

Metrology-oriented bulk-optics implementation of a stochastic NOT gate: Unified error budget and tolerance analysis

Oleg Angelsky^{1,2} , Myroslav Strynadko^{2*} , Claudia Zenkova² , B. Tymochko³ , R. Zaiats² ,
Xinzheng Zhang⁴ , Jun Zheng^{1**} , Jingxian Cai¹

¹Taizhou Institute of Zhejiang University, Taizhou, China

²Yuriy Fedkovych Chernivtsi National University, Chernivtsi, Ukraine

³Bukovinian State Medical University, Chernivtsi, Ukraine

⁴Nankai University, Tianjin, China

Article info

Article history:

Received 07 Mar 2026

Accepted 13 May 2026

Available on-line 10 Sep. 2026

Keywords:

stochastic computing;

polarisation encoding;

logical negation;

error budget;

sensitivity analysis;

calibration protocol.

Abstract

Stochastic computing represents real-valued quantities by the relative frequency of ones in a binary pulse train, enabling compact arithmetic with an accuracy that improves with stream length. Here, we report a metrology-oriented, bulk optics realisation of a stochastic NOT primitive for a polarisation-encoded bit stream ($1 \equiv H$, $0 \equiv V$), implemented as a passive exchange of two orthogonal linear polarisations. Two practical architectures are considered: (A) a folded, single-output polarising beam splitter (PBS)-based module with two PBS cubes and two half-wave plates (HWPs) and (B) a compact $3 \times$ HWP in-line chain. Using a unified Jones calculus model coupled to stochastic bit stream simulation, we quantify how finite extinction/leakage, waveplate angle, and retardance errors, arm imbalance, detector noise, and contrast threshold decision settings propagate into computation level metrics ($\hat{p}$ transfer, bias/RMSE, BER, and contrast margin). A harmonised head-to-head comparison under common decision settings ($\gamma = 0.60$, erasure window 0.35) reveals distinct sensitivity profiles for the two layouts and provides actionable tolerance guidance for teaching grade implementations and for calibrated error budget studies. The full reproducible Colab workflows and derived datasets are provided as Supplementary materials, together with an alignment/calibration protocol and a phase reference/compensation note for phase sensitive use cases.

1. Introduction

Stochastic computing represents real-valued quantities by the relative frequency of ones in a binary pulse train, enabling compact arithmetic with an accuracy that improves with stream length [1–3]. A practical bottleneck, however, is not ideal probability algebra but translating component non-idealities and alignment tolerances into computation-level errors such as bias/RMSE, bit error rate (BER), and decision robustness [1]. A metrology-oriented error budget is therefore essential for predictable cascading and for teaching grade laboratory implementations.

Polarisation encoding provides a natural physical realisation of binary states, using two orthogonal linear polarisations (H and V) as logic values. Polarisation-based

optical logic operations and gate-like functions have been demonstrated in free-space configurations [4], while the Jones-calculus description and standard H/V analysis with waveplates and polarising beam splitters provide the optical basis for controllable state transformations [5–8].

Logical negation (NOT) is a fundamental primitive because it performs probability inversion, $p_{\text{out}} = 1 - p_{\text{in}}$ and is routinely required for normalisation, complement-based arithmetic and feedback in probabilistic circuits [1]. While a polarisation swap $H \leftrightarrow V$ is conceptually simple, a practical bulk optics NOT becomes valuable only when (i) its behaviour is stated in experimentally observable metrics derived from H/V analysis and thresholding and (ii) sensitivity to realistic imperfections (finite extinction/leakage, waveplate angle/retardance errors, imbalance and detector noise) is quantified in a reproducible workflow.

Corresponding author at: *m.strinadko@chnu.edu.ua, **dbzj@netease.com

Related probabilistic hardware paradigms further motivate well characterised stochastic primitives. The probabilistic bit (p-bit) concept formalises a tunable stochastic binary element positioned between deterministic bits and quantum bits, enabling networks for optimisation and inference [9, 10]. Coherence-aware and correlation-optics perspectives are also relevant whenever phase-referenced diagnostics or coherent recombination are used to validate polarisation transformations, which motivates our dedicated phase-reference and compensation note (Supplement S6) [11–14].

Gap addressed in this work. Despite substantial literature on polarisation-encoded optical logic and on stochastic/probabilistic computing, a practical gap remains: a bulk optics stochastic bit stream NOT primitive accompanied by an explicit, computation-level error budget under realistic component tolerances. Existing optical demonstrations often focus on deterministic truth table behaviour or switching functionality, while stochastic computing works typically do not provide a component level optical uncertainty budget that can be used to predict cascaded behaviour [1, 3–5].

In this work, we report a bulk optics stochastic NOT gate for a polarisation-encoded bit stream ($1 \equiv H$, $0 \equiv V$) and provide a unified modelling workflow that maps optical non-idealities to computation-level performance. Two practical realisations are considered: (A) a folded, single output PBS-based exchange module and (B) a compact in-line $3 \times$ half-wave plate (HWP) chain. Using Jones calculus modelling coupled to stochastic bit stream simulation, we quantify transfer behaviour, bias/RMSE, BER, throughput (with an optional erasure window), and contrast margin as robustness indicators.

The main contributions are:

- A clear bulk optics abstraction of polarisation encoded stochastic NOT with two implementable layouts (A and B) suitable for laboratory teaching and component tolerance studies.
- A unified metric framework (contrast C , BER, throughput, margin) and an automated sensitivity workflow (sweeps + joint Monte Carlo) that produces an explicit error budget for finite extinction/leakage, HWP angle/retardance errors, imbalance and detection noise.
- A harmonised head-to-head comparison of A vs. B under common decision settings ($\gamma = 0.60$, erasure window 0.35) revealing distinct sensitivity profiles and practical tolerance guidance.
- A complete reproducible supplementary package: Colab notebooks for A and B (S2–S3), probability level boundary limits (S2.1–S3.1), a comparison pack (S4a), an alignment/calibration protocol (S5), and a phase reference/compensation note for phase sensitive use cases (S6) [15].

