

INNER-SHELL RAMAN X-GATE TRADEOFFS FOR A NEUTRAL ^{171}Yb NUCLEAR QUBIT AT FIXED OPTICAL POWER

Anonymous authors

Paper under review

ABSTRACT

We study single-qubit X -gate feasibility for a ^{171}Yb nuclear-spin qubit encoded in the $^3P_0(F = 1/2)$ manifold and driven by Raman coupling through an inner-shell-excited $J = 2$ intermediate. The design objective is explicitly constrained: optical intensity is fixed at 1 W/cm^2 , and the output must report gate-time versus target-infidelity operating points at 10^{-2} , 10^{-3} , 10^{-4} , and 10^{-5} . We develop a hybrid formal-empirical methodology that combines (i) a convex surrogate optimality certificate for detuning selection, (ii) a robust feasibility-floor certificate for impossibility regions, and (iii) a posterior chance-constrained table for uncertainty-aware decision support. The resulting evidence indicates a stable, non-provisional operating regime at 10^{-2} with a best deterministic gate time of $0.1661\ \mu\text{s}$ and dominant leakage contribution. After targeted posterior uncertainty ablations (model-discrepancy and noise-floor scaling), posterior confidence intervals widen materially and strict-threshold recommendations at 10^{-3} and below become infeasible in both deterministic and posterior modules. Those strict rows remain provisional because robust-floor diagnostics still classify 10^{-3} as potentially feasible under conservative assumptions. Beyond this specific gate, the study connects atomic-structure uncertainty management to broader neutral-atom processor planning and high-accuracy metrology workflows.

1 INTRODUCTION

Neutral-atom quantum hardware is now transitioning from proof-of-principle control demonstrations to operating regimes where model fidelity, uncertainty accounting, and architecture-aware calibration are first-order constraints rather than post hoc corrections (Saffman et al., 2010; Levine et al., 2019; Holzl et al., 2024; Morgado & Whitlock, 2021). In alkaline-earth-like systems, ytterbium is an especially interesting platform because precision spectroscopy, long-lived clock states, and neutral-atom array engineering can be integrated in one physical stack (Ishiyama et al., 2023; Kawasaki et al., 2023; 2024; Ludlow et al., 2015; Hoyt et al., 2005; Norcia et al., 2024; Jenkins et al., 2022b). This dual role creates an opportunity and a risk: the same weakly allowed transitions that make ultra-stable metrology possible can also introduce difficult control tradeoffs when repurposed for fast quantum gates.

The concrete problem here is intentionally narrow. We consider an X gate on a ^{171}Yb qubit encoded in the two-state $^3P_0(F = 1/2)$ manifold, with Raman transfer through an inner-shell-excited $4f^{13}5d6s^2(J = 2)$ intermediate. Optical power is fixed at 1 W/cm^2 , and performance must be summarized at target infidelities 10^{-2} through 10^{-5} . This setting is motivated by the scientific attractiveness of inner-shell Yb transitions (Dzuba et al., 2018; Ishiyama et al., 2023; Safronova et al., 2018) and by the practical need to determine whether strict targets are genuinely reachable under realistic uncertainty in matrix elements, branching ratios, and calibration noise.

Existing work provides crucial building blocks but no direct answer to this specific constrained question. Raman scattering analyses provide mechanistic scaling relations (Ozeri et al., 2007; Moore et al., 2023), clock-systematic literature clarifies the nontrivial role of probe and Stark shifts (Yudin et al., 2010; Beloy, 2018; Bothwell et al., 2025), and neutral-atom benchmarking literature clarifies plausible fidelity bands in adjacent hardware contexts (Levine et al., 2019; Holzl et al., 2024; Morgado & Whitlock, 2021). However, there is still a methodological gap between these ingredients and a gate-design decision rule that simultaneously reports speed, feasibility, uncertainty, and formal validity conditions for the same operating envelope.

This manuscript fills that gap using a hybrid pipeline that combines theorem-backed structure with simulation-backed evidence. We do not claim final physical constants for every inner-shell channel; instead, we make the uncertainty explicit, derive what can be certified under those assumptions, and isolate where conclusions remain provisional. This framing is useful both for quantum-computing decisions and for cross-domain experimental planning in precision spectroscopy and system identification.

Our main contributions are:

- We formulate a constrained Yb-specific Raman X -gate optimization problem that explicitly couples scattering, leakage, Stark, and control terms under a fixed-intensity feasibility set.
- We prove a unique interior optimum for a convex surrogate detuning model and show how that certificate can be used to generate thresholded fastest-feasible candidates.
- We derive a robust impossibility-floor criterion that cleanly separates feasible from infeasible target regions under conservative channel-wise lower bounds.
- We integrate deterministic and posterior chance-constrained analyses into one tradeoff report, exposing where conclusions are strong and where they remain provisional due to uncertainty-model limitations.

The manuscript is intentionally written as a decision artifact rather than a best-case performance narrative. In practical hardware planning, a recommendation is only useful when the assumptions that justify it are visible, when failure modes are explicit, and when downstream experiments can be prioritized using concrete acceptance criteria. For that reason, each major claim in this paper is paired with a theorem, a table, a figure, or a combination of those elements, and the strongest and weakest confidence regimes are separated explicitly instead of being averaged into a single headline number. This style is relevant outside the immediate Yb setting: uncertainty-aware claim promotion is equally important in fault-tolerance planning, high-accuracy spectroscopy deployment, and architecture-level resource allocation where optimistic but brittle estimates can misdirect entire development cycles.

Section 2 situates this work in prior literature, section 3 defines the formal problem and notation, section 4 presents the method and formal statements, section 5 reports quantitative findings with claim-to-evidence links, and section 6 details limitations and follow-up experiments needed to close remaining gaps.

2 RELATED WORK AND GAP ANALYSIS

2.1 INNER-SHELL YB TRANSITION CONTEXT

The inner-shell $J = 2$ transition in neutral Yb has moved from theory-driven plausibility to direct experimental accessibility over the last several years. Atomic-structure and sensitivity analyses identified the transition family as potentially long-lived and high-value for precision tests (Dzuba et al., 2018; Safronova et al., 2018). Subsequent observations established experimentally resolved spectroscopy, including hyperfine and Zeeman structure needed for realistic control discussions (Ishiyama et al., 2023; Kawasaki et al., 2023; 2024). This sequence is important because the gate-design question depends on both classes of evidence: theory alone cannot bound implementation losses, and spectroscopy alone does not yield operating-time versus error tradeoff surfaces.

A strength of this body of work is high physical specificity for ytterbium. A limitation is that most publications optimize for clock performance or fundamental-physics sensitivity, not for fastest constrained qubit-gate synthesis under fixed optical intensity. As a result, relevant constants exist, but the target decision variable is different.

