18/09/26

PR-8 derivation → hopefully the last one before the D4-Gate

Prelims based on main, 2026-09-18

Cell d4-family-v1

$R(x) = R_0 + kx^2$ $\quad R_0 = 1 \quad k = 0.1$ — threshold.py ln 72; operating_point.py ln 45

$\mu^* = \sqrt{R_0/k} = \sqrt{10}$ — threshold.py: ln 78

$\bar{R} = R(\mu^*) = 2R_0 = 2$

$s_p^2 = \sigma^2 + \bar{R}$ — threshold.py: 169

$\pi(x) = N(x; \mu^*, \sigma^2$ — // : 176

$\sigma \in$ geomspace$(0.005, 0.7, 15)$ — // : 81

Declare Lattice — // 121

$x_i = \mu^* - 12\sigma + i h_x \quad i = 0..1600 \quad h_x = 24\sigma/1600 = 0.015\sigma$

$y_k = \mu^* - 9 s_p + k h_y \quad k = 0..400 \quad h_y = 18 s_p/400 = 0.045 s_p$

nodes from linspace, endpoints exact — quadrature.py: 139

weights: $w = h$, halved at both ends — // : 76–79

$\Sigma_x f := \Sigma_i w_i f(x_i) \quad \Sigma_y g := \Sigma_k w_k g(y_k)$

Per observation node $y = y_k$

1) $\pi_i = \pi(x_i), \; m_\pi = \Sigma_x \pi_i, \; \hat{\pi}_i = \pi_i / m_\pi$ — gap.py: 208, quadrature.py: 335

2) $\hat{\mu} = \Sigma_x \hat{\pi}_i x_i, \; \hat{s}^2 = \Sigma_x \hat{\pi}_i (x_i - \hat{\mu})^2$ — ladder.py: 131, //: 391–394

$\hat{R} = R(\hat{\mu})$ — // : 132

3) $J_i(y) = \hat{\pi}_i N(y; x_i, R(x_i))$ — gap.py: 125, likelihood.py: 38

$m_J(y) = \Sigma_x J_i(y) =: p^*_{grid}(y)$ — quadrature.py: 335

$\tilde{p}_i = J_i / m_J$

4) $k = \hat{s}^2/(\hat{s}^2 + \hat{R})$ — ladder.py: 101–106

$m_q = \hat{\mu} + k(y - \hat{\mu}), \; v_q = (1-k)\hat{s}^2$

$q_i = N(x_i; m_q, v_q), \; m_q = \Sigma_x q_i$ — ladder.py: 140, quadrature.py: 335

$\tilde{q}_i = q_i / m_q$

5) $KL_{grid(y)} = \Sigma_x \tilde{q}_i \left(\log(\tilde{q}_i) - \log(\tilde{p}_i)\right)$
$= \frac{1}{m_a} \Sigma_x q_i \left(\log(q_i) - \log(J_i) - \log m_a + \log(\text{bin}_w)\right)$

second line is exact algebra on the finite, since $\Sigma_x \tilde{q}_i = 1$

6) value = $\Sigma_y m_J(y_k) KL_{grid(y_k)} \Big/ \Sigma_y m_J(y_k)$
gap.py: 242-250
quadrature.py: 358

7) edge = $\max_{i \in \{0, 1600\}} J_{i(y)} / \max_i J_{i(y)}$
gap.py: 127

Every mass through log sum exp. Direction $KL(\tilde{q} \| \tilde{\pi})$

Exact quantity

$p^*(y) = \int \pi(x) N(y; x, R(x)) \cdot dx$

q_y = Kalman Gaussian from $(\mu^*, \sigma^2, \bar{R})$, not from the grid moments

$KL(y) = \int q_y (\log q_y - \log \pi - \log N(y; \cdot, R)) \cdot dx + \log p^*(y)$

$E = \int p^*(y) KL(y) \cdot dy$

all integrals over $\mathbb{R}$

Target

$\delta_{req} \geq |value - E|$ at each of the 15 spreads

bar: $\delta_{req} \leq T / k_{min} = 5.96e-5$ nats, $T = 5.962e-4$

$k_{min} = 10$

Measured (Result 2026-09-11, declared-sine, two lattices, neither exact)