2. Encoding, metrics, and what “correct NOT” means for bit streams

2.1. Polarisation encoding and probability estimation

Each time slot k carries a binary stochastic variable $X_k \in \{0, 1\}$ encoded as a Jones vector in the laboratory basis:

- $X_k = 1 \Rightarrow E_{in}(k) = e_H$,

- $X_k = 0 \Rightarrow E_{in}(k) = e_V$.

For a bit stream of length N , the input probability is estimated as:

$$\hat{p}_{in} = \frac{1}{N} \sum_{\{k=1 \dots N\}} X_k. \quad (1)$$

After the optical NOT module and H/V analysis with a decision rule, the output probability is estimated analogously as:

$$\hat{p}_{out} = \frac{1}{N} \sum_{\{k=1 \dots N\}} Y_k, \quad (2)$$

where Y_k is the decided output bit. Under ideal operation:

$$\hat{p}_{out} = 1 - \hat{p}_{in}$$

2.2. Probability-level correctness (Δp criterion)

At the stochastic representation level, correctness is defined by the deviation:

$$\Delta \hat{p}(\hat{p}) = \hat{p}_{out}(\hat{p}) - (1 - \hat{p}_{in}(\hat{p})). \quad (3)$$

To quantify the worst-case behaviour across the full probability range, we use:

$$\Delta \hat{p}_{max} = \max_{\{p \in P\}} |\Delta \hat{p}(\hat{p})|, \quad (4)$$

where P is the predefined probability grid. In Supplement S2.1 and S3.1, tolerance limits are defined by: $\Delta p_{max} \leq \varepsilon$, with $\varepsilon = 1\%$ absolute probability error. This criterion captures probability inversion accuracy independently of any threshold-based decision rule.

2.3. Decision-level metrics: bias, RMSE, and BER

In practical readout, the decision is performed using the normalised contrast:

$$C = \frac{(I_H - I_V)}{(I_H + I_V + \varepsilon_0)}, \quad (5)$$

where ε_0 is a small positive regularisation constant introduced to avoid division by zero when the total detected intensity $I_H + I_V$ approaches zero.

A baseline threshold rule assigns:

- $Y_k = 1$ if $C_k > +\gamma$
- $Y_k = 0$ if $C_k < -\gamma$

with an optional erasure window for $|C_k| < \gamma_{\text{erasure}}$.

From this decision model we compute:

- bias(p) = $\hat{p}_{out}(p) - (1 - p)$
- RMSE evaluated over P
- BER($\hat{p}$) = $\frac{1}{N} \sum_{\{k=1 \dots N\}} 1(Y_k \neq (1 - X_k))$ (6)

Unlike Δp , BER is strongly threshold-dependent and may remain zero over a wide tolerance region (robust regime).

2.4. Robustness indicator: contrast margin

To detect early degradation before BER increases, we introduce the contrast-margin metric:

$$\text{margin} = 1 - q0.05(|C|), \quad (7)$$

where $q_{0.05}$ denotes the 5th percentile of $|C|$. The margin quantifies the separation of the contrast distribution from zero and therefore from the decision threshold. In regimes where $\text{BER} \approx 0$, margin degradation provides a more sensitive robustness indicator.

2.5. Three notions of “correct NOT”

For clarity, we distinguish:

1. Probability-correct NOT: $\Delta p_{\max} \leq \varepsilon$.
Table 1 summarises the nominal optical components, baseline values, and tolerance ranges adopted for the unified model and for the comparative simulations of Architectures A and B.
2. Decision-correct NOT: $\text{BER} \approx 0$ under chosen γ .
3. Robust NOT: margin remains significantly above zero.

These criteria are used consistently for both architectures and in all Supplementary analyses [15].

3. Gate architectures (two implementations)

This section presents both implementations under a shared functional abstraction and then highlights their architecture-specific features.

3.1. Common functional block diagram (shared abstraction)

We treat both realisations as instantiations of the same probability transform block:

- **Input:** Time-slot bit stream encoded as a linear polarisation state in the laboratory basis (H, V)

- **Core operation:** Polarisation exchange (ideal NOT), i.e., swap H and V, giving $p_{\text{out}} \approx 1 - p_{\text{in}}$ under the unipolar mapping.
- **Readout/decision:** H/V analysis + thresholding (optionally with an “erasure” rule).
- **Evaluation:** The same metrics (bias/RMSE/BER) and the same Monte-Carlo error-budget framework.

What is deliberately kept identical across A and B:

- encoding convention and probability estimator $p\text{-hat}$;
- detector/decision model used for BER;
- definition of non-idealities (PBS leakage, HWP angle/retardance errors, imbalance, phase drift); Colab sweep/MC pipeline and exported datasets.

3.2. Implementation A – folded/single output PBS-based NOT (baseline)

Optical idea. PBS1 splits H and V into two arms; each arm contains one HWP set for $H \leftrightarrow V$ exchange; PBS2 recombines into a single output port (folded geometry for compactness and arm access).

Figure 1 illustrates the physical implementation of the stochastic NOT operation.

Key practical constraints:

- Differential delay must satisfy $\Delta\tau \ll T$ (slot duration) to avoid inter-slot mixing.
- Dominant non-idealities expected: PBS extinction/leakage and HWP angle/retardance errors.

Main-text pointer (S5): Practical alignment and calibration (effective ER measurement, HWP angle setting, encoding verification, and decision-threshold calibration) are documented in **Supplement S5** and map directly to the Colab configuration parameters used in sections 5–6.

Table 1.
Optical components and nominal tolerances used for bulk-optics NOT implementations (A and B).