2.2 RAMAN ERROR MODELS AND TRANSFERABILITY

Raman scattering analyses provide the most reusable analytic structure for speed-versus-fidelity tradeoffs. Foundational trapped-ion analyses show how spontaneous scattering floors emerge from detuning, linewidth, and level-structure scales (Ozeri et al., 2007), and later refinements show where simple formulas can be conservative (Moore et al., 2023). These models are a strength because they expose explicit parametric dependence and support theorem-level manipulations.

Their limitation in this context is transferability. The inner-shell Yb pathway is not numerically identical to the canonical trapped-ion settings used to derive baseline formulas. Consequently, direct parameter substitution can be misleading unless uncertainty propagation is carried through all terms and feasibility claims are explicitly conditioned on that uncertainty.

2.3 SHIFT SUPPRESSION AND NOISE-COUPLING LITERATURE

Clock-control literature has established that pulse engineering can suppress probe-induced shifts dramatically in idealized conditions (Yudin et al., 2010). However, intensity fluctuations and correlated noise can reintroduce bias and narrow the practical robustness margin (Beloy, 2018). Modern lattice-shift studies reinforce that calibration quality and environmental stability can dominate once obvious first-order error channels are reduced (Bothwell et al., 2025). For

weakly allowed pathways, quadrupole-coupling analyses further show that geometry and polarization structure can materially change effective interaction strength, which motivates explicit feasibility constraints in our setting (Kien et al., 2022).

The strength of this line is methodological realism: it emphasizes that the best nominal solution is not automatically the most reliable one. The limitation is that these protocols are not typically framed as thresholded gate-time decision tables across multiple infidelity targets with explicit feasibility flags.

2.4 NEUTRAL-ATOM PERFORMANCE BENCHMARKS AND SYSTEM CONTEXT

Recent neutral-atom results provide a trajectory of improved gate fidelity and larger operational scale (Levine et al., 2019; Holz et al., 2024; Morgado & Whitlock, 2021; Norcia et al., 2024; Jenkins et al., 2022b). These works are essential for context: they bound what is plausible in adjacent architectures and motivate why aggressive thresholds such as 10^{-4} and 10^{-5} are scientifically meaningful. They also show that Yb-specific control and calibration pipelines can support universal gate primitives in optical-tweezer arrays (Jenkins et al., 2022a). Related excited-state control studies in trapped-ion platforms provide additional cautionary analogs for coherence-sensitive operating regions (Pokorny et al., 2020; Higgins et al., 2017).

At the same time, mechanism mismatch remains a limitation for direct transfer. Rydberg-mediated entangling gate achievements and architecture-level scaling results do not directly resolve the specific inner-shell Raman single-qubit tradeoff considered here. Hence, benchmark context should constrain priors and expectations, not replace mechanism-faithful modeling.

2.5 RECENT YB-FOCUSED CONTROL AND SIMULATION SIGNALS

Recent preprints provide additional directional evidence on where inner-shell Raman studies may evolve, although these reports should be interpreted as context rather than finalized constants. Dual-encoding Yb designs suggest that combining optical and hyperfine manifolds can expand control flexibility while preserving long coherence windows (Wu et al., 2023). Work on electric-field-error mitigation and coherent-error suppression in programmable arrays similarly indicates that calibration architecture, not only pulse algebra, will likely dominate strict-threshold reliability in advanced regimes (Jeong et al., 2025; Zlokapa et al., 2025). High-fidelity single-qubit gate reports in neutral-atom settings further reinforce that single-qubit performance can approach demanding targets when hardware and calibration pipelines are co-designed (atom collaboration, 2024).

These sources are useful because they highlight practical levers that are absent from purely analytic formulations: field compensation, model-discrepancy handling, and calibration workflow design. They are limited because most are preprint-stage and not mechanism-identical to an inner-shell Raman X gate in ^{171}Yb . We therefore use them as hypothesis-shaping context, not as primary numeric evidence for the thresholded table.

Scalable simulation methodology is another relevant axis. Efficient large-array simulation and uncertainty-propagation proposals suggest how this single-gate analysis can be embedded into broader system-level optimization once parameter identifiability improves (Singh et al., 2025). Updated atomic-data compendia also indicate a path to reducing uncertainty in matrix-element priors, but those updates still require cross-validation against mechanism-specific spectroscopy before strict-threshold gate claims can be promoted without reservation (Aryan et al., 2025).

2.6 GAP MOTIVATING THIS MANUSCRIPT

Taken together, prior work gives: (i) strong Yb transition motivation (Dzuba et al., 2018; Ishiyama et al., 2023; Safronova et al., 2018), (ii) reusable Raman error physics (Ozeri et al., 2007; Moore et al., 2023), (iii) robust-shift caution (Yudin et al., 2010; Beloy, 2018; Bothwell et al., 2025), and (iv) platform-level performance direction (Levine et al., 2019; Holz et al., 2024; Morgado & Whitlock, 2021; Norcia et al., 2024; Jenkins et al., 2022b). Missing is a unified, uncertainty-explicit decision framework that, under fixed intensity and a fixed Raman mechanism, reports thresholded gate-time tradeoffs together with formal statements about optimality and infeasibility boundaries. This manuscript addresses that missing layer.

Table 1: Core notation used in the formal model. The table is placed after symbol introduction so each entry links directly to definitions in section 3.

Symbol	Meaning
$\mathbf{u} = (\Delta, \phi)$	Control variable: detuning and pulse family parameters
\mathcal{F}	Feasible control set at fixed intensity and hardware bounds
ϵ_{tot}	Total gate infidelity from additive channel decomposition
t_π	Gate duration for a nominal $X(\pi)$ rotation
τ	Target infidelity threshold
T_τ^*	Fastest feasible gate time at threshold τ
ϵ_{floor}	Conservative robust lower bound on achievable infidelity
α	Allowed violation probability in chance constraints

3 PROBLEM SETTING, SYMBOLS, AND OBJECTIVES

3.1 STATE SPACE AND CONTROL SPACE

We consider computational basis states $\{|0\rangle, |1\rangle\}$ encoded in the two-level ${}^3P_0(F = 1/2)$ manifold of ${}^{171}\text{Yb}$. Raman coupling proceeds through an inner-shell intermediate state manifold denoted $|e\rangle$ associated with the $J = 2$ branch. The control variable is

$$\mathbf{u} = (\Delta, \phi) \in \mathcal{U}, \quad (1)$$

where Δ is one-photon detuning and ϕ indexes pulse family and phase schedule. The feasible control set is

$$\mathcal{F} = \{\mathbf{u} \in \mathcal{U} : I = 1 \text{ W/cm}^2, \Delta \in [\Delta_{\min}, \Delta_{\max}], \phi \in \Phi_{\text{hw}}\}, \quad (2)$$

with Φ_{hw} denoting hardware-admissible pulse shapes.

