No Entropy Without a Model

A Model-Based Discipline for Entropy Sources

Sara Malik* and Naveed A. Aun

PakCrypt NPO
*smk@pakcrypt.org

Abstract. Hardware and quantum random number generators are routinely justified by an appeal to physical unpredictability followed by a clean pass through a statistical test battery. We argue that this justification is insufficient. Unpredictability is not a property of an output string; it is a property of an adversary's uncertainty, and a cryptographic argument must *lower-bound* that uncertainty. The admissible object is therefore a proven lower bound on the conditional min-entropy of the raw source, derived from a stochastic model of the underlying physics, held invariant online by health tests whose detection power is proven against the same model, protected from side-channel leakage that would hand the adversary the very uncertainty the bound credits, and only then passed to an entropy-preserving conditioner. We make this position concrete. We treat three classical noise sources (oscillator jitter, amplified thermal noise, metastability) and two quantum sources (single-photon which-path and vacuum-fluctuation homodyne); for each we derive an analytical output model $p(x \mid \theta)$ from first principles together with an explicit falsification criterion; we give a measurement and estimation procedure aligned with the NIST SP 800-90B non-IID track and with the 2024 revision of BSI AIS 31; we bind every health test to a model parameter and validate the discipline on synthetic data drawn from the models themselves. The recurring lesson is that a conditioner cannot manufacture entropy a model never proved.

Keywords: True random number generators · Quantum RNG · Min-entropy · Stochastic model · SP 800-90B · AIS 31 · Health tests · Side-channel leakage · Conditioning.

1 Introduction

A random number generator that merely *looks* unpredictable offers a cryptographer little assurance, because essentially every deterministic generator that was eventually broken also looked unpredictable beforehand. The clearest illustration is Dual_EC_DRBG. It was standardised in NIST SP 800-90, shipped by default in widely deployed libraries, and passes statistical test batteries without complaint; yet Shumow and Ferguson showed in 2007 that an adversary who knows a secret relation between two curve points can reconstruct the generator's internal state from a short run of output and predict everything that follows [7,

8]. The output was indistinguishable from random to anyone without that side information and fully predictable to anyone with it. A purely classical example makes the same point without malice: the Mersenne Twister passes most of Dieharder and TestU01, yet its entire future is determined once an observer has seen 624 consecutive outputs. "Passes the tests" and "is unpredictable to the adversary" are simply different properties.

Randomness, then, is not a visible feature of a bit string emerging from a black box. It is the gap between what the adversary knows and what the adversary would need to know to predict the next sample. Building an entropy source is the discipline of bounding that gap from below, and of keeping it bounded while the device runs. Everything in this paper follows from taking that one sentence literally.

This view is sometimes treated as overly conservative in a market that sells *quantum* random numbers, where the word *quantum* is occasionally presented as if it settled the matter. It does not, and the reason is physical rather than rhetorical. A photon striking a balanced beam splitter is, by the Born rule, unpredictable with one full bit of entropy. But the detector that registers it has a dark-count rate, a dead time, an afterpulsing probability and an efficiency mismatch between its two arms, and every one of those quantities is classical, modellable, and—to an adversary who can measure or perturb the device—potentially known. The extractable randomness is not the entropy of the ideal photon; it is the *conditional* min-entropy of the realised detector event given that classical side information.

This distinction is the heart of the matter, so we state it directly. Quantum mechanics provides a *ceiling* on unpredictability: it guarantees that an ideal measurement of a suitable state is intrinsically random, a guarantee no classical source can claim [24, 25]. It says nothing about the *floor*, which is set entirely by how well the real apparatus approximates the ideal and by how much of the observed fluctuation is classical noise the adversary may share. A QRNG that monitors its classical imperfections can approach its ceiling; one that ignores them ships a number whose true floor is unknown. The strongest QRNG constructions make this explicit by bounding output entropy from observed data under minimal trust assumptions (source-device independence) [25, 27]; the weakest quote the ideal bit and stop.

The error we address. A common informal argument runs as follows: the source is physically unpredictable; a hash or a von Neumann corrector compresses it; the compressed stream passes NIST STS and Dieharder; therefore the output has full entropy. Each step is either a non-sequitur or false. "Physically unpredictable" is unquantified until a model attaches a number to it. Compression *preserves* min-entropy at best and usually reduces it—it can never increase it, because a deterministic function of X cannot have more min-entropy than X: any guessing strategy for the input induces at least as good a strategy for the output (formally, $\mathrm{H}_\infty(f(X)) \le \mathrm{H}_\infty(X)$ for every deterministic f, the data-processing inequality for min-entropy). And a statistical battery detects gross structure while certifying nothing about entropy: a counter encrypted under AES passes every test in STS and Dieharder and has zero entropy given the key.

Where the standards stand, and where we stand relative to them. Two evaluation traditions govern entropy sources, and they emphasise different things. NIST SP 800-90B [1] is estimator-centric: it requires a description of the noise source and then runs a battery of conservative min-entropy estimators on raw data, taking the minimum, with a lighter formal demand on a closed-form physical model. BSI AIS 31—in its 2024 revision, now Version 3.0 of the mathematical-technical reference by Peter and Schindler [5], superseding the 2011 Killmann–Schindler document [4]—is model-centric: for its higher class PTG.3 it requires an explicit *stochastic model* of the raw signal from which the entropy is derived and against which the online tests are justified. The two are converging: BSI and NIST have publicly described joint work comparing AIS 20/31 with the SP 800-90 series [6]. Our position is the practical intersection of the two, and we do not claim it as novel doctrine: it is what an evaluator who takes both seriously already does. We restate it because the community sometimes presents the model and the health tests as compliance overhead bolted onto a source that is "obviously random," when in fact they are the only content the claim has. An entropy claim is admissible if and only if it is the conclusion of the following linked obligations.

1. A *stochastic model* $p(x \mid \theta)$ of the raw source, derived from device physics, with the parameter vector θ identified with measurable electrical or optical quantities. (This obligation is the AIS 31 PTG.3 requirement; SP 800-90B accepts it as the strongest form of source description.)
2. A *measured* lower confidence bound on the (possibly conditional) min-entropy per raw sample, obtained on raw, pre-conditioning data with the most conservative applicable estimator. "Raw" is made precise in Section 4: it is the digitised noise-source output before any hashing, XOR, or von Neumann step; debiasing is permitted but only downstream of this measurement point, and it cannot create entropy.
3. *Health tests* whose alarm thresholds are computed from the model's worst-case parameters and whose detection latency is proven, by fault injection, to fire before the min-entropy falls below the claim.
4. A *conditioner* whose input-entropy margin is accounted for explicitly, so that the output entropy is a consequence of (2), not an assumption.
5. *Leakage containment*: an argument that the side information E in the conditional bound is complete—that no physical side channel (power, EM, timing, acoustic, photonic) hands the adversary information the model omitted. This obligation is implicit in the conditional bound and explicit in FIPS 140-3's non-invasive-attack requirements; we treat it in Section 5.

In one line: we replace "unpredictable, compressed, and tested" with "modeled, measured, monitored, margined, and shielded." The intuition behind why all five are needed is simple. The model turns a vague physical claim into a number; the measurement checks that the real device meets the number; the health test keeps the number true over time and temperature; the conditioner spends the number honestly; and leakage containment ensures the adversary has not already been handed the number through a side door. Drop any one and the others stop meaning anything.

