The expiry corner as a logarithmic clock: a boundary-level error law for Carr randomization of the American put

Charlie Yan*

Draft of 29 August 2026

Abstract

Carr's maturity randomization replaces a deterministic expiry by an Erlang clock, turning the American put into a ladder of stationary free boundary problems. Since 1998 the method has been used with Richardson extrapolation applied to the *exercise boundary* — while the expansion justifying that extrapolation has only ever been written down for the price. This paper supplies the missing object. Against a released verifier stack (integral-equation collocation self-converged to 5×10^{-14}, cross-checked by an independent PDE family at 1.3×10^{-7} in price and 4×10^{-6} at the boundary), we measure the boundary error law of the n-stage scheme and find it two-variable: at fixed maturity the convergence is clean $O(1/n)$, the $\ln n$ coefficient bounded by 1.2×10^{-3} and decaying, while the $1/n$ amplitude is corner-amplified — measured as $\sigma\sqrt{\tau}\,\Lambda^{3/2}$ in $\Lambda = \ln(\sigma^2/8\pi r^2 \tau)$ — and changes sign near $\tau \approx 0.06$. The logarithm of Erlangization ticks in $\ln(1/\tau)$, not in $\ln n$. Two independently written complementarity solvers agree on the law to three significant figures or better. Supporting exact structure: Carr's one-stage boundary in closed form, whence the clock law $b_1 = \frac{1}{2}\ln q + \ln(\sigma/\sqrt{2}r)$; a per-stage increment law (one telescoping step an explicitly flagged conjecture with stated numerical support) whose deep-bulk weld $c_0 \to -\frac{1}{2}\ln 2$ reproduces the continuum 8π normalization; and the scheme's exact $q \to \infty$ corner layer, differentially algebraic of order one. A negative result of independent interest: every finite-order corner expansion is differentially algebraic over $\mathbb{C}(\tau)$, so the conjectured hypertranscendence of the boundary cannot be decided by corner asymptotics at any finite order. By-products for practice: a log-basis extrapolation, an explicit corner approximant, a 1.7×10^{-7} desk approximant, and an exact-identity certification harness for arbitrary American put pricers. All claims are re-verified on the released engine; the campaign's erratum ledger, including the claims this audit retired, is published as data.

Keywords: American put; optimal exercise boundary; Carr randomization; Erlangization; Richardson extrapolation; near-expiry asymptotics; free boundary problem; differential algebra; hypertranscendence; reproducibility.

MSC 2020: 91G20, 35R35 (primary); 41A60, 65M12, 12H05 (secondary).

JEL: G13, C63.

*AI-assistance statement: numerics, derivation search, and drafting were carried out with Claude (Anthropic); all results were verified numerically against independently implemented solvers, and the verification scripts are released with this paper.

Code and data: the complete verifier stack, solver lineages, reference tables, ladder archive, campaign reports, and erratum ledger are archived at `github.com/charlieyanhx/carr-boundary-error-law` (§9.4).

1 Introduction

The exercise boundary $B(\tau)$ of the constant-parameter American put is one of the oldest unsolved objects in mathematical finance: sixty years after McKean [35], it has an integral equation [11, 19, 21], a rigorous all-orders near-expiry expansion [12], and no closed form. Carr's maturity randomization [9] — replace the deterministic expiry by an Erlang(n) clock, equivalently run n backward-Euler stages each of which is an exactly solvable stationary problem — is the natural computational and analytical bridge to it, and Carr left the convergence of the scheme as an open question. The first convergence-speed result appeared only in 2025, at the price level and to first order [31], whose author records it as the first study of convergence speed in Canadized option *pricing* — every quantity in it a value, none a boundary.

What happens at the *boundary* level — how fast the free boundary of the randomized problem approaches $B(\tau)$, with what constants, and with what logarithmic fine structure — appears never to have been settled, and the gap is a conspicuous one. Carr himself conjectured joint convergence of the pair (price, boundary) and called a formal proof "an open question"; more strikingly, his Eq. (36) transplants the Richardson weights of his Eq. (35) onto *the critical stock price*, while the Taylor expansion in the mean period length that licenses those weights — his Eq. (33), together with the smoothness argument he gives for it — concerns the put value alone. The boundary error law has thus been *presupposed by practice since 1998 and never derived*. The subsequent rigorous literature stays at the level of the value function: Bouchard et al. [6] prove consistency without an explicit order, Kyprianou and Pistorius [26] defer the convergence of the randomized stopping times to future work, and Kleinert and van Schaik [23] compute boundary points but decline to prove they converge, naming the obstruction — smooth pasting makes the implicit function theorem inapplicable at the contact point. The one rigorous justification of Richardson extrapolation for Carr randomization, Levendorskiĭ [32], concerns prices and sensitivities at an *exogenous* barrier, not a free boundary. Boundary-level convergence has been studied as an object in its own right for other schemes [27], but not for randomization. This paper maps that territory.

The results are several, but the story is one, and the title names its protagonist. The natural expectation — the put's corner is logarithmic, and the scheme's very first rung already sits at the log-anomalous depth $\frac{1}{2}\ln q$ (Corollary 5.3) — was that the boundary error law would carry a $\ln n$; a precursor of this campaign conjectured exactly that (§5.2, retired). The measurement says otherwise, and the mechanism says why: each rung advances by the clock constant plus the logarithm of the ladder's own memory (Proposition 6.2), and the scheme's prefactor welds onto the continuum corner's 8π normalization (verified to ± 0.01; its completion an open lemma), so the would-be $\ln n$ is reallocated into a constant — and the logarithm reappears as maturity structure instead, the corner amplification $\sigma\sqrt{\tau}\,\Lambda^{3/2}$ in $\Lambda = \ln(\sigma^2/8\pi r^2\tau)$. One object therefore runs the paper: the expiry corner is a *logarithmic clock*, $\tau\Lambda'(\tau) = -1$, and that single fact sets the shape of the error law (§§5–6), caps what any numerical campaign can resolve at the corner (the identifiability barrier of §4), and caps what any finite-order corner analysis can decide about transcendence of the boundary (§8). The positive results below are what the clock permits; the negative results are what it forbids; and the surviving routes to Conjecture 8.11 are the ones that read data the clock does not meter — coefficient growth, and the balayage continuation at the perpetual end of the time axis.

Our contributions, in the order they matter.

- **C1** *The boundary-level error law* (§5): the first quantitative account of how $B^{(n)}(\tau)$ approaches $B(\tau)$. It is two-variable. At fixed maturity the scheme is clean first order — the $\ln n$ coefficient is bounded by 1.2×10^{-3} and decays, though at $O(1)$ maturities its $\ln n$ term is a nonnegligible fraction of the $1/n$ amplitude, which is exactly what the log-basis extrapolation of §7 exploits — while the $1/n$ amplitude is corner-amplified, measured as $\sigma\sqrt{\tau}\,\Lambda^{3/2}$, and changes sign near $\tau \approx 0.06$: the logarithm of Erlangization ticks in $\ln(1/\tau)$, not in $\ln n$. Two independently written complementarity solvers agree on it to three significant figures or better. Carr's own closed form for the one-stage boundary (Proposition 5.1) supplies the clock law $b_1 = \frac{1}{2}\ln q + \ln(\sigma/\sqrt{2}r)$, with explicit finite-$q$ corrections rather than a fitted constant.

