Sequential Certificate Risk for Declared RF Carriers

A finite-horizon technical addendum

Prepared for Brian Theory

AI-assisted derivation and internal mathematical review

9 November 2026 Release candidate 0.1

Abstract

We extend a fixed-state RF observation model by allowing explicitly declared carriers and charging the first exclusion of the true state. This changes the admissibility contract; it does not remove states from exact posterior support. We formulate a complete finite-history linear program for a survival-conditional remaining-risk bound imposed after each action is selected. In a four-state, known-noise example, its five-observation minimax loss is exactly 267963/640625. A specified prior/posterior-ratio gate has loss 291433/640625, while a tailored time-uniform gate attains the former value. An exact bounded calculation shows immediate caps imply the full conditional bound through six observations in this example. Enlarging one carrier can nevertheless invalidate a later conditional bound. These are specialized benchmarks and contract distinctions, not new general methods of chance constraint optimization or confidence-sequence construction.

Editorial status. This is an integration-ready addendum, not a claim that a separate general manuscript is warranted. It supplements *From Observable States to Acquirable Predictive Models* [1] and cross-references *Diagnosing While Repairing*, companion v0.2.5 [2]. AI proof and hostile reviews are internal assessments, not human peer review. External priority, authorship approval and physical applicability remain unresolved. No physical safety or consciousness conclusion is drawn. Publication is a separate editorial action.

1 Increment beyond the existing manuscripts

The local RF predicate and its endpoint characterization are inherited from [1, Lemma 1]; we do not re-prove the structural classification. The preserved self-diagnosis predecessor already proves that positive likelihoods preserve exact support, distinguishes a probabilistic gate from that support rule, warns against reusing a per-use risk allowance, and notes generic time-uniform coverage [3, section "Cost and risk as parts of the contract"]. Those ideas are antecedents here.

The companion v0.2.5 studies silent repairs that change current state while diagnosing an origin. Its Sections 5–9 address charged readout, minimax error, terminal recalibration, acquisition randomization and a statewise error/read-cost problem [2]. Our state never changes, our cost has no repair component, and our failure event is carrier exclusion. Thus neither its numerical frontiers nor its positive-noise restoration obligation can be substituted for ours. The incremental material is an explicit first-exclusion contract, its complete conditional realization rows, and exact finite RF benchmarks. The general oracle/amplification characterization from the earlier noise report is omitted: its method overlaps the predecessor and is unnecessary for this addendum.

2 Fixed model and policy information

Let $H = (a, x, y, z)$ have order $a < x, y < z$, uniform prior $\pi_h = 1/4$, target labels $\ell = (0, 1, 2, 0)$, and supplied queries $q = 0110$, $r = 1100$. This same diamond, query library and target occur in the predecessor's subsection "A noiseless RF example and an explicit noisy lift" [3]; the new object here is the declared-carrier risk contract. Each execution is corrupted by fresh independent binary-symmetric noise of known error $1/5$. The state is fixed; queries have no cumulative structural effect or consumed permission. An execution costs $c = 1/50$. A policy may stop and report any label, use fresh behavioral randomization, and take at most N observations. The public history contains all past queries, carriers and answers. Future coins are independent of the state conditional on that history; unobserved survival is not extra information available to the controller.

Each query declares a nonempty carrier $C \subseteq H$. RF legality means the pre-answer order-complex inclusion under Boolean order thinning on C is a homotopy equivalence. With bit masks $(a, x, y, z) = (1, 2, 4, 8)$ the complete list is

$$\begin{aligned} q &: 1, 2, 3, 4, 5, 6, 7, 8, 9, 11, 13, 15, \\ r &: 1, 2, 3, 4, 6, 8, 12. \end{aligned} \tag{1}$$

These 19 actions follow from the endpoint criterion [1]. The maximal carriers are H for q and ax, xy, yz for r. Positive noise retains full posterior support. Declared-carrier execution is a deliberate relaxation, not a theorem licensing r under the original exact-support rule.