Contributions. This is a framework paper with a worked, model-grounded core. We contribute (i) a unifying argument that reduces "unpredictability" to a monitored lower bound on conditional min-entropy and shows why the standards' apparatus is the minimal honest content of the word rather than bureaucratic overhead; (ii) first-principles output models $p(x \mid \theta)$ for five source classes, each with the resulting per-sample min-entropy and an explicit criterion under which the model is declared false; (iii) a measurement-and-health-test protocol that binds each online test to a specific model parameter, with worked examples and a proven detection-latency criterion; (iv) an explicit treatment of side-channel leakage as a first-class entropy obligation; and (v) a conditioning and entropy-pool discipline with end-to-end accounting. Throughout, we validate the models on *synthetic* data drawn from the models themselves, which lets us show that the closed-form min-entropy expressions agree with what the SP 800-90B estimators recover, that the homodyne over-statement is quantitative rather than rhetorical, and that the health-test latency criterion is meetable.

Scope of the data. We do not report a laboratory measurement campaign. The figures are produced by simulating each stochastic model and applying the SP 800-90B estimators to the simulated output; they are representative of what the model predicts a typical device of each class would yield, not of any one fabricated part. We are explicit about this because a model is not validated until a laboratory instantiates it on silicon or optics; what synthetic data *can* establish—and what we use it for—is internal consistency between the analytical bound, the estimator, and the health-test threshold.

2 From Unpredictability to Conditional Min-Entropy

2.1 Min-entropy

For a source emitting a symbol X with distribution $\{p_i\}$, Shannon entropy measures the *average* surprise. That is the wrong average for a key, which an adversary guesses once with their single best guess, not repeatedly. The relevant quantity is the worst case, the min-entropy:

$$\mathrm{H}_\infty(X) \;=\; -\log_2 \max_i p_i. \tag{1}$$

The two diverge precisely when one outcome is much more likely than the rest. Consider a source over 129 symbols that emits one fixed symbol with probability $\frac{1}{2}$ and spreads the remaining probability uniformly over the other 128. Its Shannon entropy is $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot (1 + \log_2 128) = 4.5$ bits, which sounds comfortable; its min-entropy is $-\log_2 \frac{1}{2} = 1$ bit, because an adversary who always guesses the fixed symbol is right half the time. A key drawn from this source is broken on the first guess with probability one-half. Shannon entropy is irrelevant to that fact; min-entropy reports it exactly. This is why cryptographic entropy accounting is always a worst-case, lower-bound discipline.

2.2 Conditional min-entropy

Real sources are not isolated, and the adversary is not blind. Let E denote all side information available to or controllable by the adversary: electronic noise, supply ripple, the classical excess noise of a homodyne detector, a temperature the attacker is allowed to drift. The correct object is the average conditional min-entropy of Dodis et al. [10]:

$$\widetilde{\mathrm{H}}_\infty(X \mid E) \;=\; -\log_2 \mathbb{E}_{e \leftarrow E}\Big[\max_x \Pr(X = x \mid E = e)\Big]. \tag{2}$$

The inner $\max_x$ is the adversary's best guess once they have seen the realisation e of the side information; the outer expectation averages their success probability over the side information they will actually see; the negative logarithm turns that success probability into bits. A small worked case fixes the idea. Suppose X is a fair bit but a noisy sensor leaks $E = X \oplus N$, where $N = 0$ with probability 0.9. Unconditionally $\mathrm{H}_\infty(X) = 1$. But an adversary who reads E and guesses $X = E$ is correct with probability 0.9, so $\widetilde{\mathrm{H}}_\infty(X \mid E) = -\log_2 0.9 \approx 0.152$ bits. The bit "looks" fully random in isolation and is worth almost nothing against this adversary. Equation (2) is exactly why a vacuum-fluctuation QRNG cannot quote the variance of its digitised quadrature as entropy: part of that variance is classical excess noise E, and the adversary is permitted to know it. The extractable randomness is what survives conditioning on E, never the total.

2.3 The Leftover Hash Lemma and the price of conditioning

A conditioner cannot create min-entropy; the best it can do is concentrate existing min-entropy into a shorter, nearly uniform string, and even that costs something. The tool that prices it is the Leftover Hash Lemma [9, 10]. A family $\mathcal{H} = \{h : \{0,1\}^n \to \{0,1\}^m\}$ is *universal* if for every pair of distinct inputs $x \neq x'$ a uniformly chosen h collides with probability at most 2^{-m}; the standard low-cost constructions (multiplication in $\mathrm{GF}(2^n)$, Toeplitz-matrix products) satisfy this. Applied with a public uniform seed to a source of conditional min-entropy k, such a family yields output within statistical distance ε of uniform (even given E and the seed) provided

$$m \;\leq\; k - 2\log_2(1/\varepsilon). \tag{3}$$

This is the whole truth about lossless post-processing, and it disposes of the amplification fantasy in one line: if k is small, no choice of hash makes m large. The $2\log_2(1/\varepsilon)$ term is the *entropy gap* one pays for near-uniformity; in cryptographically hardened form it reappears as the "+64" in the NIST conditioning rule of Section 6.

This analogy needs one caveat stated precisely. Eq. 3 is unconditional: it holds for any adversary, with no assumption beyond an honestly-sampled seed independent of the source. The vetted conditioners of Section 6 (SHA-2/SHA-3-based hash df, HMAC, CMAC)are fixed, unseeded functions, not draws from a universal

family. Their entropy-preservation property is therefore not an instance of the Leftover Hash Lemma but a separate, weaker claim: that a specific deterministic function behaves like a random oracle (hash df) or a secure PRF (HMAC, CMAC) on the relevant input distribution. The +64 margin is the standards' heuristic compensation for that gap — a computational-assumption-based safety margin, not a derived information-theoretic bound. We use "cryptographically hardened form" to mean exactly this: the same shape of accounting, ported from an unconditional setting to a computational one, at the cost of the unconditional guarantee.

The practical reading is worth stating because it is the mental model an engineer should carry. What you must defend is a trustworthy lower bound $\hat{H}$ on the input min-entropy. If you hold $\hat{H}$ and you target m output bits at security level ε, Eq. (3) tells you exactly how much input you must consume. A source whose true min-entropy is, say, 7 bits per byte but that you cannot bound below 7—because you have no model, or because no health test ties to it—is not a 7-bit source for accounting purposes; it is an unbounded source, and you are in uncharted territory. The danger is never that the measured bound is conservative. The danger is claiming a bound you cannot defend, because the conditioner will faithfully whiten the output either way and hide the difference.

2.4 Why a stochastic model is necessary, not merely nice

It is tempting to ask why the model is required at all. If one captures raw data, runs the SP 800-90B non-IID estimators, and obtains a healthy number, why is that not enough? The honest answer is that an estimator measures the *sample in front of it* under the assumption that the sample is representative; a model is what licenses that assumption and what tells you when it fails. Three concrete reasons follow.

First, estimators bound the entropy of the *observed distribution*, but an entropy source must hold across an operating envelope—process, voltage, temperature, ageing—that no finite capture exhausts. The model is what extrapolates a bound measured at a few corners to the corners you did not test, because it says *which parameter* (Q, V_{os}, a clearance ratio) governs the entropy and how. Without it, a passing estimate at room temperature is silent about the cold corner.