Component / block	Parameter (symbol)	Nominal / example value	Tolerance / range used in simulations	Notes (how it enters the model)
PBS cubes (Architecture A core)	Extinction ratio (ER_{PBS})	30 dB	20–50 dB	Sets leakage power = $10^{-\text{ER}/10}$ in the exchange/recombination path
Polarisation analyser (readout, A & B)	Extinction ratio (ER_{an})	25–35 dB	20–50 dB	Controls readout leakage in analysed channels I_{H} , I_{V}
Half-wave plate(s)	Angle error ($\Delta\theta$)	0°	0–3° (sweeps); boundary via bisection	Common-mode (or per-plate in B) missetting of fast axis
Half-wave plate(s)	Retardance error ($\Delta\delta$)	0°	0–5° (sweeps); boundary via bisection	$\delta = \pi + \Delta\delta$; pincludes fabrication error and/or wavelength detuning mapping $\Delta\delta(\lambda)$
Two-arm path (Architecture A only)	Arm imbalance (ΔL)	0 dB	0–1.5 dB	Amplitude ratio $\rho = 10^{\Delta L/20}$ between arms before recombination
Detectors/acquisition	Relative intensity noise (σ_{rel})	0.01–0.06	0–0.03 (sweeps); boundary estimate	Independent multiplicative noise on I_{H} and I_{V}
Decision rule	Contrast threshold (γ)	0.60	fixed (harmonised)	Decision: $Y = 1$ if $C > +\gamma$; $Y = 0$ if $C < -\gamma$
Decision rule	Erasure window (γ_{erasure})	0.35	fixed (harmonised)	Slot valid if $ C \geq \gamma_{\text{erasure}}$
Wavelength (Architecture B emphasis)	Operating λ	633 nm (example)	400–1500 nm (boundary scan)	Enters via $\Delta\delta(\lambda)$ for HWPs (wavelength mis-tuning analysis)

The table lists the component-level parameters used as baseline inputs for the unified model and for the Supplementary notebooks. Where hardware-specific values are available, you can replace the same fields with measured specifications.




Fig. 1. Optical layout of the folded-PBS implementation of the stochastic NOT module. The stochastic bit stream is encoded in polarisation states H and V. The first polarising beam splitter (PBS1) separates the polarisation components into two optical paths, which are redirected by mirrors (M5, M6) and processed by half-wave plates (HWP3, HWP4). The second beam splitter (PBS2) recombines the beams with exchanged polarisation channels, resulting in inversion of the encoded probability, yielding $p_{out} \approx 1 - p_{in}$. Red and green lines indicate the optical paths corresponding to the H and V polarisation components. Architecture A serves as a minimal demonstrator based on standard PBS cubes and two HWPs.

3.3. Implementation B – HWP-chain NOT (HWP1–HWP3 style) with improved controllability

Optical idea. The NOT operation is implemented as a polarisation exchange using a sequence of HWPs with fixed nominal orientations (e.g., 0° – 45° – 90° in a compact in-line layout) that realises the same net swap in the (H, V) basis.

The corresponding compact in-line implementation is shown in Fig. 2.

Compared with Architecture A, the HWP-chain implementation and its variants provide greater experimental flexibility in a compact in-line configuration. In particular, they enable controlled studies of HWP angle and retardance errors, support explicit phase-reference management when a π phase shift between logical states arises under the adopted convention, and facilitate systematic calibration demonstrations for teaching and metrology-oriented laboratory practice.

Expected failure modes:

- Systematic bias from HWP angle/retardance errors in the chain.
- Phase-reference mismatch between logical states; addressed by **Supplement S6**.

Main-text pointers (S5, S6): The practical alignment/calibration workflow (ER measurement, HWP setting, and decision-threshold calibration) is provided in **Supplement S5**. A phase-reference discussion and the $45^\circ+0^\circ+90^\circ$ compensation logic (relevant only for phase-sensitive use) are provided in **Supplement S6**.

3.4. Comparative analysis of Architectures A and B

From a practical standpoint, Architecture A is better suited to a quick teaching demonstrator, whereas Architecture B is more appropriate for metrology-oriented tolerance studies, wavelength-detuning analysis, and phase-compensation investigations. The comparison focuses on




Fig. 2. In-line three-half-wave-plate ($3 \times$ HWP) implementation of the stochastic NOT module (Architecture B). The stochastic bit stream is encoded in the polarisation state with the convention $H \equiv 1$ and $V \equiv 0$. The NOT operation is realised by three HWPs oriented at nominal angles 0° , 45° , and 90° , which collectively exchange the horizontal and vertical polarisation components. As a result, the polarisation-encoded probability is inverted, yielding $\hat{p}_{out} = 1 - \hat{p}_{in}$. This compact in-line architecture is convenient for tolerance analysis with respect to angular misalignment $\Delta\theta$, retardance error $\Delta\delta$, and wavelength detuning.

tuning freedom, dominant sensitivity channels, alignment complexity, and intended use case.

4. Modelling framework (unified)

This section defines the unified optical-and-decision model used for both implementations. The goal is not to reproduce every optical nuance, but to provide a metrology-grade mapping from measurable component non-idealities to computation-level stochastic metrics under the same readout and decision settings used in the Supplementary notebooks [15].

4.1. Jones representation of the polarisation state

We represent each time-slot field by a Jones vector in the laboratory basis (H, V):

$$E(k) = [E_H(k), E_V(k)]^T. \quad (8)$$

The ideal logical swap (NOT at the polarisation-state level) is the exchange operator $X = [[0, 1], [1, 0]]$, such that $E_{out} = X \cdot E_{in}$.

4.2. Waveplate model (HWP angle and retardance errors)

A general retarder with fast-axis angle θ and retardance δ is modelled as:

$$J(\theta, \delta) = R(-\theta) \cdot \begin{pmatrix} e^{-\frac{i\delta}{2}} & 0 \\ 0 & e^{+\frac{i\delta}{2}} \end{pmatrix} \cdot R(\theta), \quad (9)$$

Where $R(\theta)$ is the 2×2 rotation matrix. An ideal HWP has $\delta = \pi$. We model two classes of deviations:

- Angular mis-setting: $\theta \rightarrow \theta + \Delta\theta$.
- Retardance error: $\delta \rightarrow \pi + \Delta\delta$, where $\Delta\delta$ may represent intrinsic fabrication error and/or wavelength detuning.