1e-16 to 5e-16 nats below $\sigma \approx 0.3$, no sign pattern
1.8e-12 nats at $\sigma = 0.7$, steady sign
box mass $\geq 1 - 1.1e-11$ edge ratio $< 1e-16$
ROUNDOFF_FLOOR $= 100 \cdot 2^{-52} \approx 2.2e^{-14}$ relative threshold pg. 87

Step ① The split

For any integrand f on the state axis,

$$E_x f - \int_{\mathbb{R}} f = \underbrace{\left[E_{box} f - \int_{box} f\right]}_{S1} + \underbrace{\left[-\int_{|x-\mu^*|>120} f\right]}_{S3}$$

and on the observation axis,

$$E_y g - \int_{\mathbb{R}} g = \underbrace{\left[\Sigma_{box} g - \int_{box} g\right]}_{S2} + \underbrace{\left[-\int_{|y-\mu^*|>4s_p} g\right]}_{S4}$$

S1 acts on six sums: m_π, $\hat{\mu}$, $\hat{s}^2$, $m_{J(y)}$, m_a, cross sum in ⑤
S2 acts on two: numerator & denominator of ⑥
S3: same six integrands, tails outside the state box, at fixed y

S4: p^{**}'s tail outside the observation box, and (6) then estimates $E[KL \mid y \in box]$

Decision: route for S1 & S2

Route A = Euler-MacLaurin, on the finite box.
Error $\propto h^{2m} \cdot \sup|f^{2m}|$. Needs 2m bounded derivatives

Route B = the exponentially convergent trapezoid rule for analytic integrands (Trefethen & Weideman, SIAM Review 56(3), 2014)
Error $\leq 2M/(e^{2\pi a/h} - 1)$. Needs analyticity in a strip of half-width a & an integral bound M inside it.

State axis (S1): Route B.
Observation axis (S2): deferred to step 2, since $\log \rho^*(\zeta)$ may have no strip
Reason for B on x: $h_x/\sigma = 0.015$, so $e^{-2\pi a/h_x} = e^{-439\, a/\sigma}$
Route A at order $2m$ buys $(0.015)^{2m}$ and needs the $2m$-th derivative of f by hand.

Route B

Let f be analytic in the strip $S_a = \{z : |\mathrm{Im}\, z| < a\}$, $a > 0$

(H1) f analytic on S_a and $f(z) \to 0$ ~~with~~ uniformly as $|\mathrm{Re}\, z| \to \infty$ in S_a

(H2) $\int_{\mathbb{R}} |f(u+iv)| \cdot du \le M$ for every $|v| < a$

then for every $h > 0$ & every shift x_0,

$$\left| h \sum_{j \in \mathbb{Z}} f(x_0 + jh) - \int_{\mathbb{R}} f \right| \le 2M / e^{2\pi a/h} - 1$$

What it's bounding: the ∞ trapezoid sum. The engine's sum is finite:

$$\Sigma_x f = h \sum_{j \in \mathbb{Z}} f(x_0 + jh) - h \sum_{\text{nodes outside box}} f(x_j) + (\text{end-weight correction})$$

The middle term is bounded by S3's tail argument.
The end-weight correction is $h \cdot f(x_0)/2 + h \cdot f(x_{100})/2$, each a value of f at the box edge, 12σ out.
So S1 = [theorem bound] + [outside nodes] + [two edge values]

The integrand, at fixed y, as a function of x

$$f_{y(x)} = \pi(x) \cdot N(y;\, x,\, R(x)) = (2\pi\sigma^2)^{-1/2} e^{-(x-\mu^*)^2/2\sigma^2} \cdot (2\pi R(x))^{-1/2} e^{-(y-x)^2/2R(x)}$$

($J_i = \hat{\pi}_i N(\cdot)$ carries the constant $1/m_\pi$; bound f_y & divide by m_π at the end.)