Second, an estimator cannot separate intrinsic randomness from deterministic structure that merely looks complex. A weakly chaotic but deterministic circuit, or one injection-locked to an attacker's signal, can yield raw data on which the estimators report high entropy while the true conditional min-entropy against an informed adversary is near zero—this is precisely the Dual_EC_DRBG situation re-expressed at the physical layer. The model, by predicting the *shape* of the distribution (Gaussian quadrature, exponential inter-arrival, a specific bias-versus-Q law) and letting you reject the data when the shape is wrong, is the only thing that catches this.

Third, the model is what makes the health test meaningful: a test threshold computed from a model parameter (Section 4.3) is a test of the entropy; a generic statistical test run on the output is a test of the conditioner. This is

why AIS 31 PTG.3 elevates the stochastic model to a requirement rather than a recommendation, and why we adopt its position. The model is not only a description of the noise source; it must also account for every deterministic component—amplifier, comparator, sampler, local oscillator—that shapes the signal before digitisation, because each of those is a place the adversary's E can enter.

Definition (certifiable entropy source). An entropy source is *certifiable* when there exist (i) a stochastic model $p(x \mid \theta)$, (ii) a measured lower bound $\hat{H}$ on $\widetilde{\mathrm{H}}_\infty(X \mid E)$ valid across the declared operating envelope, and (iii) an online test that rejects any trajectory on which the model's parameters leave the region in which $\widetilde{\mathrm{H}}_\infty \geq \hat{H}$ holds. Unpredictability that cannot be written in this form is an aesthetic judgement, not a security parameter.

3 Analytical Output Models

We study five source classes chosen to span the field rather than to win a ranking: free-running oscillator jitter, amplified thermal noise, and metastability on the classical side; single-photon which-path and vacuum-fluctuation homodyne on the quantum side. Each has at least two independent peer-reviewed silicon or optical realisations, a published stochastic model referenced by NIST or BSI material, and raw pre-conditioning data that is reproducible. We deliberately omit a formal scoring procedure: the point of the paper is the discipline applied to each model, not a beauty contest among sources, and the two we set aside in passing (PLL jitter, laser phase noise) are omitted only because their published models are less mature, not because they are unusable.

For each source we give a physics-grounded derivation of the output model $p(x \mid \theta)$, identify θ with measurable quantities, state the resulting per-sample min-entropy, and give the criterion under which we would declare the model *false*. The falsification criterion is not decoration: a model that no measurement could contradict carries no information, and in practice the falsification test *is* the acceptance test an evaluation lab runs. If a model cannot be falsified by the data a given source can produce, that source cannot be certified under the discipline of Section 2.4—one must change the measurement until it can, or change the source.

3.1 Oscillator jitter

Physics. A free-running ring oscillator accumulates phase error because each gate delay is perturbed by thermal noise in the channel of its transistors; over time scales long compared with one period these independent perturbations sum into a random walk in phase. This is the standard Wiener-process model of phase diffusion [11, 12]: the accumulated phase deviation over a sampling interval τ has variance growing linearly, $\sigma^2_{\mathrm{acc}}(\tau) = \kappa\,\tau$, with diffusion coefficient κ fixed by device thermal noise. The linear-in-τ growth is the experimental signature that

distinguishes genuine thermal (white-frequency) jitter from flicker ($1/f$) jitter and from deterministic interference, and it has been characterised repeatedly in silicon and FPGA implementations [13, 14, 11].

Output. Sampling the oscillator with a reference edge reads the phase modulo the period T_0. Writing the normalised sampling phase $\psi = (\phi/T_0) \bmod 1$, the output bit is its most significant bit. Treating the accumulated jitter as Gaussian and expanding the bit probability in a Fourier series, every harmonic of order k is suppressed by $\exp(-2\pi^2 k^2 Q)$ with the single governing parameter

$$Q \;=\; \frac{\sigma_{\mathrm{acc}}^2}{T_0^2}, \tag{4}$$

the accumulated jitter variance normalised by the squared period. The bias is dominated by the first harmonic,

$$\left|\, \Pr(b=1) - \tfrac{1}{2} \,\right| \;\le\; \frac{C}{\pi}\, e^{-2\pi^2 Q}, \tag{5}$$

with a model-dependent constant C of order unity (the precise value for the elementary and XOR-tree architectures is given by Baudet et al. [11]), and the per-bit min-entropy follows,

$$\mathrm{H}_\infty(b) \;=\; -\log_2\!\big(\tfrac{1}{2} + \big|\Pr(b=1) - \tfrac{1}{2}\big|\big) \;\approx\; 1 - \frac{2}{\ln 2}\,\frac{C}{\pi}\, e^{-2\pi^2 Q}. \tag{6}$$

Serial correlation between successive bits decays with the same exponential, so a measured Q bounds bias and dependence simultaneously: the entropy claim reduces to measuring one jitter variance and one period.

Synthetic validation. Figure 1 draws bits from this model at worst-case sampling phase and compares the SP 800-90B most-common-value (MCV) lower bound on 10^6 simulated bits against both the exact wrapped-Gaussian bit probability and the first-harmonic closed form. The MCV estimate sits on the exact curve to within its confidence interval (e.g. 0.768 versus 0.773 at $Q = 0.1$, 0.992 versus 0.996 at $Q = 0.3$). Two lessons: the closed form is accurate in the operating regime $Q \gtrsim 0.25$, and it is mildly *optimistic* at low Q, where dropped higher harmonics make it predict more entropy than the source has. That gap is itself an argument for measuring rather than trusting the formula at the edge of the envelope.

Falsification. Collect raw bits across the PVT envelope; estimate Q from a direct period/jitter measurement; predict bias and lag-1 correlation from Eq. (5); compare with measurement. If the measured bias exceeds the prediction beyond the confidence interval, or if its scaling with τ departs from $\exp(-2\pi^2\kappa\tau/T_0^2)$, the Wiener-jitter model is rejected— typically a sign of injection locking or of flicker jitter dominating, either of which changes the scaling law.

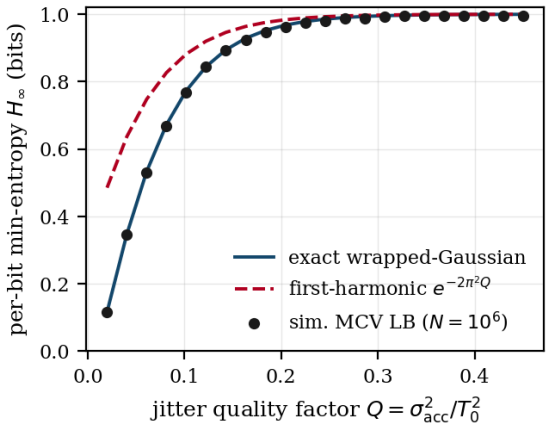

Fig. 1. Oscillator-jitter per-bit min-entropy versus the jitter quality factor Q. Markers: SP 800-90B MCV lower bound on 10^6 bits simulated from the Wiener model at worst-case sampling phase. Solid: exact wrapped-Gaussian bit probability. Dashed: first-harmonic closed form, which is optimistic at low Q.