- **C2** *The per-stage mechanism behind it* (§6): an increment law $\Delta b_m = \beta_c - S(b_m, m) - \frac{1}{2}\ln(2\pi m) + c_0$ on an exact substrate (moment recursion, resolvent action, exact pairing; one telescoping step is an explicit conjecture with 0.6–1.9% support), reverse-Bessel closed forms for the stage kernels, and the deep-bulk weld $c_0 \to -\frac{1}{2}\ln 2$ that reproduces the continuum 8π normalization — which is what explains the missing $\ln n$ term of C1. The scheme's exact $q \to \infty$ corner layer is an explicit parametric curve, differentially algebraic of order one, scoped precisely: the physical corner is its $s \to 0$ endpoint, the interior is Erlang discreteness.

- **C3** *The corner, post-audit* (§4): an explicit uniform crossover law $\xi^2 = \frac{1}{2}(\Lambda + \sqrt{\Lambda^2 + 8})$, solvable in radicals where the uniformly valid approximants already in the literature [12, 13] are transcendental; an exact dictionary onto the Chen–Chadam expansion; a term-by-term accounting that credits every coefficient to prior work, including a priority correction assigning the $2/\Lambda$ term to Goodman and Ostrov [14]; an independent numerical confirmation of the third coefficient; and a quantified identifiability barrier explaining why deeper coefficients are unreachable in principle.

- **C4** *Infrastructure and deliverables* (§§2, 3, 7): a released verifier stack (collocation solver self-converged to 5×10^{-14} at $B(1) = 0.8087505696755$, cross-checked at 1.3×10^{-7} by an independent PDE family, plus a 10^{-13}-grade reference table over 28 dyadic octaves); two exact contact relations powering a certification harness — the curvature identity $V_{SS}(B) = 2rK/\sigma^2 B^2$, which we trace to Goodman and Ostrov [14, eq. (1d)] and restate with proof, and a velocity identity that we have not found in the literature; a log-basis extrapolation replacing the $1/n$ basis; a 1.7×10^{-7} desk approximant; and an exact-identity certification harness applicable to any American put pricer.

- **C5** *A negative result of independent interest* (§8): every finite-order corner expansion is differentially algebraic over $\mathbb{C}(\tau)$ — including the full Chen–Chadam series — so the conjectured hypertranscendence of $B(\tau)$ cannot be decided by corner asymptotics at any finite order. What remains open to attack is coefficient growth, and the branch structure of Hedenmalm's balayage representation [17], which we state as a program.

The body is ordered by dependency rather than by the list above: the setup and verifier stack (§2) and the exact contact relations (§3) come first, then the continuum corner (§4) that the error law of §5 is measured against, the per-stage mechanism (§6), the practitioner deliverables (§7), and the transcendence scope (§8).

§9 is a methods contribution in its own right: the campaign was run under an adversarial audit protocol, and its erratum ledger — every retracted claim, recorded with whether it had been numerically verified at the moment it was asserted — is published as data. Its lesson is sharper than the tally it produces: the one claim that *was* verified when asserted and was retracted anyway fell to a check that was blind by construction, so what predicts survival is not verification but verification against a substrate that does not share the claim's error mode. Reported alongside it is a numerical anomaly the gate found in our own deepest runs, tracked to a double-precision cancellation floor in the contact extraction, with the mechanism and the superseded explanation it replaces.

Related work. The integral-equation characterization of B is Carr et al. [11], Jacka [19], Kim [21], with uniqueness settled by Peskir [38]; the stationary-stage ladder itself coincides with the analytic method of lines [10, 36], which computes the same rungs without the probabilistic reading. Near-expiry asymptotics: Barles et al. [5], Evans et al. [13], Goodman and Ostrov [14], Kuske and Keller [25], Stamicar et al. [44], with Chen and Chadam [12] definitive; our §4 is positioned against it term by term. Randomization asymptotics: Levendorskiĭ [32] (Lévy barrier options, price and sensitivities) is the nearest neighbor; Leduc [31] gives the price-level rate; Bouchard et al. [6] give value-level consistency; Avram et al. [4], Kyprianou and Pistorius [26] and Kleinert and van Schaik [23] develop Canadization in the Lévy setting, the last computing boundary points while explicitly leaving their convergence unproved. Log-corrected discretization rates in adjacent problems: Lamberton's binomial-tree bounds [28, 29] and the near-maturity Lévy boundary laws of Lamberton and Mikou [30] concern a different question (model class or tree discretization; order, not constants) — we discretize the fixed Black–Scholes model in maturity and resolve constants. Transcendence machinery in free-boundary problems appears to be untouched; the difference-Galois tools we show *inapplicable* are Adamczewski et al. [1], Hardouin and Singer [15].

What we do not claim. No closed form for $B(\tau)$; no hypertranscendence proof; no convergence proof of the measured error law (its two open lemmas — the $\gamma \approx 0$ mechanism and the $\Lambda^{3/2}$ amplitude — are stated as such); the kill-cancellation telescoping and the ξ^{-4} completion of the weld remain conjectures with stated numerical support.

2 Setup, scheme, and verifier stack

2.1 Model and boundary

We work in the constant-parameter Black–Scholes model with zero dividends: under the pricing measure, $dS_t = rS_t\,dt + \sigma S_t\,dW_t$, and the American put with strike K and time to expiry τ has value $V(S,\tau) = \sup_{0\le\vartheta\le\tau} \mathbb{E}\big[e^{-r\vartheta}(K - S_\vartheta)^+\big]$. The exercise region is $\{S \le B(\tau)\}$ with $B(0^+) = K$; B solves Kim's integral equation [11, 19, 21], of which it is the unique solution [38], and no closed form is known for constant parameters (time-dependent-parameter constructions that render the Volterra equation solvable exist [22] and delimit, rather than contradict, this statement).

After the scaling $S \mapsto S/K$, $\tau \mapsto \sigma^2\tau$, the boundary depends on the single invariant $\tilde{r} = r/\sigma^2$. Our canonical configuration is $\sigma = 0.2$, $r = 0.05$, $K = 1$ ($\tilde{r} = 1.25$); all quantities below are stated for it unless noted, and §4 treats the $\tilde{r}$-family.

2.2 Carr randomization / Erlangization

Carr's maturity randomization [9] replaces the deterministic expiry τ by an independent Erlang$(n, n/\tau)$ time, equivalently: the value function is computed by n backward-Euler stages in maturity, each a *stationary* free-boundary problem. (The same ladder of stationary problems, read as a numerical method rather than a randomization, is the analytic method of lines of Carr and Faguet [10] and Meyer and van der Hoek [36].) Writing $q = n/\tau$ for the stage rate, stage m solves $(1-q^{-1}\mathcal{L})V_m = V_{m-1}$ on the continuation set with $\mathcal{L}$ the Black–Scholes generator, and produces a boundary $B^{(m)}$. The randomized ("Canadized") boundary at maturity τ is $B^{(n)}(\tau) := B^{(n)}$ at rate $q = n/\tau$. Carr left the convergence rate of the scheme open; the first speed result — linear $O(1/n)$, at the price level — appeared only recently [31], with Levendorskiĭ [32] the nearest methodological antecedent (Carr randomization asymptotics near barriers, in Lévy models, at the price/sensitivity level). This paper works at the *boundary* level, where constants and logarithmic structure become visible.

Corner coordinates. For the corner analysis we use $x = \nu \ln(K/S)$, $\nu = \sqrt{2q}/\sigma$, in which stage m reads $\bigl(-\partial_x^2 + \delta\,\partial_x + 1 + r/q\bigr)v_m = v_{m-1}$, $\delta = 2(r - \frac{1}{2}\sigma^2)/(\sigma^2\nu)$, with obstacle $g(x) = 1 - e^{-x/\nu}$; the stage boundary in these units is b_m, and $\ln(K/B^{(m)}) = b_m/\nu$. The clock scale $\beta_c = \frac{1}{2}\ln q + C$, $C = \ln(\sigma/(\sqrt{2}\,r))$, recurs throughout (§5 derives it exactly at stage one).