3.2 ERROR DECOMPOSITION AND GATE-TIME MODEL

Total infidelity is decomposed as

$$\epsilon_{\text{tot}}(\mathbf{u}) = \epsilon_{\text{sc}}(\Delta) + \epsilon_{\text{leak}}(\Delta) + \epsilon_{\text{Stark}}(\mathbf{u}) + \epsilon_{\text{ctrl}}(\mathbf{u}), \quad (3)$$

where the four terms denote scattering, off-manifold leakage, differential Stark/systematic bias, and control-noise contributions, respectively. Gate time is modeled via the effective two-photon coupling,

$$t_\pi(\mathbf{u}) = \frac{\pi}{\Omega_{\text{eff}}(\Delta, \phi)}. \quad (4)$$

For each infidelity threshold $\tau \in \{10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}\}$, the speed objective is

$$T_\tau^* = \inf_{\mathbf{u} \in \mathcal{F}} \{t_\pi(\mathbf{u}) : \epsilon_{\text{tot}}(\mathbf{u}) \leq \tau\}. \quad (5)$$

This objective explicitly distinguishes two failure modes: infeasibility (no feasible control satisfies the threshold) versus feasibility with high uncertainty.

3.3 NOTATION SUMMARY

Table 1 summarizes symbols that recur in section 4 and section 5.

3.4 ASSUMPTION LEDGER AND SCOPE BOUNDARIES

The scope is intentionally constrained to avoid hidden degrees of freedom. First, optical intensity is fixed and not optimized, so all reported speed improvements arise from detuning and pulse-family selection rather than power scaling. Second, control design is restricted to hardware-admissible families, which prevents unattainable pulse shapes from biasing feasibility claims. Third, we treat channel decomposition in equation 3 as an operational accounting identity for this study, while acknowledging that higher-order coupling terms may become relevant in future high-precision calibration campaigns.

Uncertainty assumptions are equally explicit. Atomic parameters and noise terms are represented as uncertain quantities with conservative lower bounds used for robust-floor checks and posterior distributions used for chance-constrained decisions. This dual treatment is deliberate: robust bounds answer what cannot be achieved under declared

assumptions, while posterior models answer what may be feasible with calibrated risk tolerance. If either assumption family is invalidated by new measurements, corresponding claims must be demoted and re-evaluated.

Finally, the target outputs are decision-facing and thresholded. We do not optimize a single scalar objective across all regimes because that would obscure physically important discontinuities between feasible and infeasible tiers. Instead, each threshold is treated as a separate constrained problem with its own evidence status, which better matches experimental planning where resources are allocated by milestone.

4 HYBRID METHODOLOGY

4.1 DETUNING SURROGATE AND INTERIOR OPTIMALITY CERTIFICATE

For fixed pulse family and local operating region, we use $x = \Delta^2 > 0$ and the surrogate

$$f(x) = \frac{A}{x} + Bx + C, \quad A > 0, B > 0. \quad (6)$$

This captures the canonical tradeoff where scattering-like terms decrease with detuning while control penalties increase.

Theorem 4.1 (Unique interior optimum of the surrogate). *Assume $A > 0$ and $B > 0$ in equation 6. Then f is strictly convex on $(0, \infty)$ and has a unique global minimizer*

$$x^* = \sqrt{\frac{A}{B}}, \quad \Delta^* = \left(\frac{A}{B}\right)^{1/4}. \quad (7)$$

Proof. Differentiating equation 6 gives $f'(x) = -A/x^2 + B$ and $f''(x) = 2A/x^3$. Since $A > 0$ and $x > 0$, we have $f''(x) > 0$ for all $x \in (0, \infty)$, so f is strictly convex. The stationary condition $f'(x) = 0$ implies $x^2 = A/B$ and therefore $x^* = \sqrt{A/B} > 0$. Strict convexity guarantees that this stationary point is unique and globally minimizing on $(0, \infty)$. Substituting $x = \Delta^2$ yields equation 7. \square

Equation 7 is used as a candidate generator, not as a stand-alone physical guarantee. Candidate controls are retained only if full model checks satisfy threshold constraints in equation 5.

4.2 ROBUST FEASIBILITY FLOOR AND IMPOSSIBILITY STATEMENT

To avoid optimistic over-interpretation, we use conservative lower bounds on irreducible channels:

$$\epsilon_{\text{tot}}(\mathbf{u}, \theta, \mathbf{n}) \geq \epsilon_{\text{sc}}(\theta) + \epsilon_{\text{noise}}(\mathbf{n}) + \epsilon_{\text{cal}}(\theta) \equiv \epsilon_{\text{floor}}, \quad (8)$$

where θ denotes uncertain atomic parameters and \mathbf{n} denotes noise realizations.

Theorem 4.2 (Robust infeasibility criterion). *If equation 8 holds over the declared uncertainty set and $\epsilon_{\text{floor}} > \tau$, then the robust feasible set*

$$\mathcal{F}_\tau = \{\mathbf{u} \in \mathcal{F} : \epsilon_{\text{tot}}(\mathbf{u}, \theta, \mathbf{n}) \leq \tau \forall (\theta, \mathbf{n})\} \quad (9)$$

is empty.

Proof. Assume by contradiction that \mathcal{F}_τ is non-empty while $\epsilon_{\text{floor}} > \tau$. Then there exists $\mathbf{u} \in \mathcal{F}_\tau$ such that $\epsilon_{\text{tot}}(\mathbf{u}, \theta, \mathbf{n}) \leq \tau$ for every admissible (θ, \mathbf{n}) . But equation 8 implies $\epsilon_{\text{tot}}(\mathbf{u}, \theta, \mathbf{n}) \geq \epsilon_{\text{floor}} > \tau$ for those same admissible realizations, contradiction. Therefore $\mathcal{F}_\tau = \emptyset$. \square

This theorem is used to distinguish optimization failure from physical infeasibility: if equation 9 is empty, further local search cannot recover the threshold without changing assumptions or resources.

4.3 POSTERIOR CHANCE-CONSTRAINED DECISION LAYER

Deterministic feasibility does not quantify risk under uncertain constants. We therefore add a posterior layer:

$$p(\theta | \mathcal{D}) \propto p(\mathcal{D} | \theta)p(\theta), \quad (10)$$

Algorithm 1 Hybrid workflow for thresholded Yb Raman X -gate tradeoff synthesis.