In Architecture A, $\Delta\theta$ and $\Delta\delta$ can be treated as per-arm scalar errors (one effective HWP per arm). In Architecture B, we consider either common-mode errors ($\Delta\theta_i = \Delta\theta$, $\Delta\delta_i = \Delta\delta$ for all three plates) or independent per-plate errors ($\Delta\theta_i, \Delta\delta_i$) depending on the experiment (S3).

4.3. PBS/analyser leakage model (finite extinction ratio)

Finite extinction of a PBS/analyser is modelled as leakage between the two analysed components. Using an extinction ratio ER (in dB), the leakage power fraction is

$$\eta = 10^{-\text{ER}/10}, \quad (10)$$

and the corresponding leakage amplitude is $\ell = \sqrt{\eta}$. For a field $E = \begin{bmatrix} E_H \\ E_V \end{bmatrix}$, the analysed intensities are computed via a simple leakage mixing model:

$$I_H = |E_H + \ell E_V|^2, \quad I_V = |E_V + \ell E_H|^2. \quad (11)$$

This symmetric leakage-mixing approximation models a finite-extinction PBS/analyser in which each analysed output contains a small coherent admixture of the orthogonal polarisation component; it is sufficient for the intensity-based contrast decision used throughout this work.

For Architecture A, the PBS cubes inside the gate primarily contribute to effective leakage in the exchange/recombination path; for Architecture B, the dominant ER parameter is typically the external analyser used for readout.

4.4. Two-arm recombination and imbalance (Architecture A)

Architecture A splits the input into two arms and recombines them. In the unified model, we capture the leading-order effects relevant for stochastic metrics:

- Arm imbalance: A relative amplitude ratio $\rho = 10^{\Delta L/20}$ derived from an imbalance in dB (ΔL). One arm is scaled by ρ .
- Timing condition: The analysis assumes $\Delta\tau \ll T$ (slot duration), so recombination occurs within a time slot (no inter-slot mixing).

Phase-sensitive terms are not required for intensity-threshold operation; phase-reference aspects are treated separately in Supplement S6 for phase-sensitive scenarios.

4.5. Detection noise model

Detector and acquisition noise are represented as relative intensity noise (RIN-like) applied independently to I_H and I_V :

$$I' = \max(0, I \cdot (1 + n)), \quad n \sim \text{Normal}(0, \sigma_{\text{rel}}). \quad (12)$$

The parameter σ_{rel} is denoted detector_noise_rel in the Supplementary notebooks [15].

4.6. Contrast and decision rule (harmonised settings)

The readout uses the normalised contrast defined in (5).

The baseline decision rule is a symmetric threshold with an optional erasure window:

- decide $Y = 1$ if $C > +\gamma$
- decide $Y = 0$ if $C < -\gamma$,
- if erasure is enabled: mark a slot valid if $|C| \geq \gamma_{\text{erasure}}$.

Unless stated otherwise, architecture comparisons use harmonised settings consistent with S3 and S4a:

$$\gamma = 0.60 \quad \gamma_{\text{erasure}} = 0.35$$

These thresholds are deliberately strict to expose degradation under non-idealities and to make robustness margins informative.

4.7. Parameter set and mapping to the Supplementary notebooks [15]

Table 2 summarises the primary parameters and their meaning. These symbols are used consistently in the main text and in S2–S4a; Supplement S5 provides measurement procedures and practical calibration steps.

5. Experimental/simulation workflow

This section summarises the experimental/simulation pipeline used to generate all figures and tables. Full implementation details, parameter files, and exported datasets are provided in the Supplementary materials [15].

5.1. Reproducibility and configuration

All simulations are reproducible via fixed random seeds and explicit configuration blocks stored within the Colab notebooks. For each experiment, we record (i) the bit stream length N , (ii) the probability grid P for transfer curves, (iii) the decision settings (γ and optional γ_{erasure}), and (iv) the component non-ideality parameters (ER/leakage, $\Delta\theta$, $\Delta\delta$, imbalance, detector noise). Exported CSV/JSON files include a manifest with the configuration snapshot and file index.

Table 2.
Unified model parameters and mapping to Supplementary configuration fields [15].

Symbol	Physical meaning	Units	Notebook field
ER_dB	Extinction ratio of PBS/analyser	dB	pbs_ER_dB (A), analyser_ER_dB (B)
$\eta = 10^{-\text{ER}/10}$	Leakage power fraction	–	derived from ER_dB
$\Delta\theta$	HWP angular misalignment	deg	hwp_theta_err_deg, dtheta_deg
$\Delta\delta$	HWP retardance deviation from π	deg	hwp_delta_deg, ddelta_deg
ΔL	Two-arm power imbalance (A only)	dB	imbalance_dB
σ_{rel}	Relative detector intensity noise	–	detector_noise_rel
γ	Decision threshold	–	gamma
γ_{erasure}	Erasure window threshold	–	gamma_erasure
λ	Operating wavelength	nm	lambda_nm
$\Delta\lambda$	Wavelength detuning from HWP design	nm	derived via retardance model

5.2. End-to-end pipeline (common to A and B)

For a given architecture and parameter set, the workflow proceeds as:

1. **Bit stream generation:** Draw $X_k \sim \text{Bernoulli}(p_{in})$ for each p_{in} in the grid (or at p_{test} for sensitivity runs) and encode as Jones vectors $E_{in}(k) \in \{e_H, e_V\}$.
2. **Optical gate model:** Propagate $E_{in}(k)$ through the architecture-specific Jones/combiner model (section 4), including non-idealities (finite ER/leakage, HWP $\Delta\theta/\Delta\delta$ imbalance for A).
3. **H/V analysis:** Compute analysed intensities (I_H, I_V) using the leakage-based analyser model; apply detector relative noise independently to each channel.
4. **Decision and validity (optional erasure):** Compute the contrast $C = \frac{(I_H - I_V)}{(I_H + I_V + \epsilon_0)}$, decide Y_k via $\pm\gamma$, and optionally mark a slot valid if $|C| \geq \gamma_{erasure}$.
5. **Metric evaluation:** Compute $\hat{p}_{out}$, bias/RMSE, BER, throughput (= fraction of valid slots), and the robustness in margin $= 1 - q0.05(|C|)$.