3.2 Amplified thermal noise

Physics. The thermal agitation of charge carriers in a resistor produces a fluctuating voltage even with no applied signal: at the microscopic level the carriers are in constant random motion set by temperature, and the fluctuation–dissipation theorem ties the resulting open-circuit voltage to temperature and resistance. The Johnson–Nyquist law gives a Gaussian voltage of variance $\sigma_v^2 = 4k_B\mathsf{T}RB$ across a resistor R at temperature T over bandwidth B [15]. This is genuine, irreducible physics, not a circuit artefact, which is why it is an attractive entropy source; the engineering challenge is to amplify it without letting deterministic interference dominate. After gain G and a comparator at threshold V_{th} with input-referred offset V_{os}, the output bit is the sign of $(v - V_{\text{th}})$.

Output. With the threshold at the mean,

$$\Pr(b = 1) = \Phi\left(\frac{-V_{\text{os}}}{G\sigma_v}\right) \approx \frac{1}{2} - \frac{V_{\text{os}}}{G\sigma_v\sqrt{2\pi}}, \tag{7}$$

so the bias is the offset measured in units of the amplified noise standard deviation. Successive samples are independent only if taken slower than the noise correlation time; for a one-pole front end of cutoff f_c, $\tau_c \approx 1/(2\pi f_c)$ and the autocorrelation decays as $e^{-|\tau|/\tau_c}$. Sampling at several τ_c yields a near-IID Bernoulli source of min-entropy $-\log_2 \max(p, 1-p)$. The synthetic source at $V_{\text{os}}/(G\sigma_v) = 0.15$ yields an MCV bound of 0.83 bit/sample, consistent with the closed form.

Falsification. Verify the Gaussian shape of the pre-comparator distribution (Kolmogorov–Smirnov against $\mathcal{N}$), verify that σ_v^2 tracks $4k_B\mathsf{T}RB$ as temperature

is swept, and verify the exponential autocorrelation. Departure indicates amplifier saturation, $1/f$ contamination, or coupled deterministic interference—each of which the model explicitly forbids.

3.3 Metastability

Physics. A cross-coupled latch driven to its metastable point sits, in principle, on an unstable equilibrium between its two stable states; thermal noise breaks the tie. The differential voltage resolves as $\Delta V(t) = \Delta V_0\, e^{t/\tau_{\mathrm{reg}}}$, where τ_{reg} is the regeneration time constant and the initial imbalance $\Delta V_0 = V_{\mathrm{off}} + n$ combines a deterministic offset and thermal noise $n \sim \mathcal{N}(0, \sigma_n^2)$ [16, 17]. The offset V_{off} is dominated by device mismatch from manufacturing variation, so this source is intrinsically process-dependent: nominally identical latches on the same die have different biases, and the offset drifts with voltage and temperature. This is the source's characteristic weakness and the reason its honest model is rarely a clean Bernoulli.

Output. Reading the latch after a fixed window gives

$$\Pr(\mathrm{out} = 1) = \Phi\!\left(\frac{V_{\mathrm{off}}}{\sigma_n}\right), \tag{8}$$

maximised in entropy as $V_{\mathrm{off}} \to 0$. Hysteresis and incomplete settling make the present output depend on the previous one, so the honest model is a two-state Markov chain rather than a Bernoulli source; the transition probabilities are the objects to estimate, and the SP 800-90B Markov estimator is the matched tool. A synthetic chain with $V_{\mathrm{off}}/\sigma_n = 0.3$ and mild settling-induced correlation gives a Markov-estimator bound of 0.79 bit/sample, below the $\Phi(0.3)$ Bernoulli value—the dependence is real entropy loss the Bernoulli view would miss.

Falsification. The model predicts that residual offset, hence bias, is cancellable by a calibration DAC, and that the unresolved-state rate falls as the window lengthens at rate $1/\tau_{\mathrm{reg}}$. If bias persists after offset cancellation, or if the Markov dependence does not vanish as settling time grows, a deterministic coupling is present and the source is not yet an entropy source.

3.4 Single-photon which-path (quantum)

Physics. A single photon at a balanced beam splitter is detected in one of two arms; the Born rule assigns probability $\frac{1}{2}$ to each and one bit of *ideal* entropy. Reality enters through unequal detection efficiencies $\eta_1 \neq \eta_2$, dark counts at rate d, dead time, and afterpulsing probability a. Commercial which-path QRNGs are built on exactly this principle; the ID Quantique Quantis family, for instance, has undergone certification under both AIS 31 and SP 800-90B, and such certifications rest on a vendor stochastic model of precisely these imperfections rather than on the ideal bit alone [28]. The certification reports are a useful template for the evidence an evaluator expects.

Output. Conditioned on a valid detection, the effective bias is, to first order,

$$\Pr(b=1) \approx \frac{\eta_1}{\eta_1+\eta_2} \;+\; (\text{dark-count and afterpulsing corrections}), \tag{9}$$

and inter-arrival times follow the Poisson statistics of the source rate μ. The certifiable quantity is the conditional min-entropy of Eq. (2) with E the classical imperfection parameters: a detector blinding or efficiency-mismatch attack is exactly an adversary exploiting E. The ideal bit is a ceiling; the floor is set by how tightly η_1/η_2, d and a are bounded and monitored.

Falsification. Predict the bias from independently measured η_1, η_2, d, a; measure the realised bias; they must agree. Predict the exponential inter-arrival distribution; deviation signals afterpulsing or dead time outside the model, i.e. unaccounted classical correlation.

3.5 Vacuum-fluctuation homodyne (quantum)

Physics. Homodyne detection of the vacuum quadrature against a strong local oscillator yields a measured value

$$M = Q_{\mathrm{q}} + E_{\mathrm{c}}, \qquad Q_{\mathrm{q}} \sim \mathcal{N}(0,\sigma_q^2), \;\; E_{\mathrm{c}} \sim \mathcal{N}(0,\sigma_e^2), \tag{10}$$

where σ_q^2 is the shot-noise (vacuum) variance and σ_e^2 is classical electronic and excess noise [26, 27].

Output. After ADC discretisation with bin width δ, the extractable randomness is the discretised conditional min-entropy $\widetilde{\mathrm{H}}_\infty(M_\delta \mid E_{\mathrm{c}})$, *not* the entropy of the total variance $\sigma_q^2+\sigma_e^2$. Quoting the latter is the single most common over-statement in continuous-variable QRNG, because it credits the user with randomness the adversary controls through E_{c}. Figure 2 makes the size of the error concrete: as the shot-noise clearance ratio $r = \sigma_q^2/\sigma_e^2$ falls, the honest conditional min-entropy stays pinned near the value set by σ_q (here ≈ 2.35 bits/sample), while the naive total-variance figure inflates without bound—at $r=1$ the over-statement is already about half a bit per sample, and at $r=\frac{1}{4}$ it exceeds a full bit. The clearance ratio is therefore not a performance number but a monitored security parameter.

Falsification. Vary the local-oscillator power: σ_q^2 must scale linearly with it (shot noise) while σ_e^2 must not. If the "quantum" variance does not track the local oscillator, it is not quantum, and the conditional min-entropy claim collapses.