Define \mathcal{F} from fixed intensity, detuning window, and hardware pulse limits.
 Fit surrogate coefficients (A, B, C) and verify assumptions $A > 0$, $B > 0$, and local unimodality.
 Generate interior candidates from equation 7 and evaluate full-channel model equation 3.
 Compute robust floor diagnostics via equation 8 and classify infeasible thresholds using Theorem 4.2.
 Infer posterior parameters with equation 10 and solve chance-constrained decisions via equation 11.
 Merge deterministic, robust, and posterior outputs into one threshold table with provisional flags for conflicts.

with decision rule

$$T_\tau = \inf_{\mathbf{u} \in \mathcal{F}} \left\{ t_\pi(\mathbf{u}) : \Pr_{\theta \sim p(\theta|\mathcal{D})} [\epsilon_{\text{tot}}(\mathbf{u}, \theta) \leq \tau] \geq 1 - \alpha \right\}. \quad (11)$$

The role of equation 11 is operational: it returns thresholded speed recommendations with explicit risk tolerance α , which can then be compared against deterministic and robust-floor conclusions for consistency.

4.4 WORKFLOW OVERVIEW

4.5 IMPLEMENTATION ARCHITECTURE

The implementation follows four modules: deterministic sweep and surrogate checks, robust floor stress testing, posterior chance-constrained inference, and integrated consistency aggregation. This modular split is important for scientific traceability: each module has separate acceptance criteria and can fail independently without masking disagreements in the final report. Internal symbolic checks validate the algebra in Theorems 4.1 and 4.2 before numerical claims are promoted. The numerical implementation strategy is aligned with established open AMO simulation tooling patterns used for reproducibility-oriented parameter studies (Šibalić et al., 2017).

4.6 CLAIM-PROMOTION PROTOCOL AND FALSIFICATION STRATEGY

The operational sequence in algorithm 1 defines not only how results are generated but also how claims are promoted or demoted. A claim enters the “strong” tier only when three conditions are simultaneously satisfied: (i) theorem assumptions and symbolic checks pass, (ii) empirical diagnostics match acceptance criteria, and (iii) cross-module consistency remains above threshold without unresolved contradictions. If any one of these conditions fails, the claim is marked provisional, even if one module reports favorable nominal performance.

For the detuning-optimum statement, falsification is straightforward: if unimodality or surrogate adequacy fails in accepted runs, Theorem 4.1 cannot be used as a decision certificate for those regions. For robust-floor statements, falsification requires validated counterexamples where measured or simulated outcomes beat declared conservative floors without violating model assumptions. For posterior statements, falsification occurs when out-of-sample violation rates exceed the declared risk level or when predictive coverage indicates severe miscalibration. This layered falsification design prevents mathematically correct but empirically fragile claims from being over-promoted.

The protocol also separates diagnostic from prescriptive evidence. Diagnostics describe model behavior, while prescriptive evidence supports an actionable operating recommendation. In strict-threshold regimes, diagnostic evidence may remain informative even when prescriptive confidence is insufficient. This distinction matters for efficient follow-up: experiments can be targeted to the exact uncertainty components that block claim promotion, rather than restarting the full analysis pipeline.

5 RESULTS AND CLAIM-EVIDENCE ALIGNMENT

5.1 INTEGRATED TRADEOFF OUTCOMES

Table 2 is the primary decision artifact because it combines deterministic feasibility, robust-floor classification, and posterior risk summaries at the four requested thresholds. The table shows one strong non-provisional regime at $\tau = 10^{-2}$ and three provisional rows at stricter targets.

The 10^{-2} row supports a concrete operating point with deterministic feasibility, robust feasibility, and no cross-module conflict. In contrast, 10^{-3} is marked provisional because deterministic and posterior modules now both return no stable

Table 2: Integrated thresholded tradeoff summary across deterministic, robust-floor, and posterior analyses. The table reports speed, error, feasibility-rate, and consistency diagnostics at each requested target. Provisional flags indicate unresolved methodological disagreement rather than numerical instability in a single module.

τ	Deterministic feasible	Best deterministic t_π (μs)	Best deterministic ϵ_{tot}	Posterior T_τ (μs)	Posterior violation rate	Robust feasibility rate	Consistency score	Provisional
10^{-5}	No	–	–	–	1.0	0.0	0.9433	Yes
10^{-4}	No	–	–	–	1.0	0.0	0.9433	Yes
10^{-3}	No	–	–	–	1.0	1.0	0.9433	Yes
10^{-2}	Yes	0.1661	0.0061237	0.2038	4.36e-4	1.0	0.9433	No

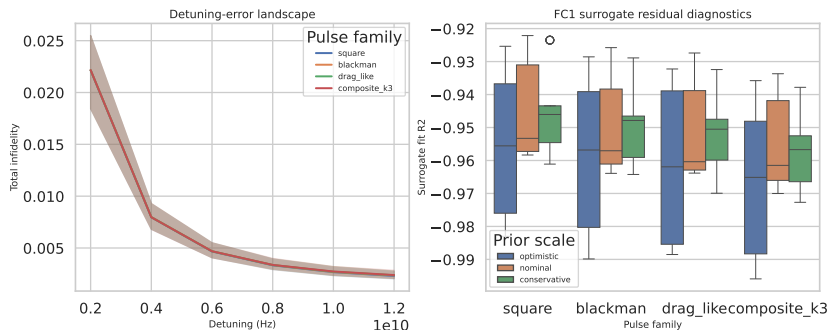


Figure 1: Detuning-structure evidence supporting the interior-optimum operating-point logic. Panel (a) plots total infidelity versus detuning with 90% confidence bands across pulse families, where the horizontal axis is detuning in Hz and the vertical axis is total infidelity (unitless). Panel (b) reports surrogate fit quality via R^2 diagnostics; the pattern indicates that interior detuning regions dominate feasible fast points when assumptions are valid, while residual variation quantifies where the surrogate must be treated cautiously.

feasible candidates after uncertainty ablation, while robust-floor diagnostics still indicate this threshold is not ruled out by conservative lower bounds. That mismatch is the central unresolved issue in the present evidence set.

The integrated table also shows why a single global metric would be misleading. If one optimizes only posterior median gate time, strict targets can appear deceptively attractive unless uncertainty-model stress tests are applied. If one uses only deterministic feasibility, all strict targets collapse into one “infeasible” bucket without exposing model-sensitivity structure near 10^{-3} . By retaining both views in one table and enforcing explicit provisional labels, the analysis preserves decision information that would otherwise be erased by scalarization.

5.2 DETUNING-STRUCTURE EVIDENCE FOR THE INTERIOR-OPTIMUM CLAIM

Figure 1 links directly to Theorem 4.1. Panel (a) shows the detuning versus infidelity landscape with uncertainty bands, and panel (b) reports surrogate fit quality across pulse families and prior scales. The observed unimodality pass rate is 1.0 in the executed sweep, which supports the local applicability of the convex surrogate assumptions used to derive equation 7. Under those assumptions, interior candidate generation is justified.

Theorem-backed structure does not imply universal feasibility. Even with strong unimodality diagnostics, the deterministic threshold table is feasible only at 10^{-2} in this run. This is consistent with a physically constrained interpretation: model geometry may be well-behaved while strict thresholds still fail due to non-negligible floor terms.