2.3 Verifier stack

All numerical claims in this paper rest on a two-family verifier stack, released with the paper (archived at `github.com/charlieyanhx/carr-boundary-error-law`; manifest in §9.4).

Primary: integral-equation collocation. We solve the smooth-pasting (delta-matching) form of Kim's equation,

$$B(\tau)\,\Phi\bigl(d_+(B(\tau)/K,\tau)\bigr) = rK\int_0^\tau e^{-ru}\,\frac{\varphi\bigl(d_-(B(\tau)/B(\tau-u),u)\bigr)}{\sigma\sqrt{u}}\,du, \tag{1}$$

$d_\pm(x,t) = [\ln x + (r \pm \frac{1}{2}\sigma^2)t]/(\sigma\sqrt{t})$, by Chebyshev–Lobatto collocation in $z = \sqrt{\tau}$ in the spirit of Andersen et al. [3], carrying the boundary as $H(z) = \ln^2(K/B(z^2))$, with the quadrature substitution $u = (\sqrt{\tau}\sin\theta)^2$ (which removes both endpoint square-root singularities) and a global Newton iteration on the log-residual. Two failure modes of simpler schemes are part of the record: the naive fixed-point map is *expansive* at small τ (error is amplified by $1/\Phi(d_+)$), and per-node scalar root-finding fails because the log-residual's leading terms cancel as the node value grows (the residual is non-monotone and eventually sign-definite). Self-convergence at $\tau = 1$ is 3.7×10^{-14} between the two finest node counts, and the result is quadrature-independent to all 14 digits shown:

$$B(1) = 0.808\,750\,569\,675\,5 \pm 5 \times 10^{-14}. \tag{2}$$

This confirms, and sharpens by three orders, the reference value $0.8087505696 \pm 3 \times 10^{-10}$ used as the campaign fingerprint.

Cross-check: CN-LCP. An independently implemented Crank–Nicolson solver for the obstacle problem (Rannacher start, mirrored Brennan–Schwartz projection [7, 20], boundary extracted by a $\sqrt{w}$-window fit rather than a contact-point fit) agrees with (2) to 4×10^{-6} at the boundary level — limited by its own contact staircase noise — and, decisively, to 1.3×10^{-7} at the price level, $P_A(K,1) = 0.0609037059$ from (1) versus 0.060903576 (Richardson) from the PDE family, with clean $O(\Delta x^2)$ decay of the PDE error.

A reference table of $B(\tau)$ on the dyadic grid $\tau = 2^0, \ldots, 2^{-28}$ at 10^{-13}-grade accuracy accompanies the released code, together with a method-independent certification harness (§7).

3 Exact relations at the free boundary

Both relations of this section are elementary consequences of value matching and smooth pasting; we record them because they are used repeatedly below and they power the certification harness of §7. The first is *not* new: the contact-curvature identity (3) appears verbatim as eq. (1d) of Goodman and Ostrov [14], who derive it en route to their integral equation and remark on its $\tau \to 0$ limit; we restate it with a proof because the increment-law derivation of §6 and the harness of §7 both consume it. The velocity identity (Proposition 3.2) and the dividend forms (Proposition 3.3) we have not found in the literature; local expansions of the value function at contact are the working currency of the near-expiry matched-asymptotics literature [12, 14, 18], so we claim for them only the statements themselves and their use as pointwise pricer certificates.

Proposition 3.1 (contact curvature; Goodman and Ostrov 14, eq. (1d))**.** *For all $\tau > 0$,*

$$V_{SS}\big(B(\tau)^+, \tau\big) \;=\; \frac{2rK}{\sigma^2 B(\tau)^2}. \tag{3}$$

Equivalently, in $x = \ln(K/S)$ the excess $w = V - (K-S)^+$ has w_{xx}(contact) $= 2rK/\sigma^2$, a τ-independent constant.

Proof. Along the boundary, value matching $V(B(\tau),\tau) = K - B(\tau)$ and smooth pasting $V_S(B(\tau),\tau) = -1$ give $\frac{d}{d\tau}V(B(\tau),\tau) = V_S B' + V_\tau = -B'$, hence $V_\tau \to 0$ at contact. The Black–Scholes equation $V_\tau = \frac{1}{2}\sigma^2 S^2 V_{SS} + rSV_S - rV$ evaluated at $S = B(\tau)^+$ then reads $0 = \frac{1}{2}\sigma^2 B^2 V_{SS} - rB - r(K-B)$, i.e. (3). □

Proposition 3.2 (velocity from the third contact coefficient)**.** *Let $W = V - (K-S)$ on the continuation side, $\kappa = W_{SS}(B^+)$ as in* (3)*, and $c_3 = \frac{1}{6}W_{SSS}(B^+)$. Then*

$$B'(\tau) \;=\; -\,\frac{3\sigma^2 B^2 c_3}{\kappa} \;-\; (r+\sigma^2)B. \tag{4}$$

Proof. W satisfies $W_\tau = \frac{1}{2}\sigma^2 S^2 W_{SS} + rSW_S - rW - rK$ with $W(B) = W_S(B) = 0$. Differentiating smooth pasting along the boundary, $0 = \frac{d}{d\tau}W_S(B(\tau),\tau) = W_{SS}B' + W_{S\tau}$, so $W_{S\tau}(B) = -\kappa B'$. Taking ∂_S of the PDE at contact (the terms containing W and W_S vanish there) gives $W_{S\tau} = \sigma^2 B\,\kappa + \frac{1}{2}\sigma^2 B^2 W_{SSS} + rB\,\kappa$. Equating and solving for B' yields (4). □

The boundary velocity is thus encoded in the *third* Taylor coefficient of the value above the payoff — the lowest coefficient not already pinned by the contact conditions. Numerically (canonical parameters, $\tau = 1$): the CN-LCP solution at $\Delta x = 2.5 \times 10^{-4}$ gives $w_{xx} = 2.500035$ against the exact 2.5 (1.4×10^{-5}; $\leq 1.5 \times 10^{-4}$ across all fit windows), and (4) predicts $B'(1) = -0.044188$ against -0.044112 from central differencing of the collocation boundary (1.7×10^{-3} relative, extraction-noise-limited: window scatter 0.2–1.8% for a third-derivative estimate of an LCP solution).

Proposition 3.1 is the physical-units form of the "contact curvature equals sink strength" identity that drives the increment-law derivation of §6.

Both relations extend to a continuous dividend yield at no cost. We record the forms because the certification harness of §7 is built on them, and because dividend-paying underlyings are the case in which a reference-free boundary check is actually wanted.

Proposition 3.3 (contact relations under a continuous yield)**.** *Let the underlying pay a continuous yield $\delta \geq 0$, so that $V_\tau = \frac{1}{2}\sigma^2 S^2 V_{SS} + (r-\delta)SV_S - rV$. Then for all $\tau > 0$*

$$V_{SS}\big(B^+,\tau\big) = \frac{2\big(rK - \delta B(\tau)\big)}{\sigma^2 B(\tau)^2}, \qquad w_{xx}(\textit{contact}) = \frac{2\big(rK - \delta B(\tau)\big)}{\sigma^2}, \tag{5}$$

and, with κ and c_3 as in Proposition 3.2,

$$B'(\tau) \;=\; -\,\frac{3\sigma^2 B^2 c_3 + \delta}{\kappa} \;-\; \big(r - \delta + \sigma^2\big)B. \tag{6}$$

Both reduce to (3) *and* (4) *at $\delta = 0$.*

Proof. $V_\tau \to 0$ at contact exactly as before, that argument using only value matching and smooth pasting. Evaluating the dividend Black–Scholes equation at $S = B^+$ with $V_S = -1$ and $V = K - B$ gives $0 = \frac{1}{2}\sigma^2 B^2 V_{SS} + \delta B - rK$, which is (5); the log form follows from $w_{xx} = Sw_S + S^2 w_{SS}$ with $w_S = 0$ at contact. For (6), $W = V - (K - S)$ now satisfies $W_\tau = \frac{1}{2}\sigma^2 S^2 W_{SS} + (r-\delta)SW_S - rW - rK + \delta S$, so taking ∂_S at contact adds the constant δ to the right-hand side of the $\delta = 0$ computation, while $W_{S\tau}(B) = -\kappa B'$ as before. □

Remark 3.4 (what the yield costs the check). Two consequences matter for §7. First, (5) is no longer a τ-independent constant: the test becomes a consistency relation between two outputs of the *same* pricer, its boundary and its contact curvature. It remains reference-free, but it is no longer a single number to compare against. Second, it loses discriminating power near expiry when $\delta > r$, since there $B(0^+) = rK/\delta$ and both sides of (5) vanish together; for $\delta \leq r$ no such degeneracy occurs.

