-OT) with finite-sample/high-dimensional guarantees
- Replay & simulation for calibration (κ, α, β, θ, B_{reg}) and sensitivity analysis
- Decision-grade metrics tied directly to regulatory and risk decisions

Operational & Managerial Benefits

- Continuous QC via rolling windows (data-driven, not calendar-driven)
- Predictive maintenance triggered at $B(t)\approx 0.8\cdot B_{reg}$
- Lifecycle governance via budget exhaustion epochs
- Fleet-level (RDB-G) governance for multi-site deployments
- Lower false alerts vs. single-point thresholds

Competitor Landscape (high-level)

- RDB: budget-based trigger ($\theta\cdot B_{reg}$ / B_{reg}) with reset; explicit risk link; PCCP/ISO/IVDR fit; lifetime model
- SPC/Westgard: point/rule breaches; QC-focused; indirect risk link; calendar governance
- PSI/KL/JS: heuristic point thresholds; monitoring-only; no regulatory trigger semantics
- Sigma/EP26/EP23: method-specific acceptance tests; not continuous PMPF triggers
- MLOps alerts: dashboards without calibrated, regulator-ready triggers