Let F be the event that some executed carrier excludes the true state. Write $r_h = \mathbb{P}_h(F)$, $e_h = \mathbb{P}_h(\widehat{\ell} \neq \ell(h))$ and $n_h = \mathbb{E}_h T$. The objective is

$$V = \max_h (e_h + c n_h), \tag{2}$$

not $\max_h e_h + c \max_h n_h$. Bayesian episode risk is $\sum_h \pi_h r_h$; statewise episode risk bounds each r_h. Neither is a bound on physical hazards. Error and certificate failure are distinct events.

3 First exclusions and the conditional contract

Before a query, let b be the full posterior and D the intersection of all previously executed carriers, initially H. At a selected carrier C, new first-exclusion mass is $b(D \setminus C)$.

Proposition 1 (First-exclusion accounting)**.** *For any finite policy,*

$$\mathbb{P}_\pi(F) = \mathbb{E}_\pi \sum_t b_t(D_t \setminus C_t). \tag{3}$$

Proof. The event of first exclusion at time t is $\{h \in D_t \setminus C_t\}$. These events are disjoint on the joint state/history space. Condition each indicator on the history including the selected action, and sum. The conditional distribution of h is still b_t because the selected action is conditionally state-independent. No independence of failure events is used. After stopping take $C_t = H$ formally. The same identity extends by monotone convergence when infinitely many gates are permitted. □

Reusing a carrier adds no new certificate failure. Reintroducing excluded states later cannot undo past failure. If D is empty, future increments vanish because failure has already occurred, not because coverage recovered.

For history η and selected test action u, define

$$\rho(\eta, u) = \mathbb{P}_\pi(\text{first exclusion now or later} \mid \eta, u, h \in D_\eta). \tag{4}$$

HC-commit requires $\rho(\eta, u) \le \alpha$ at each positive-probability commit with positive survival mass. Conditioning occurs *after* action selection and *before* execution; the current exclusion counts. HC-pre would average the current action coin and is different. A separate constraint under each surviving state would disallow all exclusions for $\alpha < 1$ and is not the Bayesian conditional contract (4).

Put $\beta = b(D \setminus C)/b(D)$. Let $p_{\mathrm{surv}}(y)$ condition on $h \in D \cap C$, and let $\bar{\rho}_y$ average future action coins. Then

$$\rho = \beta + (1 - \beta) \sum_y p_{\mathrm{surv}}(y) \bar{\rho}_y. \tag{5}$$

Equation (5) uses $\beta < 1$. If $\beta = 1$, then $\rho = 1$ and no survivor observation law is defined; such a commit is infeasible for $\alpha < 1$. This is the first-step decomposition of disjoint failure events. Survival conditioning and execution-risk accounting have close antecedents in RAO* [4, Sections 2–3]. Our derivation uses a static certificate event; it imports no assertion about physical safety. In general local $\beta \le \alpha$ does not establish (4). Also

$$\mathbb{P}_\pi(F \text{ eventually} \mid \eta) = b(D^c) + b(D)\bar{\rho}(\eta). \tag{6}$$

Full posterior mass, including D^c, remains in every error and cost calculation. None of these formulas resets a risk budget.

4 Complete finite-history linear formulation

At each possible public history η, let $l_{\eta h}$ be the product of observation likelihoods under h, excluding behavioral action factors. Let x_η be the product of those past action probabilities and $x_{\eta u}$ the product after selecting u. Include every legal test/carrier pair and each stop label, with tests removed at depth N. Impose

$$x_\varnothing = 1, \qquad \sum_u x_{\eta u} = x_\eta, \qquad x_{\eta u y} = x_{\eta u} \quad \text{for each answer } y, \qquad x \ge 0. \tag{7}$$

The children do not sum to $x_{\eta u}$: the answer probabilities already occur in l. This is the standard realization-plan construction under perfect recall [7], specialized to one controller.