4 Measurement, Estimation, and Model-Bound Health Tests

4.1 Measure the raw source, before conditioning

The cardinal rule is procedural: tap the digitiser output *upstream* of any hash or von Neumann corrector. A conditioner whitens everything, including failure,

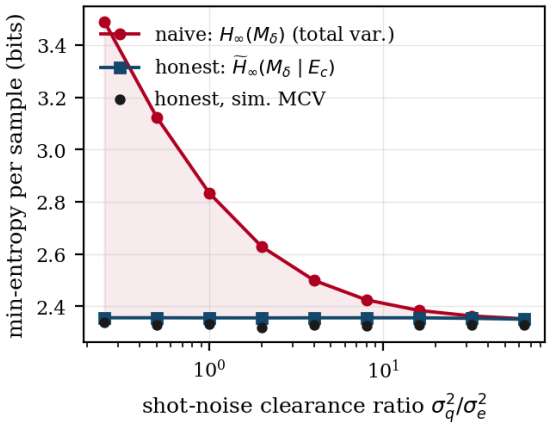

Fig. 2. Homodyne QRNG: honest conditional min-entropy $\widetilde{\mathrm{H}}_\infty(M_\delta \mid E_\mathrm{c})$ (blue, with simulated MCV markers) stays fixed by the quantum variance, while the naive figure crediting total variance (red) inflates as the shot-noise clearance ratio worsens. The shaded gap is randomness the adversary controls.

so measuring after it measures the conditioner, not the source. We can now answer precisely the question of what "raw" permits. The entropy source is a noise source, a digitiser, and an optional conditioner. The *raw random numbers* (the AIS 31 term; the digitised noise-source samples) are what you estimate on, and the choice of *which* digitised quantity is the sample—the single comparator bit, the MSB of a ring oscillator, a chosen ADC sub-range—is part of defining the source, made before estimation. What you may *not* do is apply a von Neumann corrector, an XOR combiner, an LFSR, or a hash before estimating and then credit the cleaned-up statistics. Those are conditioning steps. They are permitted in the data path, but they live downstream of the measurement point, they cannot increase min-entropy ($\mathrm{H}_\infty(f(X)) \leq \mathrm{H}_\infty(X)$), and SP 800-90B caps the entropy creditable at the output of a non-vetted conditioner accordingly. In short: debiasing is allowed, but it is never free and never measured as if it were the source. Log at least 10^6 raw samples per operating corner (nominal, hot, cold, low-voltage, high-voltage), with timestamping for jitter sources and ADC capture at $\geq 10\times$ the noise bandwidth for thermal sources.

4.2 Estimate min-entropy with the non-IID track

Do not assume independence. Run the SP 800-90B IID permutation tests (the standard's Section 5); if IID is not credibly established, take the non-IID route and report the *minimum* over the ten min-entropy estimators of Section 6 (most-common value, collision, Markov, compression, t-tuple, longest repeated substring, and the MultiMCW / Lag / MultiMMC / LZ78Y predictors). Document which estimator dominates, because the dominating estimator names the structure the

source is leaking. This is not folklore; it is mechanical. If the *Markov* estimator gives the minimum, the source has first-order serial dependence—the metastability latch with incomplete settling in Section 3 is the textbook case. If a *predictor* (LZ78Y, MultiMMC) dominates, the source has longer-range or periodic structure an order-1 model misses—injection locking on a jitter source shows up here. If the *most-common-value* estimator dominates, the source is essentially memoryless and merely biased—the thermal source with a comparator offset. Reading the dominating estimator is therefore a free diagnosis of the failure mode, which is exactly why the minimum-over-estimators rule is conservative in the right way.

4.3 Tie every health test to a model parameter

The two SP 800-90B continuous tests are mandatory, and their thresholds are not free parameters: they are computed from the claimed min-entropy $\hat{H}$ and a false-alarm probability α (commonly 2^{-20}),

$$\text{Repetition-Count cutoff } C = 1 + \left\lceil \frac{-\log_2 \alpha}{\hat{H}} \right\rceil, \tag{11}$$

$$\text{Adaptive-Proportion cutoff} = \min\left\{ c : \Pr[\mathrm{Bin}(W, 2^{-\hat{H}}) \geq c] \leq \alpha \right\}, \tag{12}$$

with window $W \in \{512, 1024\}$. These catch *catastrophic* collapse—a stuck source, a totally biased one. They do *not* catch a source that has quietly drifted from 0.9 to 0.6 bits while still looking noisy, and that slow drift is the failure mode that actually occurs in the field as a device ages or warms. The fix is to add, for each source, a direct online estimator of the very parameter θ that the model says governs entropy, with an alarm set at the value of θ where H_∞ meets the claim. The binding is the whole point: a generic output test asks "does this look random?"; a model-bound test asks "is the physical quantity that produces the randomness still in its certified range?". Concretely:

- **Oscillator jitter:** a sliding-window estimator of period variance, tracking Q; alarm when Q falls below the value at which H_∞ meets the claim. This also catches injection locking, which collapses Q abruptly.
- **Amplified thermal:** a DC-offset / bias-drift monitor and an amplifier-saturation detector—the two ways $\Pr(b = 1)$ leaves its modelled band.
- **Metastability:** the unresolved-state rate and the estimated Markov transition asymmetry.
- **Photon which-path:** count-rate bounds, dead-time and afterpulsing monitors, and a continuous efficiency-mismatch check.
- **Vacuum homodyne:** the shot-noise clearance ratio σ_q^2/σ_e^2, alarmed against the value that keeps $\widetilde{\mathrm{H}}_\infty(M_\delta \mid E_\mathrm{c})$ above the claim.

A startup (total-failure) test of at least 1024 bits, tuned to the source's worst-case statistics rather than copied from a reference design, completes the set, as AIS 31 requires a proven total-failure test alongside the online test.

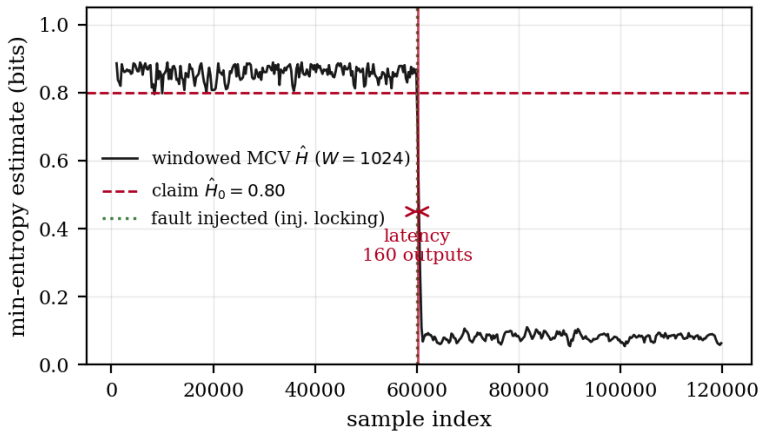

Fig. 3. Health-test effectiveness under injected fault. A sliding-window MCV estimator (W=1024) tracks the jitter source; at the marked instant injection locking collapses Q and the true min-entropy falls from ≈ 0.98 to ≈ 0.11 bit. The alarm fires when the windowed estimate crosses the claim $\hat{H}_0 = 0.8$; the detection latency bounds the number of below-claim outputs emitted.