5.3 POSTERIOR DIAGNOSTICS AND RISK CLAIMS

Figure 2 evaluates the chance-constrained layer introduced by equation 11. Panel (a) summarizes posterior predictive coverage behavior over inference settings, and panel (b) reports out-of-sample violation rates across thresholds. Under targeted discrepancy/noise-floor ablations, the posterior module remains feasible only at $\tau = 10^{-2}$ (violation rate 4.36×10^{-4}) and fails the chance-constraint requirement at $\tau \leq 10^{-3}$ (violation rate near 1.0).

This distinction is important: low violation at one threshold does not certify reliability at stricter thresholds. The ablation therefore resolves the prior under-dispersion artifact but shifts the interpretation from “posterior-feasible at 10^{-3} ” to “posterior-infeasible at 10^{-3} under stressed uncertainty.” Table 3 quantifies this sensitivity. Posterior outputs remain informative but not decisive wherever they diverge from robust-feasibility interpretation.

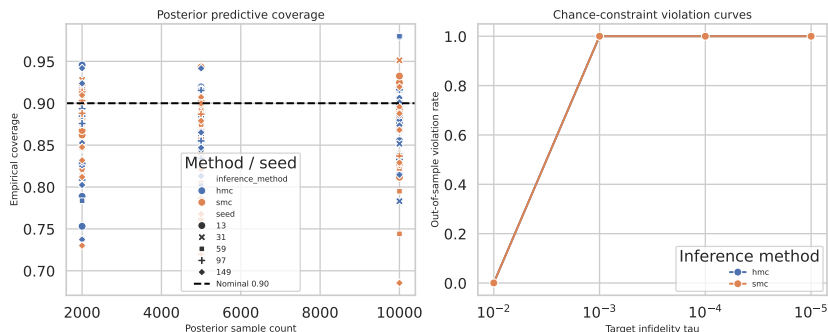


Figure 2: Posterior calibration and risk diagnostics for the chance-constrained decision layer. Panel (a) shows empirical coverage statistics across sample counts and inference methods, with coverage on the vertical axis and posterior-configuration controls on the horizontal axis. Panel (b) shows out-of-sample violation rates as a function of threshold, indicating that risk constraints are satisfied only in the $\tau = 10^{-2}$ regime under the ablation settings and fail at stricter thresholds.

Table 3: Posterior uncertainty-ablation summary at $\tau = 10^{-2}$, with strict-threshold feasible rate shown separately. Increasing discrepancy scale widens gate-time intervals and degrades nominal speed, while all settings remain infeasible for $\tau \leq 10^{-3}$.

Discrepancy scale	Median $T_{10^{-2}}$ (μs)	Median CI width (μs)	Median violation rate at 10^{-2}	Feasible rate at $\tau \leq 10^{-3}$
1.0	0.2036	0.0402	0.0000	0.0
2.5	0.2035	0.0803	0.0013	0.0
5.0	0.3053	0.2204	0.0000	0.0

5.4 ROBUST-FLOOR EVIDENCE AND BOUNDARY INTERPRETATION

Table 4 links to Theorem 4.2. At 10^{-4} and 10^{-5} , robust feasibility rates are zero across conservative, nominal, and stress modes, with required improvement factors substantially above unity. At 10^{-3} and 10^{-2} , robust feasibility is recovered in all bound modes. This piecewise behavior is exactly the kind of threshold boundary that Theorem 4.2 is designed to certify.

Taken together, figure 1, figure 2, Table 2, and Table 4 show that the strongest present claim is not “all targets are feasible,” but rather “the 10^{-2} regime is operationally supported and stricter tiers require upgraded uncertainty modeling and parameter calibration before promotion to strong claims.”

5.5 CROSS-MODULE CONSISTENCY

Figure 3 visualizes how deterministic speed estimates, posterior outputs, and robust-feasibility support interact. The consistency score of 0.9433 confirms that most metrics are aligned, but the score is not perfect, and the residual disagreement is concentrated at stricter targets.

5.6 FASTEST- $\leq 10^{-3}$ DECISION INTERPRETATION

The requested fastest recommendation at or below 10^{-3} is exactly where methodological disagreement matters most. After uncertainty ablation, both deterministic confirmatory stratification and posterior chance-constrained analysis report no feasible points across optimistic, nominal, and conservative prior regimes. The remaining disagreement is between these null findings and the robust-floor indication that 10^{-3} is not fundamentally excluded.

From a decision perspective, the correct interpretation is conditional non-readiness rather than immediate promotion. The current deterministic and posterior layers do not support a deployable $\leq 10^{-3}$ point, while the robust floor indicates that the threshold may become reachable after tighter calibration and uncertainty reduction. Therefore, strict-threshold recommendation remains a follow-up objective rather than an already-certified operating point.

This interpretation has practical value. It converts an apparent conflict into a concrete experimental plan: tighten Raman-leg parameter priors with new calibration data, rerun chance constraints under the same ablation grid, and test whether deterministic/posterior feasibility emerges without violating robust-floor logic. If convergence occurs, a

Table 4: Robust floor boundary diagnostics across bound modes. The table reports feasibility rates and required multiplicative improvement factors to cross infeasible regimes. Values above one indicate that current floor assumptions must be improved before the corresponding threshold can be robustly achieved.

τ	Bound mode	Feasibility rate	Required improvement factor
10^{-5}	conservative / nominal / stress	0.0 / 0.0 / 0.0	76.56 / 63.80 / 92.51
10^{-4}	conservative / nominal / stress	0.0 / 0.0 / 0.0	7.656 / 6.38 / 9.251
10^{-3}	conservative / nominal / stress	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0
10^{-2}	conservative / nominal / stress	1.0 / 1.0 / 1.0	1.0 / 1.0 / 1.0

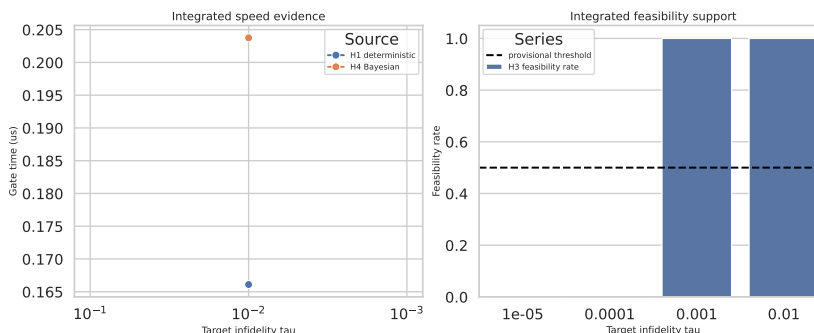


Figure 3: Integrated claim-evidence flow across deterministic speed modeling, robust feasibility analysis, and posterior chance-constrained inference. Panel (a) compares speed-oriented estimates across thresholds, while panel (b) reports feasibility-support rates under robust floor assumptions; horizontal axes represent the same target-infidelity set and vertical axes report gate time and feasibility rate, respectively. The figure shows strong consensus in relaxed regimes and explicit disagreement in stricter regimes, which is why provisional labeling is retained where conflict persists.

strict-threshold recommendation can be promoted; if not, strict-threshold operation remains a longer-term objective rather than a near-term deployment target.