For test edge $v = (\nu, u)$ with carrier C_v, put

$$f_{vh} = l_{\nu h} \mathbf{1}\{h \in D_\nu \setminus C_v\}, \qquad f_v = \sum_h \pi_h f_{vh}.$$

Episode risk is $R_h(x) = \sum_v f_{vh} x_v$. Terminal error and expected tests are the linear forms

$$E_h(x) = \sum_{(\eta, j)\,\mathrm{stop}} l_{\eta h} x_{\eta j} \mathbf{1}\{\ell(h) \ne j\}, \qquad T_h(x) = \sum_{(\eta, u)\,\mathrm{test}} l_{\eta h} x_{\eta u}. \tag{8}$$

In particular (8) does not restrict sums to surviving states.

Proposition 2 (Exact HC rows)**.** *Let* $\mathrm{Sub}(v)$ *include a committed edge* $v = (\eta, u)$ *itself and every test edge below it. The additional rows*

$$\sum_{w \in \mathrm{Sub}(v)} f_w x_w \;\le\; \alpha x_v \sum_{h \in D_\eta} \pi_h l_{\eta h} \tag{9}$$

are necessary and sufficient for HC-commit on the public-history behavioral class. Thus minimizing t *subject to (7), (9), and* $E_h + cT_h \le t$ *is a complete finite LP. Optional episode rows are* $R_h \le \delta$*.*

Proof. The left side is the joint probability of reaching the selected subtree and then first failing there. Later intersections are contained in D_η, so the disjoint-event numerator counts only states alive at η. The right side is α times the probability of reaching that commit alive. Division yields (4). At zero action weight, nonnegative flows force all descendants to zero. At zero survival mass the conditional quantity is undefined and no risk restoration is inferred. Conversely, positive-node ratios $x_{\eta u}/x_\eta$ recover behavioral probabilities; zero-node behavior is arbitrary. These probabilities reproduce all linear forms and conditional ratios. Therefore every feasible realization gives an admissible policy and every such policy gives a feasible realization. □

A known root mixture seed included in the conditioning information defines a stronger contract than (4) generally. Every HC component exhibited below separately passes (4), so the achieved value also meets that seed-aware interpretation. The lower bound applies to the larger public-history class.

Pruning an edge with $\beta > \alpha$ is necessary and sound; future risk is nonnegative. No maximal-carrier dominance is used. After this necessary pruning, the complete $N = 5$ tree has node counts $1, 2, 4, 8, 28, 128$ at depths 0 through 5, 598 action variables, 171 flow rows and 85 HC rows. Together with t there are 599 variables. The supplied floating LP is a cross-check; the exact proof of its optimum is the rational certificate below.

5 Exact benchmarks and certificate method

Fix $\alpha = \delta = c = 1/50$. All objectives in Table 1 are (2).

Class at $N = 5$	Exact optimum	Approximation
HC-commit, all 19 actions	267963/640625	0.418284
HC-commit, maximal carriers	267963/640625	0.418284
Specified raw PPR core	291433/640625	0.454920
Specified PPR plus HC	291433/640625	0.454920
Statewise episode risk only	3766683/10250000	0.367481

Table 1: Same fixed-state model and objective; distinct admissibility contracts.

For nonnegative w summing to one and $\lambda_h \geq 0$, weak duality gives

$$V \geq \min_P \sum_h \{w_h(e_h(P) + cn_h(P)) + \lambda_h r_h(P)\} - \delta \sum_h \lambda_h. \tag{10}$$

The scalar minimum is computed exactly by backward induction. At likelihood vector l, intersection D and remaining horizon n, minimize the stop cost $\sum_h w_h l_h \mathbf{1}\{\ell(h) \neq j\}$ and, for each allowed test (e, C), the sum

$$c \sum_h w_h l_h + \sum_{h \in D \setminus C} \lambda_h l_h + \sum_y V_{n-1}\big((l_h K_e(y \mid h))_h, D \cap C\big). \tag{11}$$

At $n = 0$ only stops remain. Induction proves (11) minimizes over all deterministic trees. A scalar minimum also bounds their mixtures. The HC lower-bound oracle permits every immediate-cap-feasible action, a superset of the full HC class, so its lower bound is valid without assuming immediate caps are sufficient. The HC primal is audited separately against actual suffix first-failure probabilities. LP scalarization and sequential-testing dynamic programming are classical [5]; our finite proof does not rely on their general regularity assumptions.