Numerically, on the CN–LCP solution at $\Delta x = 2.5 \times 10^{-4}$, $\tau = 1$, canonical σ, r, with the boundary Richardson-extrapolated over two grids: (5) holds to relative 3.0×10^{-6}, 5.5×10^{-4} and 1.4×10^{-4} at $\delta = 0,\ 0.03,\ 0.08$, where the $\delta = 0$ form would be wrong by 84% and 771% at the latter two; (6) holds to 3.2×10^{-3}, 1.6×10^{-2} and 4.7×10^{-3} — third-derivative extraction noise, the same limitation as at $\delta = 0$ — against 15% and 112% for the $\delta = 0$ form. The δ-dependent terms are thus large enough that these are genuine tests of (5)–(6), not of their $\delta = 0$ limits.

4 The near-expiry corner

4.1 Corner variables and the uniform law

Define

$$\xi^2(\tau) \;=\; \frac{\ln^2\big(K/B(\tau)\big)}{\sigma^2 \tau}, \qquad \Lambda(\tau) \;=\; \ln \frac{\sigma^2}{8\pi r^2 \tau}. \tag{7}$$

The near-expiry cusp $\ln(K/B) \sim \sigma\sqrt{\tau \ln(1/\tau)}$ [5, 25, 44] says $\xi^2 \sim \ln(1/\tau)$; the corner refinement is the statement that ξ^2 is a universal function of Λ alone. Balancing the Mills-ratio decay of the survival mass against the ladder drift (the scheme-side derivation is in §6; the balance is the fixed point of $\xi^2 = \Lambda + 2/\xi^2$) and solving the resulting quadratic gives the closed *uniform* form

$$\boxed{\xi^2 \;=\; \tfrac{1}{2}\Big(\Lambda + \sqrt{\Lambda^2 + 8}\Big)} \tag{8}$$

— parameter-free, defined for *all* $\Lambda \in \mathbb{R}$ (in particular through the crossover $\Lambda \approx 0$, i.e. $\tau \approx \sigma^2/8\pi r^2$, where truncated corner series diverge), and reducing to $\xi^2 = \Lambda + 2/\Lambda + O(\Lambda^{-3})$ as $\Lambda \to \infty$.

4.2 Relation to Chen–Chadam: the exact map and a term-by-term accounting

Chen and Chadam [12] proved the corner expansion to all orders. In their variables ($t = \sigma^2\tau/2$, $k = 2r/\sigma^2$, $\alpha = \ln^2(S_f/E)/(4t)$, $\xi_{\mathrm{CC}} = \frac{1}{2}\ln(4\pi k^2 t)$) their Theorems 4.1 and 6.5 give

$$\alpha = -\xi_{\mathrm{CC}} - \frac{1}{2\xi_{\mathrm{CC}}} + \frac{1}{8\xi_{\mathrm{CC}}^2} + \frac{17}{24\xi_{\mathrm{CC}}^3} - \frac{51}{64\xi_{\mathrm{CC}}^4} - \cdots, \tag{9}$$

Their Theorem 4.1 establishes the four-term truncation with a proven $O(\xi_{\mathrm{CC}}^{-4})$ remainder, and their Theorem 6.5 upgrades this to all orders: by induction, the n-term truncation u_n satisfies $|u - u_{n+1}| \le M_{n+1}|\xi_{\mathrm{CC}}|^{-n-2}$ for $\xi_{\mathrm{CC}} \le \xi_{n+1}$, for every n. The exact dictionary to (7) is

$$\alpha = \tfrac{1}{2}\,\xi^2, \qquad \xi_{\mathrm{CC}} = -\tfrac{1}{2}\,\Lambda, \tag{10}$$

under which (9) becomes $\xi^2 = \Lambda + 2/\Lambda + 1/\Lambda^2 - \frac{34}{3}\Lambda^{-3} + \cdots$. Table 1 records the accounting that results.

Table 1: Term-by-term comparison of (8) with the mapped Chen–Chadam series. The leading two orders of our law are re-derivations; the third CC term is *not* captured by (8).

order	CC (proven)	law (8)	verdict
Λ	1	1	re-derivation (SSC 1999 leading order)
Λ^{-1}	2	2	re-derivation; first obtained by Goodman and Ostrov [14]
Λ^{-2}	1	0	CC's $1/(8\xi_{\mathrm{CC}}^2)$; not captured
Λ^{-3}	$-34/3$	-4	subsumed by the Λ^{-2} miss

Priority on the coefficients is not ours and we claim none: the leading order is Stamicar et al. [44], the $2/\Lambda$ term is Goodman and Ostrov [14] — whose eqs. (1f)–(1g) give $b = K\big(1 - \sigma\sqrt{-t\ln(Ctv^2)}\big)$, $C = 8\pi r^2/\sigma^2$, with $v = 1 + 1/\ln(Ct) + O\big(\ln^{-2}(Ct)\big)$, i.e. exactly $\xi^2 = \Lambda + 2/\Lambda + O(\Lambda^{-2})$, five years before Chen and Chadam [12] — a priority that appears to have gone unrecorded in the subsequent literature[1] — and every higher coefficient is Chen–Chadam's.

Nor is the *existence* of a uniformly valid, parameter-free corner law new, and we want to be exact about this because it is easy to overstate. Chen–Chadam's own §7.2 gives an implicit approximant, defined by $\xi_{\mathrm{CC}} = -\alpha - \ln\big\{1 - \frac{1}{2(\alpha+1)} - \frac{1}{2(\alpha+1)^2}\big\}$, which they show is analytic with a unique root for every $\xi_{\mathrm{CC}} \in \mathbb{R}$ and which is more accurate than (8) deep in the corner; Evans et al. [13] give another implicit law of the same genus. Both are transcendental: they define the boundary through an equation that must be solved numerically.

What (8) adds is therefore narrow, and we state it narrowly. (i) It is *explicit*: the balance $\xi^2 = \Lambda + 2/\xi^2$ is a quadratic, so the uniform law is solvable in radicals — evaluable without root-finding and differentiable in closed form, which is what §8 needs when it places the law in a differential-algebraic class, and what makes it usable as the zeroth-order object of the scheme-side analysis. (ii) The derivation route is new: the same balance falls out of the discrete Erlang increment law, welded to the continuum in §6, rather than from the continuum integral equation. (iii) The verification path below is independent of both prior derivations. Measured against

[1] Concretely: the introduction of Chen and Chadam [12] credits prior asymptotics to Kuske and Keller [25], Bunch and Johnson [8] and Stamicar et al. [44] — all leading order — and its reference list does not contain Goodman and Ostrov [14].