Aggregation and export: Average metrics across repeated runs and export figures (300 dpi) and datasets (CSV/JSON) for traceability.

5.3. Automated one-parameter sweeps

To quantify first-order sensitivities, we run automated sweeps over one parameter at a time around a realistic baseline. Each sweep uses multiple stochastic repetitions (n_{rep}) to suppress finite- N variability. The sweep outputs are used to generate paper-ready sensitivity curves and to identify parameter ranges where the module enters a “robust regime” ($BER \approx 0$) vs. a threshold-limited regime.

5.4. Joint Monte-Carlo and impact-based ranking

To obtain factor rankings under simultaneous tolerances, we perform joint Monte-Carlo sampling of all non-idealities according to specified distributions (typically Normal with measured/assumed μ, σ). As the primary ranking for the main paper, we use an impact proxy computed as the absolute change in a target metric under $\mu \pm 1\sigma$ perturbations (impact $\pm 1\sigma$), evaluated at $p_{test} = 0.5$ under the harmonised decision settings ($\gamma, \gamma_{erasure}$). Correlation-based rankings are retained as supplementary diagnostics because they are less robust under strongly non-linear threshold responses.

5.5. Boundary analysis for probability-correct operation (Δp_{max})

In addition to decision-level metrics, we quantify probability-level correctness using the worst-case deviation $\Delta p_{max} = \max_{p \in P} |\hat{p}_{out}(p) - (1 - \hat{p}_{in}(p))|$. Supplements S2.1 and S3.1 implement a scan + bisection boundary search that reports the tolerance limits of key parameters (e.g., $\Delta\theta_{max}, |\Delta\delta|_{max}, \text{noise}_{max}, \text{ER}_{min}$, and $\Delta\lambda_{max}$ where relevant) such that $\Delta p_{max} \leq \epsilon$ with $\epsilon = 1\%$.

5.6. Role of S1 vs. S2/S3 (main-text provenance)

The numerical values in Supplement S1 are provided as an illustrative baseline obtained for nominal parameters. The accompanying Colab workflow is fully reproducible and designed to accept measured specifications and tolerances of specific PBS/HWP/detector components, thereby producing application-specific performance curves and a validated error budget.

To obtain informative sensitivity trends and factor rankings, Supplement S2 introduces measured/assumed component non-idealities and runs automated parameter sweeps and joint Monte-Carlo sampling. The main-text figures and tables are therefore derived from S2 outputs, while S1 serves as a sanity-check reference for idealised behaviour.

Note: unless stated otherwise, architecture comparisons use harmonised decision settings consistent with S3 and S4a: $\gamma = 0.60$ and $\gamma_{erasure} = 0.35$.

6. Results

This section reports representative results derived from the non-ideal, harmonised workflows (S2–S4a). The objective is twofold: (i) verify probability inversion at the stochastic level, and (ii) make architecture-dependent sensitivities operational in terms of measurable component parameters.

6.1. Transfer behaviour: probability inversion under realistic baselines

We first evaluate the transfer curve $\hat{p}_{out}(p_{in})$ for both architectures under a realistic non-ideal baseline (finite extinction/leakage, detector noise, and nominal alignment). Both implementations reproduce the expected inversion trend $p_{out} \approx 1 - p_{in}$ across the probability grid P , with deviations summarised by bias/RMSE and by the worst-case probability error Δp_{max} . The corresponding probability-transfer characteristics are shown in Fig. 3.




Fig. 3. Probability transfer curves for the two bulk-optics NOT implementations. Estimated output probability $\hat{p}_{out}$ vs. input probability $\hat{p}_{in}$ for Architecture A (folded PBS exchange) and Architecture B ($3 \times$ HWP chain), evaluated over the probability grid P . The dashed line indicates the ideal inversion $\hat{p}_{out} = 1 - \hat{p}_{in}$. Unless stated otherwise, both architectures are evaluated using identical decision settings $\gamma = 0.60$ and $\gamma_{erasure} = 0.35$ with $N = 4096$ time slots per probability point.

6.2. Architecture comparison under harmonised decision settings

To compare implementations under a common readout rule, we use the harmonised decision settings introduced in sections 4–5 ($\gamma = 0.60$, $\gamma_{\text{erasure}} = 0.35$) and evaluate metrics at $p_{\text{test}} = 0.5$. This strict rule intentionally exposes degradation when the contrast distribution approaches the decision threshold.

6.3. Sensitivity to HWP angular misalignment ($\Delta\theta$): BER response

Figure 4 shows BER as a function of a common-mode HWP angular offset $\Delta\theta$ (applied consistently within each architecture). BER remains near zero while the contrast distribution stays well separated from the threshold; once the separation collapses, BER rises sharply, reflecting the non-linear nature of a threshold decision.

Table 3 summarises the impact-based ranking of the dominant sensitivity factors for both implementations. It shows impact-based sensitivity ranking (impact $\pm 1\sigma$) for the two bulk-optics stochastic NOT implementations at $p_{\text{test}} = 0.5$ under harmonised decision settings $\gamma = 0.60$ and $\gamma_{\text{erasure}} = 0.35$. Two criteria are reported:

- BER impact (direct change in bit-error rate) and BER,
- margin impact for “margin = $1 - q(0.05(|C|))$ ”, which captures early degradation of the contrast distribution before BER increases. Zero BER impact indicates a robust regime where BER remains ≈ 0 under $\pm 1\sigma$ perturbations within the explored tolerance range; in this regime, the margin-based ranking provides the more informative robustness indicator.




Fig. 4. BER vs. HWP angular misalignment ($\Delta\theta$), Architecture B. BER vs. a common mode HWP angular offset $\Delta\theta$ for Implementation A and Implementation B, evaluated with identical decision settings ($\gamma = 0.60$; and $\gamma_{\text{erasure}} = 0.35$) at $p_{\text{test}} = 0.5$. The sharp transition indicates the onset of threshold-limited behaviour.