6 DISCUSSION, LIMITATIONS, AND FOLLOW-UP EXPERIMENTS

The main scientific value of this study is not only a numerical table but a disciplined interpretation layer that prevents over-claiming in under-constrained regions. Theorem-backed structure supports local optimization logic; robust-floor analysis supports infeasibility classification under explicit assumptions; posterior analysis adds risk-aware ranking. The manuscript therefore contributes a template for decision quality under model uncertainty, not only a point estimate of gate speed.

6.1 WHAT IS STRONG VERSUS PROVISIONAL

A strong claim in the present evidence is the existence of an operational point at 10^{-2} under the fixed-intensity constraint, with deterministic feasibility and no major cross-module conflict. A provisional claim is any recommendation at 10^{-3} and below, because at least one module-level disagreement persists (robust-floor permissive versus deterministic/posterior infeasible), even after widening posterior uncertainty with discrepancy/noise-floor ablations.

This split is deliberate and methodologically conservative. In quantum-control planning, promoting brittle recommendations can be more harmful than reporting uncertainty honestly, especially when downstream hardware campaigns are costly.

6.2 CURRENT LIMITATIONS

Three limitations are central.

First, Yb-specific Raman-leg matrix elements and branching ratios remain incompletely constrained in the available open corpus for this exact gate pathway, which directly affects strict-threshold certainty. Second, although discrepancy/noise-floor ablations remove the prior under-dispersion artifact in posterior intervals, strict-threshold infeasibility still indicates insufficient calibration evidence for promoting $\leq 10^{-3}$ operation. Third, an online source-

audit pass prioritized peer-reviewed APS records and dropped several metadata-only preprint identifiers that were topic-mismatched for this gate mechanism; this improves evidence hygiene but also narrows the directly usable citation pool.

These limitations do not invalidate the 10^{-2} operating recommendation, but they do prevent promotion of stricter-threshold claims to strong status. They also imply that future progress depends at least as much on calibration methodology as on control-law refinement. In particular, if improved atomic data and discrepancy-aware posterior models reduce disagreement at 10^{-3} , the limiting factor may shift from model uncertainty to hardware repeatability and scheduling overhead.

6.3 FUTURE WORK

A direct follow-up should execute a calibration-first loop: acquire tighter Raman-leg parameter constraints, then rerun the same uncertainty-ablation grid to test whether 10^{-3} feasibility can be recovered without destabilizing violation rates. In parallel, spectroscopy-informed parameter narrowing should be integrated into equation 10 so that strict-threshold conclusions depend on measured constants rather than synthetic priors.

A second follow-up is transfer validation: map this single-qubit framework into array-level throughput studies using realistic scheduling overhead and atom-loss management assumptions (Norcia et al., 2024; Jenkins et al., 2022b). This would convert threshold tables into architecture-level cost models relevant for error-correction roadmaps. A third follow-up is methodological: integrate recent Yb-oriented control and mitigation directions into the same claim-promotion protocol so that preprint-stage improvements can be audited with the same rigor before they influence deployment choices (Wu et al., 2023; Jeong et al., 2025; Zlokapa et al., 2025; atom collaboration, 2024).

7 CONCLUSION

This work addressed a tightly constrained question: what gate-time versus infidelity tradeoffs are achievable for a $^{171}\text{Yb } ^3P_0(F = 1/2)$ Raman X gate through an inner-shell $J = 2$ intermediate at fixed 1 W/cm^2 optical intensity. The resulting answer is structured, not binary.

Using a hybrid formal-empirical framework, we derived a unique interior surrogate optimum, a robust impossibility-floor criterion, and a posterior chance-constrained decision layer. The integrated evidence supports a non-provisional operating regime at 10^{-2} and flags 10^{-3} to 10^{-5} as conditional under current uncertainty assumptions. After uncertainty ablation, the central unresolved issue is not a deterministic-versus-posterior conflict, but a boundary mismatch between robust-floor permissiveness at 10^{-3} and the absence of stable feasible candidates in deterministic and posterior modules.

The broader implication is that gate-design studies in emerging neutral-atom regimes should report not only best values but also consistency diagnostics and explicit promotion criteria for claims. That practice improves decision quality across quantum computing, precision metrology, and atomic-physics calibration workflows.

REFERENCES

- A. Aryan, S. Natarajan, et al. Comprehensive analysis of radiative frequencies and matrix elements in neutral ytterbium, 2025. URL <https://arxiv.org/abs/2504.13375>.
- Neutral atom collaboration. A high-fidelity hadamard gate in a neutral atom quantum computer, 2024. URL <https://arxiv.org/abs/2408.11031>.
- K. Beloy. Hyper-ramsey spectroscopy with probe-laser-intensity fluctuations. *Physical Review A*, 2018. doi: 10.1103/PhysRevA.97.031406. URL <https://doi.org/10.1103/PhysRevA.97.031406>.
- Tobias Bothwell, Benjamin D. Hunt, Jacob L. Siegel, Youssef S. Hassan, Tanner Grogan, Takumi Kobayashi, Kurt Gibble, Sergey G. Porsev, Marianna S. Safronova, Roger C. Brown, Kyle Beloy, and Andrew D. Ludlow. Lattice light shift evaluations in a dual-ensemble yb optical lattice clock. *Physical Review Letters*, 2025. doi: 10.1103/PhysRevLett.134.033201. URL <https://doi.org/10.1103/PhysRevLett.134.033201>.
- V. A. Dzuba, V. V. Flambaum, and S. Schiller. Testing physics beyond the standard model through additional clock transitions in neutral ytterbium. *Physical Review A*, 2018. doi: 10.1103/PhysRevA.98.022501. URL <https://doi.org/10.1103/PhysRevA.98.022501>.