Chen–Chadam's implicit law ours is the cruder object, and we recommend theirs wherever corner accuracy is what matters — ours where an explicit differentiable formula is.

Independent numerical verification of the CC coefficients. The 10^{-13}-grade solver of §2.3, swept over $\tau = 2^{-14}, \ldots, 2^{-28}$ ($\Lambda = 9.25, \ldots, 18.96$), gives residuals against the mapped four-term CC form decaying from $+7.3 \times 10^{-3}$ to $\sim 10^{-5}$, and a free fit of $\xi^2 - \Lambda - 2/\Lambda$ on $\{\Lambda^{-2}, \Lambda^{-3}, \Lambda^{-4}\}$ returns leading coefficient 0.88 against the CC value 1 (10–15%, limited by basis collinearity) — an independent numerical confirmation of the third Chen–Chadam coefficient. (A recent high-precision computation of the boundary, by double-exponential quadrature of the early-exercise-premium integral, is given by Sugita and Kobayashi [45]; we have consulted its bibliographic record and published summary only, not the full text, and so do not compare against it here.) Discriminating their ξ^{-4}–ξ^{-6} terms would require $\lesssim 10^{-10}$ relative boundary accuracy at $\tau \leq 2^{-28}$, at the edge of our stack; we do not claim it.

4.3 Scheme-side verification over five certified decades

Independently of the continuum solver, the Erlang ladder itself converges to the refined law: writing $\varepsilon = \xi^2 - \ln(1/\tau)$, measured minus predicted ε stays within ± 0.025 for τ from 2.5×10^{-4} down to 10^{-6} ($q = 10^7, 10^{10}$; grid-refinement-certified on the released engine — refinement fan ≤ 0.001 in ε at $q = 10^7$ and ≤ 0.005 at $q = 10^{10}$), verifying the constant $\sigma^2/(8\pi r^2)$ inside Λ — equivalently the $\sqrt{4\pi k^2 t}$ normalization proven by Chen and Chadam [12] — to $\pm 2.5\%$, by a solver family (backward-Euler ladder in corner coordinates) with no ingredient in common with either continuum method. The two deepest archived rows ($q = 3 \times 10^{13}$, $\tau = 2.5 \times 10^{-10}, 10^{-9}$) reproduce exactly at their archived resolution but are excluded from the certified claim: at that depth the ladder exhibits a reproducible anti-convergence under refinement — traced in §9 to a double-precision cancellation floor in the contact extraction, which caps the usable dx from below rather than above at this intensity.

4.4 Universality across $\tilde{r}$, stated honestly

Across $\tilde{r} = r/\sigma^2 \in \{0.35, 0.625, 1.25, 2.5\}$ (a factor 7.1) and $\tau = 2^{-2}, \ldots, 2^{-12}$, the deviation $\xi^2 - \frac{1}{2}(\Lambda + \sqrt{\Lambda^2 + 8})$ is described by a common smooth curve in Λ at the 2.5×10^{-2} rms level (per-family 1.6–3.2×10^{-2}), decaying with depth; in B-units the uniform law is already a 10^{-5}–10^{-6}-grade approximant for $\tau \leq 2^{-10}$. We emphasize the negative result as well: the residual curves do *not* collapse at the 10^{-3} level in any variable we tested (Λ, ξ^2, or amplified combinations), so universality beyond the shared Λ-trend holds only at the few-percent level at these depths.

4.5 The identifiability barrier

Beyond the $2/\Lambda$ term, corner coefficients accrue information only like $1/\ln(1/\tau)$: even six decades of τ move Λ by only $\ln(10^6) \approx 13.8$, so an $O(\Lambda^{-2})$ term changes by under 2×10^{-3} across the entire measurable window. This is why the ± 0.025 scheme-side verification cannot see CC's third term (it is 0.002 at $\Lambda = 21.7$), why our deepest sweeps resolve it only at 10–15%, and why deeper coefficients are numerically unreachable in principle at fixed precision: the corner is a *logarithmic clock*, and every further coefficient costs an exponential factor in τ-depth. The same barrier reappears at the scheme level in §5.

5 The boundary-level error law for Carr randomization

This section measures, at the boundary level and with a 10^{-13}-grade reference, how the Carr-randomized boundary $B^{(n)}(\tau)$ approaches $B(\tau)$. The state of the art for Carr randomization is an $O(1/n)$ price-level rate [31]; at the boundary level the question has been posed and set aside rather than answered (§1). Of what follows the one-stage statements are exact — (11) recalled from Carr [9], its corner form (12) derived here — and the rest are measurements with quantified systematics (grid Richardson over three spacings; window-fit extraction whose bias is bounded by a three-window scan; domains sized so that image-charge leakage is below e^{-35}; details in §9).

5.1 The one-stage law and its clock constant

Proposition 5.1 (one-stage boundary, closed form; Carr 9)**.** *Let $\theta_+ > 0 > \theta_-$ be the roots of $\frac{1}{2}\sigma^2\theta^2 + (r - \frac{1}{2}\sigma^2)\theta - (r+q) = 0$. The exercise boundary B_1 of the single-stage (exponentially randomized, rate q) put with strike $K = 1$ is determined by*

$$B_1^{\theta_+} \;=\; \frac{|\theta_-|\,r}{r+q+|\theta_-|\,r}, \qquad \text{that is,} \qquad B_1 = \Big(\frac{|\theta_-|\,r}{r+q+|\theta_-|\,r}\Big)^{1/\theta_+} \tag{11}$$

— note that the left-hand side is B_1 raised to the power θ_+, which is what makes the corner form (12) *immediate.*

Proof. On $(B_1, 1)$ the stage equation $\frac{1}{2}\sigma^2 S^2 V'' + rSV' - (r+q)V = -q(1-S)$ has particular solution $q/(r+q) - S$; on $(1,\infty)$ the decaying homogeneous solution is cS^{θ_-}. Imposing C^1 matching at $S = 1$, and value matching plus smooth pasting at $S = B_1$, gives four linear conditions; eliminating the three homogeneous amplitudes yields (11). □

Remark 5.2 (attribution). Equation (11) is not new: it is Carr [9, eq. (18)], whose $S_1 = K\big(pRrT/(\hat{p} - Rp)\big)^{1/(\gamma+\epsilon)}$ coincides with it identically under his dictionary $\gamma \pm \epsilon = \theta_\pm$, $\lambda = q$, $T = 1/q$, $R = q/(q+r)$, $p = |\theta_-|/2\epsilon$ and $\hat{p} = (|\theta_-|+1)/2\epsilon$, which reduce his ratio to $|\theta_-|r/(|\theta_-|r + q + r)$. We restate it with a proof because what this section uses is not (11) itself but its corner form, Corollary 5.3: the depth of the very first rung, in the units in which the rest of the ladder is measured. That form does not appear to have been extracted before, and it is the reason the clock constant below is derived rather than fitted.

Corollary 5.3 (the b_1-law and its clock constant)**.** *In corner units $b_1 = \nu \ln(K/B_1)$, $\nu = \sqrt{2q}/\sigma$,*

$$b_1 \;=\; \tfrac{1}{2}\ln q + \ln\frac{\sigma}{\sqrt{2}\,r} \;+\; O\Big(\frac{\ln q}{\sqrt{q}}\Big), \tag{12}$$

with the $O(\ln q/\sqrt{q})$ term explicit from (11)*.*

The correction term quantitatively closes the gap between the empirical clock constant $C_{\text{self}} = 1.134$ measured at $q = 32$ and the asymptotic $C = \ln(\sigma/(\sqrt{2}r)) = 1.0397$: the predicted finite-$q$ shift is $+0.096$ against the observed $+0.094$, and the residual of (12) against (11) falls to 2.3×10^{-8} by $q = 10^{16}$. The log-anomalous clock — stage one already sits at depth $\frac{1}{2}\ln q$ — is thus derived, not fitted.