6.4. Leakage/ER sensitivity: contrast-margin as an early robustness indicator

In many practical tolerance ranges, the gate operates in a robust regime where $\text{BER} \approx 0$, so BER-based ranking can appear flat even though robustness is degrading. To capture early degradation, we analyse the margin metric based on the lower tail of $|C|$. Figure 5 plots margin against leakage power (equivalently ER) on a shared x-axis. Leakage primarily reduces contrast separation before it produces a noticeable probability bias, making margin a sensitive indicator for comparing dominant error channels.

Table 3.

Dominant sensitivity factors for Implementations A and B under harmonised decision settings ($\gamma = 0.60$, $\gamma_{\text{erasure}} = 0.35$) at $p_{\text{test}} = 0.5$, quantified using the impact $\pm 1\sigma$ method.

Implementation	Evaluation criterion	Parameter (physical meaning)	Impact ($\pm 1\sigma$)
A (folded PBS exchange)	BER impact	PBS extinction ratio ER (reported via leakage power ($\eta = 10^{-\text{ER}/10}$))	0.000000
		Two-arm imbalance ΔL (dB)	0.000000
		Relative detector noise (σ_{rel})	0.000000
	Margin impact	PBS extinction ratio ER (via leakage η)	0.053154
		Common-mode HWP angular offset $\Delta\theta$ (deg)	0.004897
		Relative detector noise (σ_{rel})	0.003803
B ($3 \times$ HWP chain)	BER impact	Analyser extinction ratio ER (via leakage η)	0.000000
		Relative detector noise (σ_{rel})	0.000000
		Wavelength detuning $\Delta\lambda$ (nm)	0.000000
	Margin impact	Analyser extinction ratio ER (via leakage η)	0.010222
		Relative detector noise (σ_{rel})	0.000507
		Wavelength detuning $\Delta\lambda$ (nm)	0.000172

Note. The margin metric is defined as margin = $1 - q(0.05(|C|))$. The reported impact ($\pm 1\sigma$) is computed at $p_{\text{test}} = 0.5$ by perturbing the selected parameter from its baseline value μ to $\mu \pm 1\sigma$ while keeping all other parameters fixed. Unless otherwise stated, σ values correspond to the default tolerances used in the S4a notebook and may be replaced by measured component specifications.




Fig. 5. Contrast-margin robustness vs. analyser ER for the stochastic NOT module (Architecture B). The robustness margin $= 1 - q0.05(|C|)$ is plotted vs. analyser ER, expressed as leakage power $10^{-ER/10}$. Lower ER (higher leakage) reduces the available contrast margin under harmonised decision settings ($\gamma = 0.60$, $\gamma_{\text{erasure}} = 0.35$). The margin degradation precedes observable BER increase and therefore provides an early robustness indicator. ER is expressed in dB on the horizontal axis.

6.5. Probability-correct operating limits ($\Delta p_{\max}$ boundary analysis)

Decision-level metrics depend on the chosen thresholds. To provide a threshold-independent notion of correctness at the stochastic representation level, we report boundary limits derived from the worst-case probability error $\Delta p_{\max}$ across P . Supplements S2.1 and S3.1 implement a scan + bisection search that determines tolerance limits such that $\Delta p_{\max} \leq \varepsilon$ with $\varepsilon = 1\%$. These limits quantify the “safe” region where the module performs probability inversion to within 1% absolute error over the full probability range.

The tolerance limits reported in Table 4 quantify **probability-level correctness** and are therefore inherently threshold-independent. As a result, these $\Delta p_{\max}$ boundaries

can appear substantially broader than robustness limits inferred from decision-level metrics such as BER or contrast margin under strict threshold settings. This difference is expected: $\Delta p_{\max}$ evaluates the worst-case deviation of the estimated probability transfer curve from ideal inversion over the full probability grid, while decision-level metrics are sensitive to the proximity of the contrast distribution to the applied thresholds (γ , γ_{erasure}). Consequently, probability-correct operation may still be guaranteed even in parameter regimes where margin degradation is already observable but has not yet translated into significant decision errors.

6.6. Practical interpretation: when A or B is preferable

The combined results support the following operational interpretation.

- Implementation A is the simplest teaching-grade demonstrator but is typically more exposed to PBS leakage/ER and two-arm imbalance; these effects first appear as margin degradation and later as threshold-limited BER.
- Implementation B concentrates the transform into a small number of waveplates and is therefore well suited to systematic HWP tolerance studies and to wavelength-detuning analysis [via $\Delta\delta(\lambda)$]; phase-reference aspects are relevant only in phase-sensitive scenarios and are addressed in Supplement S6.

In both cases, the recommended workflow is the following: calibrate ER and HWP angles (Supplement S5), set the decision thresholds (γ , γ_{erasure}) to the desired strictness, and then use the provided Colab notebooks to compute application-specific sensitivity curves and an explicit error budget.

7. Discussion

The results highlight a key point for bulk-optics stochastic gates: decision-level performance (BER under a contrast

Table 4.
Probability-correct operating limits $\Delta p_{\max}$ (boundary analysis) for Architectures A and B.

Parameter (worst-case over probability grid P)	Architecture A (folded PBS)	Architecture B (3 × HWP chain)	Units	Notes
ε	0.01	0.01	–	Absolute probability error
λ_0 (HWP design)	–	532	nm	Used only for B (wavelength detuning model)
$\Delta\lambda_{\max}$	–	509.36	nm	B only: wavelength-detuning limit derived from the HWP retardance model $\Delta\delta(\lambda)$.
ER _{min} (PBS/analyser)	not reached	not reached	dB	Limit not reached within tested ER range
$\Delta\theta_{\max}$ (common mode)	21.03	22.01	deg	Common-mode angular offset applied consistently within each architecture
$ \Delta\delta _{\max}$ (common mode)	88.08	88.06	deg	Retardance error w.r.t. ideal π
noise _{max} (σ_{rel})	0.424	0.431	–	Relative intensity noise in both H/V channels
imbalance _{max}	not reached	–	dB	A-only: two-arm amplitude /throughput imbalance limit not reached within tested range

Tolerance limits obtained from a scan + bisection boundary search using the worst-case probability deviation $\Delta p_{\max} = \max_{p \in P} |p_{\text{out}}(p) - (1 - p_{\text{in}}(p))|$ over the full probability grid P . All reported limits correspond to the criterion $\Delta p_{\max} \leq \varepsilon$ with $\varepsilon = 1\%$ absolute probability error. Entries marked “not reached” indicate that no violation of $\Delta p_{\max} \leq \varepsilon$ occurred within the tested sweep range. Boundary limits were validated near the transition using increased bit stream length N and multi-seed Monte-Carlo repeats.

threshold) and probability-level correctness ($\Delta p_{\max}$) are complementary and can lead to different practical conclusions. Below we interpret the dominant sensitivities in physically transparent terms and translate them into guidance for laboratory implementations and metrology-oriented studies.