For the first row take

$$w = (0, 193/410, 193/410, 12/205), \qquad \lambda = 0.$$

The exact oracle value is 267963/640625. A three-tree feasible mixture has weights

$$115/246, \qquad 46363/92496, \qquad 2893/92496.$$

Its metrics are

h	r_h	e_h	n_h
a, z	1/125	130423/384375	303548/76875
x, y	1/125	134209/384375	265688/76875

Every state's loss equals the lower bound. Every component's maximum HC risk is 1/65. Complete trees, rather than just these aggregated metrics, are supplied in `data/certificates.json`. The maximal-carrier optimum has the same errors/costs, endpoint risks 1/125 and middle-state risks zero.

For PPR, $w = (12/205, 193/410, 193/410, 0)$ and $\lambda = 0$ give the matching bound 291433/640625. Its weights are $281/984, 225055/450672, 96919/450672$; maximum component HC risk is 1/257. Because the HC and PPR dual risk multipliers vanish, their optima remain valid with the optional statewise episode rows removed. The episode-only optimum cannot satisfy HC because its value is strictly below the proven full-HC lower bound. Some components have conditional risk 1/2, which alone would only disprove seed-aware component feasibility, not public-history mixture feasibility.

Randomization is substantive under the declared information. Already at $N = 3$, a deterministic first r would have positive path probability at least 1/25 under an excluded state, violating a 1/50 statewise budget. Thus every feasible deterministic acquisition uses only q, which cannot distinguish x and y; its worst error is at least 1/2. The exact randomized composite optimum is $352229/780000 < 1/2$. This does not say randomization improves a fixed scalar DP, or that the companion's repair-order randomization theorem proves this example.

6 Two exclusions and bounded immediate-cap equivalence

Before the fourth query, no proper carrier passes the immediate cap: all earlier queries must be q/H, and through two answers the smallest posterior mass is at least $1/34 > 1/50$. Two distinct exclusions can first occur at $N = 5$. For example ask q/H three times. Only after 111, ask q/axy; only if its answer is 1 ask r/xy. Stop elsewhere. The fourth remaining risk is $1/130 + (1/5)(1/130) = 3/325$, and the fifth is 1/513. The earlier committed risks along that prefix are $3/1250, 3/625, 3/425$. The sequential exclusions of z and then a are both counted.

Here is an exhaustive finite proof method that tests, rather than assumes, immediate-cap sufficiency. For integer likelihood numerators L, put $m(L, D) = \sum_{h \in D} L_h$. Let $B_n(L, D)$ be the maximum unnormalized future first-failure mass over all immediate-feasible continuations. Then

$$\begin{aligned} &B_0(L, D) = 0, \\ &B_n(L, D) = \max\left(0, \max_{(e,C) \in \mathcal{A}(L,D)} \left[m(L, D \setminus C) + \frac{1}{5} \sum_y B_{n-1}(LK_e^{\mathrm{num}}(y), D \cap C) \right] \right), \end{aligned} \tag{12}$$

where $\mathcal{A}$ is the full legal list (1) restricted by $50m(L, D \setminus C) \le m(L, D)$, and K^{num} has entries 1 or 4. Induction and maximization over the current action prove (12); randomization cannot increase the maximum. Enumerate every reachable (L, D) and check $B_{N-d}(L, D)/m(L, D) \le 1/50$ at depth d.

The exact maxima over reachable states at each decision depth are

N	$d=0$	1	2	3	4	5
5	1/125	1/125	1/85	1/65	1/257	$--$
6	37/3125	37/3125	29/2125	1/65	1/85	1/65

The aggregated state counts are $1, 2, 3, 4, 17$ for $N = 5$, with 27 additional states at depth 5 for $N = 6$. Every entry is below $1/50$. This exact finite enumeration proves immediate caps imply HC through these horizons *in this model at this allowance*. No horizon above six was searched; no general local-cap theorem is asserted.