4.4 Prove effectiveness by injecting the failures you fear

A health test you have not tried to fool is decoration. The effectiveness obligation—explicit in AIS 31—is met by adversarial fault injection: supply droop, electromagnetic injection, thermal shock, local-oscillator manipulation for the homodyne source, controlled illumination for the single-photon source. The acceptance criterion is a *latency* statement: for each fault, the relevant test must alarm before the measured min-entropy crosses below $\hat{H}$, and the number of below-claim outputs emitted between entropy collapse and alarm must be quantified and bounded.

Figure 3 shows the criterion met for the jitter source under a simulated injection-locking fault. The source runs at $\hat{H} \approx 0.98$; at a chosen instant Q collapses, dropping the true min-entropy to ≈ 0.11 bit. A sliding-window MCV estimator ($W = 1024$) tracks the drop and the alarm fires when the windowed estimate crosses the claim $\hat{H}_0 = 0.8$; here the detection latency is 160 outputs, which is the number an integrator must discard or buffer behind the alarm. The procedure an evaluator follows is exactly this: instrument the windowed estimator, inject each fault in the threat model, record the latency distribution and the false-positive rate at the chosen α, and report both as first-class results—because an entropy claim without a detection-latency bound is a statement about the past, not a guarantee about the next sample. Note that a *slow* drift (say $0.9 \to 0.6$ over minutes) would not trip the catastrophic continuous tests at all; only the model-bound Q-monitor of Section 4.3 catches it, which is the concrete payoff of the binding.

5 Leakage: Unpredictability Must Also Be Confidential

The conditional min-entropy $\widetilde{\mathrm{H}}_\infty(X \mid E)$ is only as honest as the side information E it conditions on. If a physical side channel hands the adversary information the model omitted from E, the true conditional min-entropy is lower than the certified one, and every downstream guarantee silently fails. Leakage is therefore not a separate concern bolted onto entropy estimation; it is the question of whether the E in the bound is complete. This is why we make it the fifth obligation rather than an appendix.

The threat is sharper for an entropy source than for a cipher. Leaking a key bit costs one bit; leaking the *state of the noise source* can collapse the entropy of every subsequent output, because the adversary who learns the oscillator phase, the comparator's instantaneous input, or the local-oscillator amplitude can predict samples rather than merely recover one secret. Power and electromagnetic emanations carry exactly these quantities: the switching of the sampling logic, the comparator decision, and the ADC conversion are all correlated with the raw sample, and modern EM probes and template attacks recover such correlations at fine spatial resolution. The same emanations that classical TEMPEST analysis studied for compromising-emanation leakage are now accessible with cheaper instruments and machine-learning-assisted analysis, which lowers the attacker's cost rather than changing the principle. Worse, several of the injection faults of the previous section are *also* side channels run in reverse: frequency injection into a ring-oscillator TRNG can both reduce its entropy and synchronise it to the attacker's reference, and contactless electromagnetic injection has been demonstrated against oscillator-based sources [18, 19]. A device that is monitored for entropy but not shielded for leakage can be driven into a low-entropy, attacker-synchronised regime that its health test was never designed to see.

The standards reflect this. FIPS 140-3, by adopting ISO/IEC 19790 and the SP 800-140F test metrics, makes mitigation of non-invasive (side-channel) attacks a formal requirement area and—unlike its predecessor—calls for side-channel consideration across security levels rather than only at the top tiers [22, 23]. Concretely, an entropy-source design that takes the fifth obligation seriously should: (i) place the noise source and its first amplifier inside the same shielding and power-conditioning boundary as the conditioner, so the raw sample is never exposed on an external rail; (ii) treat the sampling-clock and conversion events as leakage sources and decorrelate or mask them, for example by randomised conversion timing or by balanced/dummy sampling paths that emit the same emanation regardless of the sampled value; (iii) include, where the threat model warrants, a decoy or masking source whose emanations are indistinguishable from the live source, so that an EM attacker cannot tell which signal carries the entropy; and (iv) extend the fault-injection campaign of Section 4 to *emission* measurement, verifying that the recoverable mutual information between an external probe and the raw sample is below a stated bound across the operating envelope. The leakage budget is then reported alongside the entropy budget, and the conditional min-entropy claim is annotated with the E it assumed. A claim that does not state its E is not yet a claim.

6 Conditioning and Entropy Accounting

6.1 The conditioner

The Leftover Hash Lemma (3) establishes, unconditionally, that a conditioner's output min-entropy is bounded by its input min-entropy minus a security margin — but only for the universal-hash-with-public-seed construction. The cryptographic standards do not instantiate that construction; they substitute vetted conditioning components — hash derivation functions over SHA-2 or SHA-3, HMAC, AES in CMAC or CBC-MAC mode — whose entropy-preservation rests on a computational assumption (random-oracle or PRF behavior) rather than on LHL's proof. The accounting below mirrors the LHL margin in structure and is the standards' adopted analogue of it, with the caveat in force throughout this section.

For a vetted component producing n_{out} bits of *full* entropy, SP 800-90B and the now-final SP 800-90C [3] require the input min-entropy to satisfy

$$h_{\text{in}} \;\geq\; n_{\text{out}} + 64. \tag{13}$$

(The full SP 800-90C `Output_Entropy` accounting is more refined than the single inequality, but the +64 margin is the design rule a practitioner uses and is the conservative reading of it.)

Two arithmetic notes. One arithmetic point and one common scoping confusion. First, the $n + 64$ rule applies to the security strength being targeted, not to the conditioner's native output width: n_{out} is the number of full-entropy bits the system actually needs, and SP 800-90B permits truncating a vetted conditioner's output to that figure. For a system that asks SHA-256 to deliver a full-entropy 256-bit block, $n_{out} = 256$ and $h_{in} \geq 320$ follows directly; this is the case we use throughout this section. A practitioner note citing 192 bits is not in error if its target is a 128-bit security strength (e.g. seeding an AES-128 key, or an RBG2 construction needing only a 128-bit seed) — there, $n_{out} = 128$ and $128 + 64 = 192$ is the correct figure for that target. A practitioner must be explicit about which n_{out} their system requires in context of Eq. (13).

Second, min-entropy bits are not raw bits. If the raw source delivers ρ bits of min-entropy per raw bit, obtaining 320 min-entropy bits requires $\lceil 320/\rho \rceil$ *raw* bits—e.g. 534 raw bits at $\rho = 0.6$, or 640 at $\rho = 0.5$.

6.2 Worked accounting, and why it is the operational core

The purpose of the accounting is not bookkeeping for its own sake; it is the bridge between the online health test and the output guarantee. The health test of Section 4.3 maintains a live lower bound ρ on the per-bit min-entropy. The accounting turns that ρ into the compression ratio the conditioner must apply: consume $\lceil (n_{\text{out}} + 64)/\rho \rceil$ raw bits per output block. If ρ degrades but stays above the claim, the conditioner simply consumes more raw bits per block; if ρ

Table 1. Entropy accounting for a 256-bit full-entropy output block via a vetted SHA-256 conditioner. "Required h_{in}" applies the $n + 64$ rule; "raw bits" divides by the conservative per-bit min-entropy ρ (lower confidence bound).