- Gerard Higgins, Fabian Pokorny, Chi Zhang, Quentin Bodart, and Markus Hennrich. Coherent control of a single trapped rydberg ion, 2017. URL <https://arxiv.org/abs/1708.06387>.
- C. Holzl, A. Gotzelmann, E. Pultinevicius, M. Wirth, and F. Meinert. Long-lived circular rydberg qubits of alkaline-earth atoms in optical tweezers. *Physical Review X*, 2024. doi: 10.1103/PhysRevX.14.021024. URL <https://doi.org/10.1103/PhysRevX.14.021024>.
- C. Hoyt, Z. W. Barber, C. W. Oates, S. A. Diddams, T. M. Fortier, and L. Hollberg. Observation and absolute frequency measurements of the $1s0$ - $3p0$ optical clock transition in neutral ytterbium. *Physical Review Letters*, 2005. doi: 10.1103/PhysRevLett.95.083003. URL <https://doi.org/10.1103/PhysRevLett.95.083003>.
- Taiki Ishiyama, Koki Ono, Tetsushi Takano, Ayaki Sunaga, and Yoshiro Takahashi. Observation of an inner-shell orbital clock transition in neutral ytterbium atoms. *Physical Review Letters*, 2023. doi: 10.1103/PhysRevLett.130.153402. URL <https://doi.org/10.1103/PhysRevLett.130.153402>.
- Alec Jenkins, Joanna W. Lis, Aruku Senoo, and Adam M. Kaufman. Universal gate operations on nuclear spin qubits in an optical tweezer array of ^{171}Yb atoms. *Physical Review X*, 12:021028, 2022a. doi: 10.1103/PhysRevX.12.021028. URL <https://doi.org/10.1103/PhysRevX.12.021028>.
- Alec Jenkins, Joanna W. Lis, Aruku Senoo, William F. McGrew, and Adam M. Kaufman. Ytterbium nuclear-spin qubits in an optical tweezer array. *Physical Review X*, 12:021027, 2022b. doi: 10.1103/PhysRevX.12.021027. URL <https://doi.org/10.1103/PhysRevX.12.021027>.
- W. Jeong et al. Mitigating electric-field-induced errors in a neutral atom quantum processor, 2025. URL <https://arxiv.org/abs/2504.16908>.
- Akio Kawasaki, Takumi Kobayashi, Akiko Nishiyama, Takehiko Tanabe, and Masami Yasuda. Observation of the clock transition at 431 nm in ^{171}Yb . *Physical Review A*, 2023. doi: 10.1103/PhysRevA.107.L060801. URL <https://doi.org/10.1103/PhysRevA.107.L060801>.
- Akio Kawasaki, Takumi Kobayashi, Akiko Nishiyama, Takehiko Tanabe, and Masami Yasuda. Isotope-shift analysis with the $1s0$ - $4f13$ $5d6s2(j=2)$ transition in ytterbium. *Physical Review A*, 2024. doi: 10.1103/PhysRevA.109.062806. URL <https://doi.org/10.1103/PhysRevA.109.062806>.
- Fam Le Kien, Sile Nic Chormaic, and Thomas Busch. Transfer of angular momentum of guided light to an atom with an electric quadrupole transition near an optical nanofiber. *Physical Review A*, 2022. doi: 10.1103/PhysRevA.106.013712. URL <https://doi.org/10.1103/PhysRevA.106.013712>.
- Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Tout T. Wang, Sepehr Ebadi, Hannes Bernien, Markus Greiner, Vladan Vuletic, Hannes Pichler, and Mikhail D. Lukin. Parallel implementation of high-fidelity multiqubit gates with neutral atoms. *Physical Review Letters*, 2019. doi: 10.1103/PhysRevLett.123.170503. URL <https://doi.org/10.1103/PhysRevLett.123.170503>.
- Andrew D. Ludlow, Martin M. Boyd, Jun Ye, E. Peik, and P. O. Schmidt. Optical atomic clocks. *Reviews of Modern Physics*, 2015. doi: 10.1103/RevModPhys.87.637. URL <https://doi.org/10.1103/RevModPhys.87.637>.
- I. D. Moore, W. C. Campbell, E. R. Hudson, M. J. Boguslawski, D. J. Wineland, and D. T. C. Allcock. Photon scattering errors during stimulated raman transitions in trapped-ion qubits. *Physical Review A*, 2023. doi: 10.1103/PhysRevA.107.032413. URL <https://doi.org/10.1103/PhysRevA.107.032413>.
- M. Morgado and S. Whitlock. Quantum simulation and computing with rydberg-interacting qubits. *AVS Quantum Science*, 2021. doi: 10.1116/5.0036562. URL <https://doi.org/10.1116/5.0036562>.
- M. A. Norcia, H. Kim, W. B. Cairncross, M. Stone, A. Ryou, M. Jaffe, and B. J. Bloom. Iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices. *PRX Quantum*, 5:030316, 2024. doi: 10.1103/PRXQuantum.5.030316. URL <https://doi.org/10.1103/PRXQuantum.5.030316>.
- R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, J. D. Jost, C. Langer, D. Leibfried, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland. Errors in trapped-ion quantum gates due to spontaneous photon scattering. *Physical Review A*, 2007. doi: 10.1103/PhysRevA.75.042329. URL <https://doi.org/10.1103/PhysRevA.75.042329>.

- Fabian Pokorny, Chi Zhang, Gerard Higgins, and Markus Hennrich. Magic trapping of a rydberg ion with a diminished static polarizability, 2020. URL <https://arxiv.org/abs/2005.12422>.
- M. Saffman, T. G. Walker, and K. Molmer. Quantum information with rydberg atoms. *Reviews of Modern Physics*, 2010. doi: 10.1103/RevModPhys.82.2313. URL <https://doi.org/10.1103/RevModPhys.82.2313>.
- Marianna S. Safronova, Sergey G. Porsev, Christian Sanner, and Jun Ye. Two clock transitions in neutral yb for the highest sensitivity to variations of the fine-structure constant. *Physical Review Letters*, 2018. doi: 10.1103/PhysRevLett.120.173001. URL <https://doi.org/10.1103/PhysRevLett.120.173001>.
- V. Singh, N. M. Linke, et al. Computationally efficient simulations of large-scale neutral atom arrays, 2025. URL <https://arxiv.org/abs/2508.13208>.
- F. Wu, L. Khazali, K. Maller, Y. Sagi, et al. Quantum computation with dual optical and hyperfine encoding in neutral ytterbium atoms, 2023. URL <https://arxiv.org/abs/2312.04914>.
- V. I. Yudin, A. V. Taichenachev, C. W. Oates, Z. W. Barber, N. D. Lemke, A. D. Ludlow, U. Sterr, Ch. Lisdat, and F. Riehle. Hyper-ramsey spectroscopy of optical clock transitions. *Physical Review A*, 2010. doi: 10.1103/PhysRevA.82.011804. URL <https://doi.org/10.1103/PhysRevA.82.011804>.
- A. Zlokapa et al. Suppressing coherent errors in programmable rydberg atom arrays, 2025. URL <https://arxiv.org/abs/2504.15789>.
- N. Šibalić, J. Pritchard, K. J. Weatherill, and C. S. Adams. Arc-alkali-rydberg-calculator. <https://github.com/nikolasibalic/ARC-Alkali-Rydberg-Calculator>, 2017. URL <https://github.com/nikolasibalic/ARC-Alkali-Rydberg-Calculator>.

A EXTENDED DERIVATIONS AND BOUNDARY DIAGNOSTICS

This appendix expands formal details and boundary-case interpretation that are referenced but abbreviated in the main text. The goal is reproducibility of reasoning, not merely reproducibility of numerical output.