Example 3 (Carrier expansion need not preserve HC)**.** *Ask q/H three times. Only after* 111 *ask q/xy, and only if its answer is 0 ask r/xy. Stop elsewhere. The fourth risk is* $1/65$*; the fifth is zero because $D = xy$. Earlier risks are* $1/250, 1/125, 1/85$*. Enlarge only the fourth carrier to H, keeping all future queries and stopping rules. After* 1110*, the fifth r/xy now excludes endpoint mass* $1/17 > 1/50$*. Thus enlargement breaks HC despite unchanged observations and performance.*

This refutes a nodewise dominance argument. It does not prove an optimized maximal-carrier value gap: Table 1 establishes equality at the point tested. Similarly, this example does not contradict (12): the expanded policy fails the immediate cap at its fifth query.

7 The specified PPR comparison and a tailored gate

Our PPR comparator uses the *raw current* core

$$A_t = \{h : b_t(h) > \alpha\pi_h\}, \qquad \alpha = 1/50, \tag{13}$$

with original prior throughout, and requires each carrier to contain A_t. It is not a running-intersection or restarted-prior variant. The prior/posterior-ratio confidence-sequence method is established [6]. For completeness, under each true h, $\pi_h/b_t(h) = \sum_g \pi_g L_t(g)/L_t(h)$ is a nonnegative mean-one martingale: tests are predictable and positive likelihood ratios have conditional expectation one. Ville's inequality bounds ever reaching or exceeding $1/\alpha$ by α. Hence legal carrier completion preserves statewise simultaneous coverage. This episode guarantee alone does not prove HC at every history.

After 111, each endpoint mass is $1/130 > \alpha\pi_h = 1/200$; PPR still retains all of H. After 1111, each is $1/514 < 1/200$, so r/xy can first execute fifth. HC already permits it fourth. The optimized composite penalty is

$$V_{\mathrm{PPR}} - V_{\mathrm{HC}} = 4694/128125 \simeq 0.0366361, \tag{14}$$

about 8.759% of V_{HC}. This is neither an extra-samples estimate at fixed error nor evidence that all confidence-sequence methods are suboptimal.

Indeed, use q/H for up to three queries. If the initial answers are 111 and a fourth query is executed, require its actual carrier to be *exactly* xy, using either q/xy or r/xy. All later carriers contain xy. On other branches keep carrier H forever. Early stopping is allowed. The intersection after executed contraction is xy and never shrinks again. Under a, z ever losing coverage requires three erroneous ones, with probability at most $1/125$; under x, y exclusion is impossible. Holding D fixed after stopping gives time-uniform statewise coverage. The committed risk bounds before and at contraction are $1/250, 1/125, 1/85, 1/65$; all later first-failure risks vanish. Thus this gate family satisfies HC even with indefinite permitted continuation. Shrinking a displayed confidence core while executing H would not suffice: the actual intersection must contract. Every component of the maximal-carrier optimum in Table 1 fits this family. A tailored time-uniform gate therefore attains the unrestricted HC optimum at $N = 5$.

8 Scope, correction and release limits

The first research report gave an invalid intersection illustration: q/axz is legal but q/xyz is not, so these cannot witness two legal carriers with illegal intersection xz. The valid non-heredity example is q/H legal versus q/xz illegal. The mask list (1) always excluded xyz, so numerical certificates are unaffected. Exhaustive same-query pair checking shows these two particular legal-carrier families are closed under nonempty intersection. No general intersection theorem follows.

Verification is exact rational arithmetic plus finite enumeration, not proof-assistant certification. The public checker covers 47 saved certificates: 42 earlier frontier/contract certificates and five five-query certificates. A separate fresh closeout check rebuilt the five-query scalar lower bound, leaf metrics and conditional risks without importing the optimizer, and reconstructed the HC rows independently. The full LP cross-check is floating point and is not the proof of exact equality. Search ends at $N = 5$ for optimization and $N = 6$ for immediate-cap feasibility. Unknown sensor laws, shared noise, physical interventions and continuous timing belong to different models and are not investigated here.

The overlap assessment found a suitable technical addendum, not a reason to duplicate the companion. General chance constraints, PPR constructions, sequence form and scalar dynamic programming retain classical attribution. Priority of the particular specialization is not established by internal review. The release should remain linked supporting material unless an editor judges its specialized increment sufficient for a separate paper.