Source (conservative ρ)	ρ	n_{out}	Req. h_{in}	Raw bits	Conditioner
Oscillator jitter	0.50	256	320	640	SHA-256 df
Amplified thermal	0.80	256	320	400	SHA-256 df
Metastability	0.60	256	320	534	SHA-256 df
Vacuum homodyne	0.95	256	320	337	SHA-256 df

crosses the claim, the health test has already alarmed. This is what makes the chain end-to-end: the same measured ρ that the model predicts, the estimator confirms, and the health test defends, is the number that sizes the conditioner. Table 1 carries the discipline through for a 256-bit output. The per-bit ρ values shown are conservative design targets; the nominal-corner synthetic estimates of Section 3 (0.99 jitter, 0.83 thermal, 0.79 metastability, and a multi-bit homodyne sample) sit above them, which is the margin a real design should carry between its certified floor and its typical operating point.

Non IID Warning. This simple division in Table 1 mathematically assumes that min-entropy is strictly additive, which is only true if the source is IID. However, we explicitly notes that sources like metastability follow a non-IID Markov model with serial dependence. For Markov chains, calculating the total min-entropy of a finite block requires evaluating the joint distribution (accounting for the initial state and transition probabilities) rather than naively multiplying an average per-sample marginal rate.

6.3 Stateful entropy pools and the draw discipline

A real generator does not consume raw bits at the rate applications request output. The source delivers entropy at a variable rate (jitter and photon sources especially), while consumers draw asynchronously and in bursts. The standard resolution is a stateful pool sitting between them, and its design is governed, not folkloric. Two regimes must be kept distinct.

In the *information-theoretic* regime—a pool that is itself the entropy reservoir, with no cryptographic generator after it—entropy is a conserved, additive quantity. The pool must accumulate at least $n_{\text{out}} + 64$ min-entropy bits before releasing any output, must *debit* the entropy it spends on each draw, and must block (or signal not-ready) when its accounted entropy is insufficient. Drawing 256 bits genuinely removes 256+64 min-entropy bits from the account; those are not replaced until fresh raw entropy is absorbed. This is the conservative model and the one to use when no vetted DRBG is present.

In the *computational* regime—the construction the SP 800-90C RBG2 and RBG3 classes actually standardise—the pool seeds a DRBG built from a vetted mechanism such as those in SP 800-90A [2], and the security argument shifts

from information-theoretic to computational. Here a single seeding with enough min-entropy ($\geq n_{\text{out}} + 64$ for a full-entropy seed) licenses the DRBG to produce many output blocks, because the DRBG's outputs are computationally indistinguishable from random to a bounded adversary. One does not debit and invalidate positions per draw, as the information-theoretic model would, because forward and backward secrecy are provided by the DRBG's one-way state update, not by entropy depletion. What one *must* do is reseed before the DRBG's reseed interval and after any health-test alarm, and *prediction resistance*, when required, is obtained by reseeding with fresh entropy before the generate call rather than by draining a finite pool.

For accumulation, a robust pool absorbs new raw entropy by mixing it into state with a function that cannot lose previously accumulated entropy even when the new input is adversarially chosen or zero—the formal "robustness" property of Dodis et al., whose analysis showed that a naive Linux-style pool can fail to be robust under state compromise [20]. When the pool is already "full" (the DRBG is seeded to capacity), additional input is not discarded but folded in to refresh state and provide prediction resistance, never to increase a counter past full. Mature reference designs to emulate are the Fortuna accumulator, which sidesteps entropy-estimation fragility by round-robining inputs across staged pools [21], and the current Linux RNG, whose construction BSI has repeatedly analysed against AIS 20 [5]. The single rule that spans both regimes: never emit before the accounted (or seeded) min-entropy meets $n_{\text{out}} + 64$, and let the health test gate the pool's input, so that a failing source stops contributing rather than diluting the account.

6.4 Implementation hygiene

Three rules close the gap between a correct accounting and a correct device. Separate the entropy-source and conditioner clock domains, so the conditioner cannot lock to and launder a source artefact. Zeroise intermediate buffers, so a fault does not leak partially-conditioned state. Run the continuous health tests on the conditioner *input*, not merely its output: testing the output tests the hash, which is precisely the component designed to hide the failure you are trying to catch.

6.5 The meaning of a "near eight bits per byte" claim

One may assert an output of, say, 7.976 bits per byte—a figure seen in particular PTG.3-class certification reports—only when the model, the measured ρ, and Eq. (13) jointly imply it at every operating corner. We stress that 7.976 is a value appearing in specific certification documents, not a constant written into AIS 31, which mandates cryptographic post-processing and computational indistinguishability rather than a single magic number. The honest statement is conditional: *given* $\rho \geq \rho_0$ held by the health tests, the conditioner output meets the target; absent that guarantee, the byte is decorative.

7 The Quantum Case

We single out the quantum case because it is where the word *quantum* does the most rhetorical work relative to its logical content, and because several distinct questions are routinely conflated; comprehensive surveys of quantum entropy sources exist [29], but the conflation persists in practice. We separate them.

Does a QRNG carry a real advantage? Yes. Its ideal model has a first-principles entropy that is a consequence of quantum measurement rather than an estimate of a complicated classical process. That advantage is genuine and worth paying for. But it is an advantage in the *ceiling*; the security parameter is the *floor*, and the floor is classical. The single-photon device leaks through detector imperfections; the homodyne device leaks through classical excess noise. In both, the certifiable output is the conditional min-entropy of Eq. (2), and a credible QRNG measures and monitors its classical side channel—efficiency mismatch, dark counts, shot-noise clearance—with the same rigour a classical TRNG applies to jitter or offset.

Do quantum computers threaten hardware RNGs? Not at the entropy-source layer. A quantum computer attacks the hard-problem assumptions behind public-key cryptography and, via Grover, halves the effective key length of symmetric primitives; it does not predict a well-characterised physical entropy source, whose security is information-theoretic at the raw layer and does not rest on a computational assumption. The post-quantum transition is a story about algorithms downstream of the RBG, not about the noise source. The one caveat is the conditioner: a hash or block cipher used for conditioning should retain adequate security under Grover, which the standard 256-bit choices already do.

Do quantum sensors make some sources more predictable? This is the more interesting threat. Improved quantum-limited sensing lowers the cost of measuring the very classical quantities that constitute E—the local-oscillator phase, a detector's afterpulsing, a faint emanation—and so can *enlarge* the adversary's side information against a source whose leakage budget was set against weaker instruments. This is a leakage question (Section 5), and it is an argument for stating the assumed measurement capability of the adversary in the model, not for or against any particular source.

Certified versus uncertified, third-party versus in-house. The discriminating question to ask any QRNG vendor is therefore not "is it quantum?" but "what is your conditional min-entropy given your own classical noise, which online test holds it, and what adversary measurement capability did your leakage budget assume?" A well-understood, transparently certified classical TRNG with a defended conditional min-entropy bound is preferable to an opaque QRNG that ships the ideal bit; an opaque QRNG is preferable to nothing only if one trusts the vendor's unstated model, which is exactly the trust this paper argues against extending. Transparency of the model is not a marketing nicety; under the definition of Section 2.4 it is what "certified" means.