Check H1 on f_y

Gaussian factor: entire

$R(z) = R_0 + kz^2$ vanishes at $z = \pm i\sqrt{R_0/k} = \pm i\sqrt{10}$

For $z = u + iv$ on S_a: $R(z) = (R_0 + ku^2 - kv^2) + i \cdot 2kuv$

$\mathrm{Re}[R(z)] = R_0 + k(u^2 - v^2) > 0$ whenever $v^2 < R_0/k = 10$

So on S_a with $a < \sqrt{10}$, $R(z)$ lies in the right half-plane, the principal branch of $R(z)^{-1/2}$ is analytic, and $1/R(z)$ is analytic.

Composition: exponential of an analytic function is analytic

$\Rightarrow f_y$ analytic on S_a for $a < \sqrt{10}$

Decay: $|f_y(u+iv)| \le$ [kill from H2 bound] $\to 0$ as $|u| \to \infty$

Check H_2 on f_y: bound $|f_y(u+iv)|$ factor by factor

Notation: $z = u + iv$, $|v| < a < \sqrt{10}$ $\quad d_i := y - u$

$R(z) = A + iB$, $A = R_0 + k(u^2 - v^2)$, $B = 2kuv$

$A \ge R_0 - kv^2 > 0$ on S_a $\quad |R|^2 = A^2 + B^2$

Factor 1, the prior: $\left| e^{-(z-\mu^*)^2/2\sigma^2} \right| = e^{-(u-\mu^*)^2/2\sigma^2} \cdot e^{v^2/2\sigma^2}$

since $(z-\mu^*)^2 = (u-\mu^*)^2 - v^2 + 2i(u-\mu^*)v$ and only the real part survives the modulus

Factor 2, the normaliser:

$$|R(z)|^{-1/2} \le A^{-1/2} \le (R_0 - \kappa v^2)^{-1/2} \le (R_0 - \kappa a^2)^{-1/2}$$

using $|R| \ge \mathrm{Re}[R] = A$

Factor 3, the likelihood exponent. Need a LOWER bound on $Q := \mathrm{Re}[(y-z)^2 / 2R(z)]$.

$$(y-z)^2 = (d - iv)^2 = (d^2 - v^2) - 2i \cdot dv =: p + iq$$

$$p = d^2 - v^2, \quad q = -2dv$$

$$(y-z)^2/R = (p+iq)(A - iB)/|R|^2$$

$$\mathrm{Re}[(p+iq)(A-iB)] = pA + qB = (d^2 - v^2)A - 2dvB$$
$$= (d^2 - v^2)A - 4\kappa u v^2 d$$

So $Q = [A d^2 - 4\kappa u v^2 d - v^2 A] / (2|R|^2)$

Complete the square in d:

$$A d^2 - 4\kappa u v^2 d - v^2 A$$
$$= A(d - 2\kappa u v^2/A)^2 - 4\kappa^2 u^2 v^4 / A - v^2 A$$
$$\ge -(v^2/A) \cdot (4\kappa^2 u^2 v^2 + A^2)$$
$$= -(v^2/A) \cdot (B^2 + A^2)$$
$$= -(v^2/A) \cdot |R|^2$$

Hence $Q \ge -v^2/2A \ge -v^2 / 2(R_0 - \kappa v^2)$, uniformly in u and z

So $|e^{-(y-z)^2/2R(z)}| = e^{-Q}$

$$\le e^{v^2/2(R_0 - \kappa v^2)}$$
$$\le e^{a^2/2(R_0 - \kappa a^2)}$$

Assemble.

$$C(a) := (2\pi(R_0 - \kappa a^2))^{-1/2} \, e^{a^2/2(R_0 - \kappa a^2)}$$

Then for every $|v| < a$, every u, every y:

$$|f_y(u+iv)| \le C(a) \cdot e^{v^2/2\sigma^2} \cdot (2\pi\sigma^2)^{-1/2} \, e^{-(u-\mu)^2/2\sigma^2}$$

which also closes the decay fitting from H1: Gauss in u uniform in v

$$M(v) = \int_{\mathbb{R}} |f_\sigma(u+iv)| \cdot du \le C(a) \cdot e^{v^2/2\sigma^2}$$

$$M = \sup_{|v|<a} M(v) \le C(a) \cdot e^{a^2/2\sigma}$$

Where the a cap comes from: $R_0 - \kappa a^2 > 0$, i.e. $a < \sqrt{R_0/\kappa} = \sqrt{10}$
and $C(a) \to \infty$ as $a \to \sqrt{10}$

Where the cost comes from $e^{a^2/2\sigma^2}$, the prior Gaussian growing off the real axis

verify in σ.