7.1. Robust regime vs. threshold-limited regime

Over a wide range of realistic tolerances, the NOT modules operate in a **robust regime** where the contrast distributions for logical 1 and 0 remain well separated from the decision thresholds. In this regime, BER can remain ≈ 0 even as robustness degrades. The transition to a **threshold-limited regime** is abrupt because the decision rule is non-linear: once a significant fraction of samples approaches $\pm \gamma$ (or enters the erasure window), small further changes in leakage, noise, or alignment rapidly increase BER and/or reduce throughput. This explains why BER-based rankings can appear flat/empty for moderate perturbations and then change sharply near the transition.

7.2. Why the contrast-margin metric is an “early warning”?

The margin metric $\text{margin} = 1 - q0.05(|C|)$ probes the lower tail of the contrast magnitude and therefore quantifies how close the system is to the decision boundary even when BER is still zero. Physically, the tail of $|C|$ is sensitive to any mechanism that reduces the H/V separation in a subset of samples (e.g., leakage mixing, relative intensity noise, or partial polarisation exchange due to $\Delta\theta/\Delta\delta$). As a result, the margin provides a more informative indicator of “how much robustness is left” in the robust regime, while BER becomes informative once the contrast tail crosses the effective threshold.

7.3. Architecture-dependent sensitivities: what dominates and why?

Although both implementations realise the same ideal polarisation swap, their error channels differ.

- **Implementation A (folded PBS exchange)** is most exposed to PBS extinction/leakage in the core split/recombine path and to two-arm imbalance. Leakage mixes the orthogonal components before and after recombination, directly reducing contrast separation. Imbalance changes the effective recombination weights, which can bias the contrast distribution asymmetrically. These mechanisms can degrade margin well before they manifest as a noticeable probability-level error.
- **Implementation B ($3 \times$ HWP chain)** concentrates the transform into birefringent elements; therefore, HWP angle ($\Delta\theta$) and retardance errors ($\Delta\delta$) are the primary internal sensitivities. In addition, wavelength mistuning can be mapped into $\Delta\delta(\lambda)$, making B a natural platform for studying broadband operation of a nominally single-wavelength HWP design. The external analyser ER mainly affects readout; when high-quality analysers are available, B can be made comparatively insensitive to leakage in the gate itself.

These qualitative trends explain why the head-to-head comparison in S4a separates the architectures more strongly under margin/BER criteria than under probability-level $\Delta p_{\max}$ boundaries.

7.4. Why probability-level boundaries ($\Delta p_{\max}$) can look similar for A and B?

The boundary analysis uses $\Delta p_{\max}$, i.e., the worst-case deviation of the estimated probability transfer curve from ideal inversion across the full p-grid. This criterion is **insensitive to phase** and, importantly, can remain small even when the decision margin is eroding, because probability estimation averages over many slots and does not explicitly penalise near-threshold samples as long as their average decision outcomes remain balanced.

Consequently, A and B can exhibit similar $\Delta p_{\max}$ limits if (i) the analysed intensities are modelled with symmetric leakage mixing, (ii) perturbations are applied in a common-mode fashion, and (iii) the probability estimator uses sufficiently large N . In practice, $\Delta p_{\max}$ is best interpreted as a conservative guarantee for probability inversion accuracy, whereas margin/BER quantify operational robustness under a chosen decision policy.

7.5. Practical guidance and recommended workflow

For teaching-grade demonstrations and calibrated studies, we recommend the following workflow.

1. **Calibrate the analyser/PBS ER and mapping H/V $\leftrightarrow$ 1/0** (Supplement S5). ER is often the dominant “hidden” parameter controlling contrast.
2. **Align the exchange** (A: minimise leakage and imbalance; B: set HWP angles and verify swap) and verify transfer $\hat{p}_{out} \approx 1 - \hat{p}_{in}$ using the S1/S2/S3 notebooks.
3. **Choose decision strictness** (γ and optional γ_{erasure}) according to the desired trade-off between robustness and yield; report these settings explicitly (as done here with $\gamma = 0.60$ and $\gamma_{\text{erasure}} = 0.35$ for harmonised comparisons).
4. **Quantify tolerances** using both layers: (i) $\Delta p_{\max}$ boundaries (S2.1/S3.1) for probability-correct operation guarantees, and (ii) margin/BER sensitivity curves (S2/S3/S4a) for operational robustness.

This two-layer reporting makes the results transferable across laboratories by separating physics-limited probability inversion from decision-policy-dependent performance.

8. Conclusions

We presented a metrology-oriented, bulk-optics realisation of a stochastic NOT primitive for a polarisation-encoded bit stream ($1 \equiv H$, $0 \equiv V$), framed explicitly in computation-level metrics and tolerances rather than in idealised truth-table behaviour.

1. **Unified modelling-to-metrics bridge.** A single Jones-calculus model coupled to stochastic bit stream simulation maps measurable optical non-idealities (finite ER/leakage, $\Delta\theta/\Delta\delta$ of HWPs, two-arm imbalance, detector noise) into computation-level outcomes ($\hat{p}$ transfer, bias/RMSE, BER, and a contrast-margin robustness indicator).