5.2 The measured error law at fixed maturity

For τ on the dyadic grid $2^0, \ldots, 2^{-18}$ and $n = 32, \ldots, 4096$ (half-octave), fitting

$$B^{(n)}(\tau) - B(\tau) \;=\; \frac{a(\tau) + \gamma(\tau)\ln n}{n} \tag{13}$$

gives the following (standard errors 1–6 $\times 10^{-5}$ on γ; window systematics $\leq 1.5 \times 10^{-4}$ on $n\,[B^{(n)} - B]$):

τ	1	2^{-2}	2^{-4}	2^{-6}	2^{-8}	2^{-10}	2^{-12}	2^{-14}	2^{-16}	2^{-18}
a	$+.0215$	$+.0098$	$-.0005$	$-.0058$	$-.0071$	$-.0062$	$-.0045$	$-.0031$	$-.0020$	$-.0012$
$10^3\gamma$	1.13	.93	.58	.27	.13	.03	$-.04$	$-.05$	$-.04$	$-.04$

Independent confirmation. Because these two findings are the paper's principal claims, they were re-measured with a second, independently written solver. The ladder above advances each stage by a single mirrored Brennan–Schwartz sweep; the check replaces that by Howard policy iteration, solving each stage's linear complementarity problem to convergence and so assuming nothing about how the contact set moves within a stage (`policy_ladder.py`). It reproduces the closed form (11) to ten digits at $n = 1$, and on $(\tau, n) \in \{1, 2^{-8}\} \times \{64, 256, 1024\}$ the two lineages agree on $n\,[B^{(n)} - B]$ to between 9×10^{-9} and 6×10^{-6} — three to six orders below the quantities being fitted (≈ 0.026 and ≈ -0.0066). Both the sign change of $a(\tau)$ and the corner amplification below therefore survive a change of solver family. We also verified that the single-sweep solution satisfies the complementarity conditions to 2.4×10^{-16}, i.e. it was solving the stage LCP exactly.

Basis discrimination. Over $n \in [32, 4096]$ the elements $\ln n/n$ and $1/n^2$ are notoriously collinear, and a γ of this size could in principle absorb genuine second-order curvature rather than a logarithm. The discriminating evidence is the head-to-head extrapolation test of §7, run on the same campaign data: the three-point extrapolant annihilating $\{1/n,\, \ln n/n\}$ beats the classical one annihilating $\{1/n,\, 1/n^2\}$ by an order of magnitude. Were the second component actually $1/n^2$, the classical extrapolant would annihilate it exactly and win; that it loses, and by the margin a residual logarithm predicts, is the model-selection test the fit alone cannot provide.

Finding 1 (clean first order at fixed τ). The logarithmic anomaly in n is bounded by $|\gamma(\tau)| \leq 1.2 \times 10^{-3}$ throughout $\tau \in [2^{-18}, 1]$, decaying as $\tau \downarrow 0$: at the boundary level and fixed maturity, Carr randomization is $a(\tau)/n$ plus a small but genuinely logarithmic correction. The smallness is in coefficient units, and the relative statement should be made honestly: at $\tau = 1$, $n = 4096$, the fitted $\gamma \ln n$ is 44% of the amplitude (0.0094 against $a = 0.0215$), falling to $\approx 15\%$ by $\tau = 2^{-8}$, while at depth γ sits at the edge of its own standard error ($|\gamma| \approx 4 \times 10^{-5}$ against s.e. 1–6×10^{-5}). The scheme therefore admits the classical Richardson basis — no change of *order* occurs — but the $\ln n$ term is worth annihilating too, which is the log-basis variant of §7 and the source of its order-of-magnitude gain at $O(1)$ maturities.

Finding 2 (the logarithm lives in the corner variable). The amplitude $a(\tau)$ changes sign at $\tau \approx 0.06$ (the randomized boundary sits *above* the true boundary at long maturities and *below* it near expiry) and is corner-amplified:

$$\frac{a(\tau)}{\sigma\sqrt{\tau}} = -0.57,\ -1.00,\ -1.45,\ -1.96,\ -2.56,\ -3.09 \quad \text{at} \quad \Lambda = 5.10, \ldots, 12.03,$$

consistent with $a(\tau) \approx -c\,\sigma\sqrt{\tau}\,\Lambda^{3/2}(1+O(1/\Lambda))$. The raw effective constant $-a/(\sigma\sqrt{\tau}\,\Lambda^{3/2})$, computed directly from the values above, rises $0.050 \to 0.074$ across the window — the drift being the $O(1/\Lambda)$ correction itself. Allowing the one-parameter correction $a = -c\,\sigma\sqrt{\tau}\,\Lambda^{3/2}(1-d/\Lambda)$, adjacent-depth pairs return c falling from ≈ 0.10 at the shallow end to ≈ 0.077 at the deepest pair, so the asymptotic constant is measured only to $c \approx 0.08(1)$; equivalently, in corner units, $n\,[\xi_n^2 - \xi^2] \approx 0.15\,\Lambda^2$ (i.e. $c \approx 0.075$) at accessible depths. The error law of Erlangization is logarithmically corrected through $\Lambda = \ln(\sigma^2/8\pi r^2\tau)$ — the expiry corner as a logarithmic clock — rather than through $\ln n$.

Both the amplitude constant and the exponent refinement beyond $\Lambda^{3/2}$ are left as measurements: the identifiability barrier of §4 applies verbatim (effective exponents drift with Λ at accessible depths), and we do not fit what we cannot pin.

An earlier conjecture, retired. A precursor of this campaign conjectured a fixed-τ law with $\ln n$ coefficient $-\frac{1}{4}$. The measurement refutes this in B-units, $\ln B$-units, and ξ^2-units (where the effective coefficient sweeps from -0.05 to $+0.66$ with depth, stabilizing nowhere); we record the retirement in the erratum ledger (§9). The surviving "$\frac{1}{8}$" of that conjecture admits two candidate origins, neither a stage-one error-law coefficient: Chen–Chadam's $1/(8\xi_{\mathrm{CC}}^2)$ (Table 1), or the level-one balayage constant $\beta_1 = \Gamma(3/2)\lambda_0 = 0.12490\ldots$ of the companion program [46], which is $\frac{1}{8}$ to within 8×10^{-4} but has been tested against exact equality and rejected.

5.3 Why no $\ln n$ at fixed τ: the weld

The vanishing of γ is not an accident. The per-stage increment law of §6 carries the scheme prefactor $-\frac{1}{2}\ln(2\pi m) + c_0(m)$, and its deep-bulk limit $c_0 \to -\frac{1}{2}\ln 2$ turns this into $-\frac{1}{2}\ln(4\pi m)$ — exactly the continuum corner's 8π normalization. The scheme's logarithmic bookkeeping thus *welds* onto the continuum's, and the would-be $\ln n$ mismatch at fixed τ cancels to the order of the weld's accuracy (verified to ± 0.01; its ξ^{-4} completion is open, §6). A derivation of Finding 1 from the weld, and of the $\Lambda^{3/2}$ amplitude in Finding 2, are the two open lemmas this section leaves; a Stefan-type obstruction — the boundary lag is not a local functional of the continuum solution, by the accounting identity of §6 — is why neither is a routine truncation-error computation.

5.4 Poisson closure and its limits

For completeness we summarize the closure structure established on the rung archive (released; see §9): the inner (memory-free) closure has a unique root for all m (no fold; $\partial G/\partial\beta > 1$), tracks the ladder to $\lesssim 2\%$ only for $m \le 8$, is pinned to the ceiling $\hat{B} = 1$ by $m \approx 20$, while the outer law reaches 2% accuracy only for $m \ge 24$ (at $q = 10^6$): stages $m \in [9, 23]$ are a genuine no-man's-land in which the boundary departs from every memoryless closure — the quantitative form of the statement that the lag functional has memory.