Numbers at the cell ($R_0 = 1$, $\kappa = 0.1$):

$a = 1.58$: $R_0 - \kappa a^2 = 0.750$, $C(a) = 0.461 \cdot e^{1.664} = 2.43$

$a = 10\sigma \le 1.58$ ($\sigma \le 0.158$: $C(a) \le 2.43$, $e^{a^2/2\sigma^2} = e^{50}$ $M \le 1.3e^{22}$

$a = 1.58$, $\sigma = 0.7$: $e^{a^2/2\sigma^2} = e^{2.55} = 12.8$ $M \le 31$

Route B bound on the infinite sum, $2M/(e^{2\pi a/h_{sc}} - 1)$:

$\sigma \le 0.158$: $2\pi a/h_x = 2\pi \cdot 10/0.015 = 4189$,
bound $\le 2.6e^{22} \cdot e^{-4189}$

$\sigma = 0.7$: $2\pi a/h_x = 2\pi\, 1.58/0.0105 = 945$, bound $\le 6^2 \cdot e^{-945}$

Strip width, as a rule of σ

$a(\sigma) = \min(10\sigma, a_{cap})$, $a_{cap} = \sqrt{10}/2 = 1.58$

Reason for the cap: any $a_{cap} < \sqrt{10}$ keeps H1; $\sqrt{10}/2$ keeps $R_0 - \kappa a^2 = 3/4$ so $C(a)$ stays $O(1)$.

Reason for $c = 10$: gain e^{4189}, per spread against cost e^{50}; the cap binds above $\sigma = 0.158$.

11/10/26 Check H2 on the four prior-side sums

Step 1.4: the trapezoid error on the four prior-side sums

The engine computes four sums over the state grid that depends only on the prior or on the polynomial's Gaussian. Each sum stands in for an integral. This section shows that the trapezoid rule's error on each one is too small to matter. It uses the same theorem as the likelihood section.

We're using our Route B which is Trefethen and Weideman, SIAM Review 56(3), 2014, Thm 5.1. Take a function f of a real variable x. Extend it to a complex input $z = u + iv$, where u is the real part and v the imaginary. Fix a width $a > 0$ and call the band $|v| < a$ "the strip". We the define two conditions for the theorem to hold.

1) Smoothness: f is analytic (complex differentiable) at every point of the strip and $f(z) \to 0$ as $|u| \to \infty$ anywhere in that strip.
2) Size: Integrating $|f|$ along any horizontal line inside the strip gives at most some number M

If both hold, the trapezoid sum with spacing h over the whole real line differs from $\int f$ by at most

$$2M/(e^{2\pi a/h} - 1)$$

In english this means (as I interpret it), a function that stays smooth a distance a off the real axis is integrated almost exactly, with error falling like $(e^{2\pi a/h} - 1)$. M measures how much the function grows as you move off the axis. Widening the strip shrinks the exponential and grows

'M'. The strip width $a = \min(10\sigma, 1.58)$ chosen earlier balances the two. We call the H1 & H2 for brevity.

Two choices that we made:
1) Unnormalised sums. The engine divides by 'm_x' to get '$\hat{\mu}$' and '$\hat{s}^2$'. A ratio's error depends on both numerator and denominator so all ratios are handled together in step 2. Meaning each sum is bounded by its own.
2) Order. The sums are bounded in the order (a), (b), (c) (d). Sum (d) needs a lower bound on '$\hat{s}^2$', which comes from (a)-(c) ∴ the argument is not circular, which would suck!

Stating an underpinning fact, plainly: how a Gaussian grows off the real axis.

$$(z-\mu^*)^2 = (u-\mu^*)^2 - v^2 + 2i(u-\mu^*)v$$

$$|e^w| = e^{\mathrm{Re}(w)} \quad \text{for any complex } w$$

$$\Rightarrow |\pi(u+iv)| = \pi(u)\cdot e^{v^2/2\sigma^2}$$

The thing to remember is moving a Gaussian of variance σ^2 a distance of v off the axis multiplies it size by $e^{v^2/2\sigma^2}$. This works for all Gaussians.