8 Conclusion

The community has the right standards and, sometimes, the wrong emphasis. The model, the health test, the leakage budget and the entropy accounting are routinely treated as compliance burden laid over a source that is "obviously random." The order of dependence is the reverse: those four are the only content the word *random* has in a cryptographic setting, and the source is nothing without them. Unpredictability is necessary, and—once written as a proven, monitored, leak-contained lower bound on conditional min-entropy—it is also sufficient, because at that point the Leftover Hash Lemma supplies the rest for free. The error worth eradicating is the belief that the conditioner can rescue a source the model never characterised. It cannot: a hash function is an excellent way to hide that one does not know how much entropy one has, and a poor way to obtain any.

References

1. M. S. Turan, E. Barker, J. Kelsey, K. A. McKay, M. L. Baish, M. Boyle: *Recommendation for the Entropy Sources Used for Random Bit Generation*. NIST Special Publication 800-90B (2018).
2. E. Barker, J. Kelsey: *Recommendation for Random Number Generation Using Deterministic Random Bit Generators*. NIST Special Publication 800-90A Rev. 1 (2015).
3. E. Barker, J. Kelsey, K. A. McKay, A. Roginsky, M. S. Turan: *Recommendation for Random Bit Generator (RBG) Constructions*. NIST Special Publication 800-90C (final, 2025). *[Specifies the RBG1/RBG2/RBG3/RBGC constructions and the* `Output_Entropy` */* $n + 64$ *conditioning accounting.]*
4. W. Killmann, W. Schindler: *A Proposal for: Functionality Classes for Random Number Generators*, Version 2.0. BSI AIS 20/AIS 31 (2011). *[Superseded historical reference; original PTG.2/PTG.3 definitions.]*
5. M. Peter, W. Schindler: *A Proposal for: Functionality Classes for Random Number Generators*, Version 3.0. BSI, mathematical-technical reference for AIS 20/AIS 31 (10 September 2024). *[Current BSI reference; PTG.2_2024 / PTG.3_2024, stochastic-model requirement; first-time certifications under v2.0 close 31 March 2026.]*
6. W. Schindler: *Overview of AIS 20/31* (NIST/BSI comparison of AIS 20/31 with the SP 800-90 series). NIST RNG workshop presentation (2023).
7. D. Shumow, N. Ferguson: *On the Possibility of a Back Door in the NIST SP 800-90 Dual_Ec_Prng*. CRYPTO 2007 Rump Session (2007).
8. D. J. Bernstein, T. Lange, R. Niederhagen: Dual EC: A Standardized Back Door. In: *The New Codebreakers*, LNCS 9100, pp. 256–281. Springer (2016).
9. R. Impagliazzo, L. A. Levin, M. Luby: Pseudo-random Generation from One-way Functions. In: *STOC 1989*, pp. 12–24. ACM (1989). *[Origin of the Leftover Hash Lemma.]*
10. Y. Dodis, R. Ostrovsky, L. Reyzin, A. Smith: Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM Journal on Computing* **38**(1), 97–139 (2008). *[Average conditional min-entropy and the extraction price.]*

11. M. Baudet, D. Lubicz, J. Micolod, A. Tassiaux: On the Security of Oscillator-Based Random Number Generators. *Journal of Cryptology* **24**(2), 398–425 (2011). *[Canonical Wiener-process jitter model and entropy bound.]*
12. W. Killmann, W. Schindler: A Design for a Physical RNG with Robust Entropy Estimators. In: *CHES 2008*, LNCS 5154, pp. 146–163. Springer (2008).
13. P. Kohlbrenner, K. Gaj: An Embedded True Random Number Generator for FPGAs. In: *FPGA 2004*, pp. 71–78. ACM (2004).
14. M. Bucci, L. Germani, R. Luzzi, A. Trifiletti, M. Varanonuovo: A High-Speed Oscillator-Based Truly Random Number Source for Cryptographic Applications on a Smartcard IC. *IEEE Transactions on Computers* **52**(4), 403–409 (2003).
15. C. S. Petrie, J. A. Connelly: A Noise-Based IC Random Number Generator for Applications in Cryptography. *IEEE Transactions on Circuits and Systems I* **47**(5), 615–621 (2000). *[Amplified-thermal-noise source in silicon.]*
16. C. Tokunaga, D. Blaauw, T. Mudge: True Random Number Generator with a Metastability-Based Quality Control. *IEEE Journal of Solid-State Circuits* **43**(1), 78–85 (2008).
17. I. Vasyltsov, E. Hambardzumyan, Y.-S. Kim, B. Karpinskyy: Fast Digital TRNG Based on Metastable Ring Oscillator. In: *CHES 2008*, LNCS 5154, pp. 164–180. Springer (2008).
18. A. T. Markettos, S. W. Moore: The Frequency Injection Attack on Ring-Oscillator-Based True Random Number Generators. In: *CHES 2009*, LNCS 5747, pp. 317–331. Springer (2009). *[Active attack that synchronises and de-randomises an oscillator TRNG.]*
19. P. Bayon, L. Bossuet, A. Aubert, V. Fischer, F. Poucheret, B. Robisson, P. Maurine: Contactless Electromagnetic Active Attack on Ring Oscillator Based True Random Number Generator. In: *COSADE 2012*, LNCS 7275, pp. 151–166. Springer (2012).
20. Y. Dodis, D. Pointcheval, S. Ruhault, D. Vergnaud, D. Wichs: Security Analysis of Pseudorandom Number Generators with Input: /dev/random is Not Robust. In: *ACM CCS 2013*, pp. 647–658. ACM (2013).
21. N. Ferguson, B. Schneier, T. Kohno: *Cryptography Engineering*, Chapter on the Fortuna generator. Wiley (2010).
22. National Institute of Standards and Technology: *Security Requirements for Cryptographic Modules*. FIPS PUB 140-3 (2019); adopts ISO/IEC 19790:2012 and ISO/IEC 24759:2017.
23. National Institute of Standards and Technology: *CMVP Approved Non-Invasive Attack Mitigation Test Metrics*. NIST Special Publication 800-140F (latest revision).
24. M. Herrero-Collantes, J. C. Garcia-Escartin: Quantum Random Number Generators. *Reviews of Modern Physics* **89**, 015004 (2017).
25. X. Ma, X. Yuan, Z. Cao, B. Qi, Z. Zhang: Quantum Random Number Generation. *npj Quantum Information* **2**, 16021 (2016).
26. C. Gabriel, C. Wittmann, D. Sych, R. Dong, W. Mauerer, U. L. Andersen, C. Marquardt, G. Leuchs: A Generator for Unique Quantum Random Numbers Based on Vacuum States. *Nature Photonics* **4**, 711–715 (2010).
27. J. Y. Haw, S. M. Assad, A. M. Lance, N. H. Y. Ng, V. Sharma, P. K. Lam, T. Symul: Maximization of Extractable Randomness in a Quantum Random-Number Generator. *Physical Review Applied* **3**, 054004 (2015). *[Conditional min-entropy against classical noise; supports Section 3.5.]*
28. ID Quantique: *Quantis QRNG — AIS 31 / SP 800-90B Certification and Stochastic-Model Documentation*. Vendor certification documentation.
29. M. Stipčević, Ç. K. Koç: True Random Number Generators. In: *Open Problems in Mathematics and Computational Science*, pp. 275–315. Springer (2014).