2. **Architecture-dependent sensitivity is operational.** Under harmonised decision settings ($\gamma = 0.60$, $\gamma_{\text{erasure}} = 0.35$), Implementations A (folded PBS exchange) and B ($3 \times$ HWP chain) exhibit distinct dominant error channels, enabling a practical choice between a teaching-grade demonstrator (A) and a compact tolerance-study platform with wavelength detuning capability (B).
3. **Two-layer correctness reporting.** Probability-level correctness, quantified by Δp_{max} , provides a threshold-independent guarantee for probability-inversion accuracy, while decision-level metrics (BER/throughput) and the margin metric quantify operational robustness under a chosen readout policy. Reporting both layers makes the conclusions transferable across laboratories.
4. **Reproducible supplementary package [15].** The full Colab workflows, exported datasets, boundary/tolerance analyses, and practical protocols (alignment/calibration and phase-reference compensation for phase-sensitive scenarios) are provided as Supplementary materials to support reuse and extension.

Outlook. The same methodology can be extended to additional stochastic primitives (AND/OR), to cascaded gate chains, and to laboratory curricula where measured component tolerances are injected into the provided notebooks to obtain experiment-specific error budgets.

Authors' statement

All authors contributed substantially to the conception, design, and completion of this work: M.S. – research concept and design; collection and assembly of data; data analysis and interpretation; writing the article. O.A., C.Z., R.Z., B.T., X.Z., J.Z., C.J. – critical revision of the article and final approval of the article.

All authors discussed the results, revised the manuscript collaboratively, and approved the final submitted version.

Acknowledgements

This research was supported by the Department of Correlation Optics at Yuriy Fedkovych Chernivtsi National University and of the National Research Foundation of Ukraine (NRFU), project no. 2025.06/0086 ‘Innovative Methods and Systems of Laser Communication in Turbulent Atmospheric Environments Based on Phase Singularities’.

Funding

No external funding was received for the research. The Taizhou Institute of Zhejiang University, Taizhou, China, covered the open-access publication charges.

Ethics statement

This study did not involve human participants, human data, or animal subjects, and therefore did not require ethical approval.

Supplementary material

The supplementary package includes reproducible Colab notebooks, parameter sweep datasets, Monte-Carlo outputs, boundary-analysis results, and practical alignment/calibration protocols used to generate the figures and tables in this work. The materials contain the workflows corresponding to Supplements S1–S6 and the associated CSV/JSON data files. The data, scripts, and processed outputs supporting the findings of this study are publicly available in the supplementary repository [15].

Data and code availability statement

All data and code supporting the findings of this study are publicly available in the supplementary repository [15]. The package includes the Colab notebooks used for simulation, exported datasets (CSV/JSON), and configuration manifests enabling full reproducibility of the reported results. This study used no external datasets.

Reproducibility statement

All numerical results reported in this study can be reproduced using the provided Colab notebooks and configuration files included in the Supplementary materials. The workflows use fixed random seeds and documented parameter settings, ensuring that the figures and tables can be regenerated from the distributed datasets and scripts.

References

- [1] Alaghi, A. & Hayes, J. P. Survey of stochastic computing. *ACM Trans. Embed. Comput. Syst.* **12**, 92 (1–19) (2013). <https://doi.org/10.1145/2465787.2465794>
- [2] Gaines, B. R. Stochastic Computing Systems. in *Advances in Information Systems Science* 2 (ed. Tou, J. T.) 37–172 (Plenum Press, New York, 1969).
- [3] Alaghi, A., Qian, W. & Hayes, J. P. The promise and challenge of stochastic computing. *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **37**, 1515–1531 (2018). <https://doi.org/10.1109/TCAD.2017.2778107>
- [4] Lohmann, A. W. Polarization and optical logic. *Appl. Opt.* **25**, 1594–1597 (1986). <https://doi.org/10.1364/AO.25.001594>
- [5] Jones, R. C. A new calculus for the treatment of optical systems. *J. Opt. Soc. Am.* **31**, 488–493 (1941). <https://doi.org/10.1364/JOSA.31.000488>
- [6] Saleh, B. E. A. & Teich, M. C. *Fundamentals of Photonics*, 2nd edition. (Wiley, Hoboken, 2007).
- [7] Goldstein, D. H. *Polarized Light*, 3rd edition. (CRC Press, Boca Raton, 2011).
- [8] Hecht, E. *Optics*, 5th edition. (Pearson, Boston, 2016).
- [9] Camsari, K. Y., Faria, R., Sutton, B. M. & Datta, S. Stochastic p -bits for invertible logic. *Phys. Rev. X* **7**, 031014 (2017). <https://doi.org/10.1103/PhysRevX.7.031014>
- [10] Camsari, K. Y., Salahuddin, S. & Datta, S. Implementing p -bits with embedded MTJ. *IEEE Electron. Device Lett.* **38**, 1767–1770 (2017). <https://doi.org/10.1109/LED.2017.2768321>
- [11] Mandel, L. & Wolf, E. *Optical Coherence and Quantum Optics*. (Cambridge University Press, Cambridge, 1995). <https://doi.org/10.1017/CBO9781139644105>
- [12] Angelsky, O. V., Bekshaev, A. Y., Zenkova, C. Y., Ivansky, D. I. & Zheng, J. Correlation optics, coherence and optical singularities: Basic concepts and practical applications. *Front. Phys.* **10**, 924508 (2022). <https://doi.org/10.3389/fphy.2022.924508>

- [13] Angelsky, O. V. *et al.* Optical phase singularities: Physical nature, manifestations and applications. *Front. Phys.* **10**, 1060787 (2022). <https://doi.org/10.3389/fphy.2022.1060787>
- [14] Angelsky, O. V. *et al.* Random object optical field diagnostics by using carbon nanoparticles. *Opt. Express* **29**, 916–928 (2021). <https://doi.org/10.1364/OE.411118>
- [15] Strynadko, M. Supplementary Materials for Metrology-Oriented Bulk-Optics Implementation of A Stochastic NOT Gate: Unified Error Budget and Tolerance Analysis. *figshare* <https://doi.org/10.6084/m9.figshare.31688230> (2026).