A Earlier finite frontiers retained as exact supporting data

For $N = 3, c = 1/50$, statewise episode cap δ, the exact minimax loss has the following six pieces. Adjacent endpoints agree.

Interval	$V(\delta)$
$[0, 296/8925]$	$25817/48750 - (20281/5200)\delta$
$[296/8925, 16/425]$	$983/1875 - (187/50)\delta$
$[16/425, 4/85]$	$2603/6375 - (33/50)\delta$
$[4/85, 1517/10200]$	$2549/6375 - (12/25)\delta$
$[1517/10200, 14/75]$	$29828/80625 - (12/43)\delta$
$[14/75, 1]$	$596/1875$

At statewise episode cap $1/50$, the minimum $\max_h n_h$ with $e_h \le \eta$ is infeasible below $463/1125$, zero for $\eta \ge 2/3$, and otherwise has seven pieces:

Interval	Minimum $\max_h n_h$
$[463/1125, 3474/8375]$	$46453/1200 - (355/4)\eta$
$[3474/8375, 1717/4125]$	$15187/1200 - (415/16)\eta$
$[1717/4125, 11059/25875]$	$20309/2880 - (2395/192)\eta$
$[11059/25875, 493/1125]$	$86213/13500 - (131/12)\eta$
$[493/1125, 47/90]$	$7639/1260 - (285/28)\eta$
$[47/90, 119/225]$	$319/60 - (35/4)\eta$
$[119/225, 2/3]$	$10/3 - 5\eta$

For the cost dual replace $w_h(e_h + cn_h)$ in (10) by $w_h n_h + \gamma_h e_h$, $\gamma_h \ge 0$, and subtract $\eta \sum_h \gamma_h$. Exact endpoint mixtures attain each interpolated line and a single feasible dual supports it throughout the interval. The supplied checker verifies endpoints and dual coefficients. This proves the displayed lines for all intervening parameters, not only a grid.

At $N = 4$, imposing $e_h \leq 2/5$, the minimum worst-state expected reads is 295163/139600 under statewise episode risk 1/50, versus 1459/560 with HC. The extra 239907/488600 reads is a fixed-error cost comparison; raw PPR is infeasible at that error target. These benchmarks use a static RF-carrier model, not the companion's state-changing repair model.

References

[1] Brian Theory. *From Observable States to Acquirable Predictive Models*. Revised technical report, 7 October 2026. Lemma 1 and Sections 2, 7–8. https://www.b-theory.com/manuscripts-html/acquirable-predictive-models/.

[2] *Diagnosing While Repairing: Information Loss, Adaptive Order, and Noisy Readouts*. Companion v0.2.5, 8 November 2026. AI-assisted research draft prepared for Brian Theory. https://www.b-theory.com/manuscripts-html/diagnosing-while-repairing/.

[3] *Causal Self-Diagnosis with Noisy Tests: Observation Laws, Irreversible Access, and Task Recovery*. Preserved reviewed v0.3, 5 November 2026. AI-assisted predecessor prepared for Brian Theory; not claimed to be an independently published primary-literature result.

[4] P. Santana, S. Thiebaux and B. Williams. RAO*: An Algorithm for Chance-Constrained POMDP's. AAAI, 2016. Sections 2–3. https://doi.org/10.1609/aaai.v30i1.10423.

[5] M. Fauss and A. M. Zoubir. A Linear Programming Approach to Sequential Hypothesis Testing. 2015. Section 4, Theorems 4.1–4.2. https://arxiv.org/abs/1501.05870.

[6] I. Waudby-Smith and A. Ramdas. Confidence sequences for sampling without replacement. NeurIPS, 2020. Section 2.1, Proposition 2.1. https://papers.nips.cc/paper_files/paper/2020/file/e96c7de8f6390b1e6c71556e4e0a4959-Paper.pdf.

[7] B. von Stengel. *Extensive Games and the Sequence Form*. Author lecture slides, sections on realization plans and behavior probabilities. The sequence-form method predates this addendum. https://conferences.mpi-inf.mpg.de/adfocs-24/material/Bernhard/2ext-adfocs.pdf.