6 The uniform increment law and the corner layer

Section 5 established the boundary-level error law for Carr randomization; this section supplies the per-stage mechanism behind it. Its subject is one backward-Euler rung: the exact operator substrate, the collapse of its memory, and the uniform law for the stage-to-stage increment that follows. The direct payoff for §5 is the weld $c_0 \to -\frac{1}{2}\ln 2$ of Section 6.6, which turns the scheme's natural prefactor $-\frac{1}{2}\ln(2\pi m)$ into the continuum $-\frac{1}{2}\ln(4\pi m)$ and so explains why the fixed-τ $\ln n$ anomaly of the error law is a constant reallocation rather than a change of order. Every numerical claim below carries its status; the two statements on which the chain is conditional are Conjecture 6.11 and Remark 6.28.

Because this section carries the campaign's working vocabulary, Table 2 collects it in one place; each term is also defined where it first occurs.

Table 2: Vocabulary and standing notation of §6.

term	meaning
rung	one backward-Euler stage m of the ladder, with boundary b_m in corner units
β_c	the clock scale $\frac{1}{2}\ln q + C$, $C = \ln(\sigma/\sqrt{2}\,r)$: the depth of the first rung
$\eta = e^{-\beta_c}$	corner-normalized sink amplitude; equals the contact curvature at the free boundary (Lemma 6.4)
$\hat{\mu} = m/\beta_c$	layer variable
deck (corner layer)	the stages $m = O(\beta_c)$, i.e. $\hat{\mu} = O(1)$; the *bulk* is $\hat{\mu} \to \infty$ at fixed τ
memory integral I_m	the e^y-moment of the stage time value; carries the ladder's history
pinning	the exact advance identity $b_m = \beta_c + \ln(1 + \eta + I_{m-1})$ (Proposition 6.2)
kill	the Dirichlet condition at the moving boundary inside the stage propagator
first-kick collapse	Conjecture 6.11: the killed all-history sum collapses at order one onto free propagation of the first kick
weld	the deep-bulk limit $c_0 \to -\frac{1}{2}\ln 2$ matching the scheme prefactor to the continuum 8π normalization (§6.6)
pinned-trace feedback	the $s = 1$ trace of the solution re-entering its own forcing; the structural obstruction to a closed form (§6.1, §8)

6.1 The corner ladder and its exact substrate

Here $q > 0$ is the randomization intensity of the Carr scheme [9]: stage m carries $\tau = m/q$, and $B^{(q)}$ is the randomized boundary. In the corner variables of §4, $y = \sqrt{2q}\,x/\sigma$ with $x = \ln(K/\text{spot})$, the rungs b_m are given by $\ln\bigl(K/B^{(q)}(m/q)\bigr) = \sigma b_m/\sqrt{2q}$. The single scheme parameter of the layer is β_c, and the layer variable and scheme-side dictionary are

$$\beta_c = \ln(1/\eta) = \tfrac{1}{2}\ln q + C, \qquad C = \ln\bigl(\sigma/(\sqrt{2}\,r)\bigr), \tag{14}$$

$$\hat{\mu} = \frac{m}{\beta_c} = \frac{q\tau}{\beta_c}, \qquad \xi^2 = \frac{b_m^2}{2m} = \frac{\ln^2\bigl(K/B^{(q)}\bigr)}{\sigma^2\tau}, \qquad \Lambda = \ln\Bigl(\frac{\sigma^2}{8\pi r^2\tau}\Bigr) = 2\beta_c - \ln(4\pi m), \tag{15}$$

with η the corner-normalized rate sink of Definition 6.1. All three identities in (15) are exact consequences of (14) and $\tau = m/q$, with no asymptotics. The *corner layer* (*deck*) is $m = O(\beta_c)$, i.e. $\hat{\mu} = O(1)$; the *bulk* is $\hat{\mu} \to \infty$ at fixed τ. All numbers below use the canonical parameters $\sigma = 0.2$, $r = 0.05$, $K = 1$ ($\tilde{r} = r/\sigma^2 = 1.25$), whence $C = \ln(2\sqrt{2}) \approx 1.0397$; in particular $\ln(2/\pi)$ in Remark 6.31 is $\ln\bigl(\sigma^2/(8\pi r^2)\bigr)$, specific to $\sigma/r = 4$.

Definition 6.1 (stage problem)**.** At leading corner order the time value $h_m(y) \geq 0$ of stage m solves

$$h_m - h_{m-1} = h_m'' + \delta_0(y) - \eta\,\mathbf{1}_{y>0} \quad \text{on } (-\infty, b_m), \tag{16}$$

with $h_m(b_m) = h_m'(b_m) = 0$, $h_m \equiv 0$ on $[b_m, \infty)$, decay as $y \to -\infty$, and $h_0 \equiv 0$; δ_0 is the unit kink source at the payoff corner, $-\eta\mathbf{1}_{y>0}$ the exercise-region sink. The *memory integral* is the e^y-moment $I_m = \int_{-\infty}^{b_m} e^y h_m(y)\,dy$.

Proposition 6.2 (pinning identity)**.** *For every $m \geq 1$, $\eta\, e^{b_m} = 1 + \eta + I_{m-1}$, equivalently*

$$b_m = \beta_c + \ln(1 + \eta + I_{m-1}) : \tag{17}$$

the ladder advances by β_c plus the logarithm of its own accumulated memory.

Proof. Multiply (16) by the adjoint zero mode e^y, which satisfies $(1 - \partial_y^2)e^y = 0$, and integrate over $(-\infty, b_m)$: parts leaves only the boundary flux $-e^{b_m} h_m'(b_m) = 0$, by smooth fit. The kink gives 1, the sink $-\eta(e^{b_m} - 1)$, and the memory I_{m-1} (extending h_{m-1} over the fresh strip (b_{m-1}, b_m), where it vanishes, is free). □

Pinning is the $s = 1$ member of a family of exact moment relations.

Proposition 6.3 (moment recursion)**.** *With $J_m(s) = \int_{-\infty}^{b_m} e^{sy} h_m(y)\,dy$ on the strip of convergence, for $m \geq 1$*

$$(1 - s^2)J_m(s) = J_{m-1}(s) + 1 - \frac{\eta}{s}\big(e^{sb_m} - 1\big), \tag{18}$$

which at $s = 1$ degenerates to (17)*.*

Proof. Multiply (16) by e^{sy} and integrate; the boundary terms vanish by $h_m(b_m) = h_m'(b_m) = 0$, and a direct computation gives the three terms. □

Iterating from $J_0 = 0$ gives the *history representation*

$$J_m(s) = \sum_{k=1}^{m} (1 - s^2)^{-(m-k+1)} \Big[1 - \frac{\eta}{s}\big(e^{sb_k} - 1\big)\Big] : \tag{19}$$

unit kink kicks injected at every past stage, plus sinks $-\eta$ over the historically exercised intervals $(0, b_k)$, each propagated by the free backward-Euler kernel $\widehat{K}_j(s) = (1 - s^2)^{-j}$. Thus I_m is an explicit functional of the history $\{b_k\}_{k \leq m}$ — but one in which, by (17), each b_k is itself a transcendental function of the earlier memory. This closed loop (the *pinned-trace feedback*: the $s = 1$ trace re-enters its own forcing) blocks a closed form for the discrete tower; see Section 6.7 and §8.