Before continuing I want to make explicit the condition for (a)–(d).

Each integrand is a Gaussian times a polynomial. No division, square root, log, or non-integer power appears, so each is analytic at every complex z. The size bounds below are a Gaussian in u times a polynomial, so each integrand tends to 0 as $|u| \to \infty$ throughout the strip.

(a) The prior's total mass, $m_\pi = \sum_{x \in} \pi_i$
Integrand $f(z) = \pi(z)$.

$$\int |\pi(u + iv)| \cdot du = e^{v^2/2\sigma^2} \cdot \int \pi(u) \cdot du$$
$$= e^{v^2/2\sigma^2},$$

$M \leq e^{a^2/2\sigma^2}$ Largest over $|v| < a$

(b) The first-moment sum, $\sum_x x_i \pi_i$
Integrand $f(z) = z\,\pi(z)$. (The numerator of $\hat{\mu}$)

$|z| \leq |u| + |v| < |u| + a$ (triangle inequality)

$U \sim N(\mu^*, \sigma^2)$:
$\int |u| \pi(u) \cdot du = E|U|$
$\leq \mu^* + E|U - \mu^*|$
$= \mu^* + \sigma\sqrt{2/\pi}$

$M \leq (\mu^* + \sigma\sqrt{2/\pi} + a) \cdot e^{a^2/2\sigma^2}$

$E|U - \mu^*| = \sigma\sqrt{2/\pi}$ is the mean absolute deviation of a Gaussian.

(c) The second-moment sum, $\Sigma_x (x_i - \hat{\mu})^2 \pi_i$
Integrand $f(z) = (z - \hat{\mu})^2 \pi(z)$, with $\hat{\mu}$ held fixed as a real number (the numerator of $\hat{s}^2$).

$$|(z - \hat{\mu})^2| = (u - \hat{\mu})^2 + v^2$$
$$\leq (u - \hat{\mu})^2 + a^2$$

$$\int (u - \hat{\mu})^2 \pi(u) \cdot du = \sigma^2 \tau (\mu^* - \hat{\mu})^2 \quad \text{(variance plus squared offset)}$$

$$|\mu^* - \hat{\mu}| \leq 12\sigma$$

$$M \leq (145\sigma^2 + a^2) \cdot e^{a^2/2\sigma^2}$$

Why 12σ: $\hat{\mu}$ is a weighted average of the grid nodes. The weights $w_i \pi_i / m_\pi$ are non-negative and sum to 1, so $\hat{\mu}$ lies inside the grid box, which runs from $\mu^* - 12\sigma$ to $\mu^* + 12\sigma$. The true offset is far smaller and using it would mean assuming $\hat{\mu} \approx \mu^*$ which is part of what we have to prove. The rubbish bound costs a factor of about 145 in M, which changes nothing against e^{-945}.

– – – – – – –

(d) The plug-in posterior's mass, $m_q = \Sigma_x q_i$
Integrand $f(z) = q(z)$, a Gaussian with real mean m_q and variance v_q. The growth fact applied with v_q in place of σ^2.

$$M \le e^{a^2/2v_q}$$

$$v_q = (1-K)\hat{s}^2 = \hat{s}^2\hat{R}/(\hat{s}^2+\hat{R})$$
$$\Rightarrow 1/v_q = 1/(\hat{s}^2 + 1/\hat{R})$$

$$\hat{R} \ge R_0 = 1$$
$$\hat{s}^2 \ge \sigma^2/2 \qquad \leftarrow \text{uses (a)-(c)}$$

$$M \le e^{a^2/\sigma^2} \cdot e^{a^2/2R_0}$$

Why $\hat{R} \ge R_0$: the noise $R(x) = R_0 + Kx_j^2$ is at least R_0 for every real x, $\hat{\mu}$ included. Why $\hat{s}^2 \ge (\sigma^2/2)$: once (a)-(c) and their tail terms are bounded (S3, the part of each integral outside the state box), $\hat{s}^2$ differs from σ^2 by far less than $\sigma^2/2$. This is the only place a later bound leans on an earlier one.