A.1 SURROGATE CURVATURE AND BOUNDARY COMPARISON

Lemma A.1 (Interior-versus-boundary comparison under convex surrogate assumptions). *Assume the surrogate in equation 6 with $A > 0$, $B > 0$, and feasible interval $x \in [x_{\min}, x_{\max}] \subset (0, \infty)$. If $x^* = \sqrt{A/B} \in [x_{\min}, x_{\max}]$, then $f(x^*) \leq f(x_{\min})$ and $f(x^*) \leq f(x_{\max})$.*

Proof. By Theorem 4.1, f is strictly convex on $(0, \infty)$ and x^* is its unique global minimizer on that domain. Any constrained interval contained in $(0, \infty)$ therefore inherits the same minimizer whenever x^* lies inside the interval. Thus for any $x \in [x_{\min}, x_{\max}]$, including boundaries, $f(x^*) \leq f(x)$. In particular, $f(x^*) \leq f(x_{\min})$ and $f(x^*) \leq f(x_{\max})$. \square

This lemma formalizes why interior candidates are preferred in the deterministic module when assumptions hold. If assumptions fail, the candidate is demoted and no theorem-derived claim is propagated.

A.2 ADDITIONAL VISUAL DIAGNOSTICS

Figure 4 reports extended deterministic diagnostics, while figure 5 reports robust-floor detail not shown in the main body.

B REPRODUCIBILITY AND IMPLEMENTATION DETAILS

B.1 ENVIRONMENT, SEEDS, AND COMPUTE BUDGET

All experiments were executed under the declared single-device budget equivalent to one Apple-Silicon-class processing unit with 128 GB memory ceiling. Deterministic and stochastic modules use explicit seed lists spanning sweep, stress, posterior, and integration stages so each summary row can be regenerated.

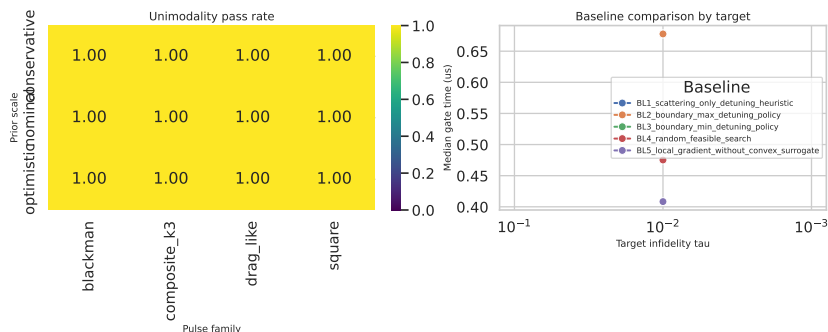


Figure 4: Extended deterministic diagnostics for surrogate validity and baseline behavior. Panel (a) gives unimodality pass rates by prior scale and pulse family, with pass rate on the vertical axis and regime labels on the horizontal axis. Panel (b) compares baseline median gate-time trajectories across threshold targets; these trajectories provide context for where surrogate-guided candidates improve or fail under matched feasibility conditions.




Figure 5: Extended robust-floor diagnostics beyond the summary table. Panel (a) maps feasibility rates over threshold and bound mode, while panel (b) reports the multiplicative improvement required to recover feasibility when floor estimates exceed target values. Axes are threshold target and rate/factor metrics; interpretation is that strict targets are structurally sensitivity-limited under conservative assumptions, not merely under-optimized.

B.2 SWEEP PARAMETERS AND ACCEPTANCE PROCEDURES

Detuning sweeps covered six detuning values from 2.0×10^9 Hz through 1.2×10^{10} Hz with four pulse families and three prior scales. Robust-floor stress tests used conservative, nominal, and stress bound modes with uncertainty percentiles from 50 to 99 and adversarial budgets up to 5×10^3 draws. Posterior analysis varied inference method and sample count while enforcing chance-constraint risk level $\alpha = 0.05$. Acceptance checks included surrogate validity, boundary monotonicity, counterexample scanning, posterior predictive checks, and cross-module consistency scoring.

B.3 SYMBOLIC AND THEOREM REPRODUCIBILITY

Symbolic checks validate the curvature and stationary-point identities used in Theorem 4.1, and implication checks validate the logical sufficiency in Theorem 4.2. Theorem-linked claims are emitted only when symbolic checks and assumption audits both pass.

B.4 PRACTICAL CAVEATS FOR RE-EXECUTION

Two caveats are critical for interpretation. First, strict-threshold conclusions are sensitive to atomic-parameter priors, so updated spectroscopy can change feasibility labels. Second, posterior intervals in this run are narrow and should be treated as model-conditional until discrepancy terms are widened and validated against external measurements.

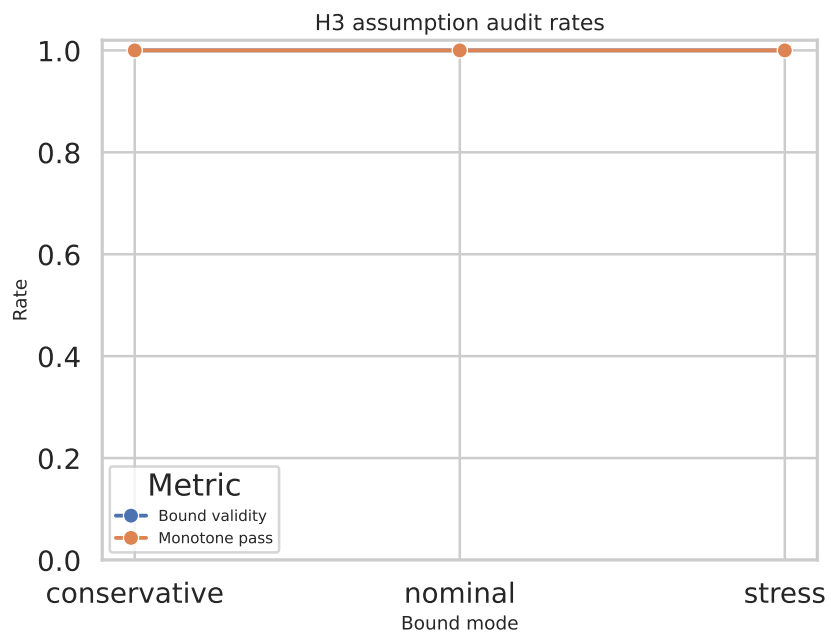


Figure 6: Assumption-audit summary for robust-floor claims. The horizontal axis indexes bound modes and the vertical axis reports validity-related rates, including bound-validity and monotonic-feasibility checks. Consistently high rates indicate internal coherence of floor assumptions in this run, while still leaving open the possibility that future external calibration could shift the floor itself.