Lemma 6.4 (fresh-strip identity)**.** *On the strip (b_{m-1}, b_m) freshly uncovered at stage m, exactly at corner order, $h_m(y) = \eta\big[\cosh(y - b_m) - 1\big]$; in particular $h_m''(b_m) = \eta$.*

Proof. There the memory term vanishes, so (16) reads $h_m'' - h_m = \eta$ with $h_m(b_m) = h_m'(b_m) = 0$, whose unique solution is as stated. □

Remark 6.5 (verification status of the substrate)*.* On the stored Brennan–Schwartz LCP fields of the ladder archive (§9): (17) holds to relative residual 2×10^{-6}; (18) to $\leq 5 \times 10^{-4}$ at $q = 10^4$ and $\leq 2 \times 10^{-4}$ at $q = 10^6$ ($s \in \{0.3, 0.6\}$, $m \leq 40$); the fresh-strip profile to 5×10^{-4}–5×10^{-3} at strip midpoints ($q = 10^4, 10^6$, $m = 10$–40). Each is exact at corner order $O(1/\sqrt{q})$, the accuracy at which (16) itself holds. This is the substrate on which everything below is built.

Lemma 6.4 identifies η with the contact curvature at the free boundary. With Proposition 6.2 it fixes how β_c enters: $\eta = e^{-\beta_c}$ occurs in the substrate *multiplicatively only*, as the sink amplitude in (19) and the prefactor of (17). No mechanism there can put β_c into a Gaussian width; this excludes a priori prefactors $\sqrt{m\beta_c}$ and forces the width bookkeeping of Section 6.5 onto $\sqrt{2\pi m}$ alone.

6.2 The killed resolvent and the exact pairing

Stage m inverts $1 - \partial_y^2$ with a Dirichlet kill at the current boundary. Let $\mathcal{R}_b := (1 - \partial_y^2)^{-1}$ on $(-\infty, b)$, Dirichlet at b, decaying at $-\infty$: self-adjoint, and well defined because the Dirichlet condition removes the would-be kernel element αe^y. Then $h_m = \mathcal{R}_{b_m}\big[h_{m-1} + \delta_0 - \eta \mathbf{1}_{y>0}\big]$ with b_m selected by the Neumann condition $h_m'(b_m) = 0$, and iterating,

$$h_m = \sum_{k=1}^{m} P_{m,k}\big[\delta_0 - \eta \mathbf{1}_{y>0}\big], \qquad P_{m,k} = \mathcal{R}_{b_m}\mathcal{R}_{b_{m-1}}\cdots\mathcal{R}_{b_k}, \tag{20}$$

the backward-Euler propagator killed along the moving boundary. The e^y-projection used above says exactly that e^y is the adjoint zero mode of each stage, Neumann pinning being equivalent to (17). That mode flows elementarily.

Lemma 6.6 (zero-mode flow)**.** *$\mathcal{R}_b\big[e^y\big] = \frac{1}{2}(b-y)e^y$.*

Proof. $\partial_y^2\big[(b-y)e^y\big] = (b-y)e^y - 2e^y$, so $(1-\partial_y^2)\big[\frac{1}{2}(b-y)e^y\big] = e^y$; the left side vanishes at $y = b$ and decays at $-\infty$, and $\mathcal{R}_b$ is uniquely defined by those two conditions. □

By self-adjointness of each factor $\langle e^y, P_{m,k} f\rangle = \langle P_{m,k}^\dagger e^y, f\rangle$ with the boundary order reversed, and Lemma 6.6 shows the adjoint flow stays in $\{W(y)e^y : W \text{ polynomial}\}$: decay excludes the growing admixture e^{-y}, and the kill acts by fixing the constant of integration at each stage's historical boundary. Since $(1-\partial_y^2)(We^y) = -(W'' + 2W')e^y$, this gives:

Lemma 6.7 (adjoint polynomial flow)**.** *Fix m, set $c_j = b_{m-j+1}$, $j = 1, \ldots, m$, and define W_j by $W_0 = 1$ and*

$$W_j'' + 2W_j' = -\,W_{j-1}, \qquad W_j(c_j) = 0. \tag{21}$$

Then $\deg W_j = j$ and $P^\dagger_{m,\,m-j+1}\big[e^y\big] = W_j(y)e^y$ for each j.

Proposition 6.8 (exact pairing)**.** *For every $m \geq 1$,*

$$I_m = \sum_{j=1}^{m} \Big[\, W_j(0) - \eta \int_0^{c_j} W_j(y)\, e^y\, dy \Big]. \tag{P$\star$}$$

Each term is an exponential polynomial in the archived boundary values — polynomial in the b_k and in e^{b_k}, the sink integrals carrying e^{c_j} — so I_m is computable from the boundary history in exact arithmetic.

Proof. Pair e^y against (20), giving $I_m = \sum_k \langle P_{m,k}^\dagger e^y, \delta_0 - \eta \mathbf{1}_{y>0}\rangle$, and apply Lemma 6.7 with $j = m-k+1$: the δ_0 term evaluates $W_j e^y$ at 0, the sink term integrates it over $(0, b_k) = (0, c_j)$. □

Remark 6.9 (verification and base-case anchor). Lemmas 6.6–6.7 and Proposition 6.8 are proved, so their numerical standing measures only the corner reduction: in exact arithmetic on the archived boundaries (P$\star$) reproduces the pinning value $\eta e^{b_{m+1}} - 1 - \eta$ to 0.09–0.7% at $q = 10^6$ and 0.9–3.6% at $q = 10^4$, pure corner-order truncation consistent with $O(1/\sqrt{q})$; against an exactly reconstructed ladder it reproduces pinning to 10^{-88} or better, as it must. At $m = 1$ the pairing gives the exact $I_1 = \big((1+\eta)b_1 - 1\big)/2$ — i.e. $(1+\eta)(b_1-1)/2$ plus exactly $\eta/2$ — whence $b_2 = \beta_c + \ln\big((1+b_1)/2\big) + O(\eta)$, matching the archive to 8×10^{-4} at $q = 10^6$.

6.3 First-kick collapse: the kill-cancellation conjecture

The pairing sums over the whole kick history, and free propagation of old kicks grows geometrically: after the saddle evaluation of Section 6.5, the k-sum in (19) carries an accumulation factor $1/\sigma_*^2$ varying by three to four across the layer. The central discovery of the campaign is that the Dirichlet kill cancels this accumulation *identically at order one*: the killed all-history sum collapses onto the free propagation of the first kick alone.

Definition 6.10 (free kernel and first-kick integral)**.** Let K_m be the free m-step backward-Euler kernel, $\widehat{K}_m(s) = (1-s^2)^{-m}$ in the bilateral convention $\widehat{K}_m(s) = \int e^{+sy} K_m$ (the m-fold convolution of $\frac{1}{2}e^{-|y|}$), and $A_m(b) = \int_{-\infty}^{b} e^y K_m(y)\,dy$. Let $\varsigma(c) = \frac{1}{6}c^3 e^{-c}$ be the fresh-strip term, the leading part of the exact strip integral $\eta e^b \chi(\Delta b)$ with $\chi(c) = \frac{c}{2} + \frac{1}{4}(1-e^{-2c}) - (1-e^{-c}) = \frac{c^3}{6} + O(c^4)$ furnished by Lemma 6.4.

Conjecture 6.11 (kill cancellation; first-kick collapse)**.** *In the corner layer $\hat\mu = m/\beta_c = O(1)$, $\beta_c \to \infty$, there is a correction $\mathcal{E}(\hat\mu;\beta_c)$, bounded on the deck and vanishing as $\beta_c \to \infty$, with*

$$I_m = A_m(b_m)\,\exp\big[\mathcal{E}(\hat\mu;\beta_c) + \varsigma(\Delta b_m)\big]; \tag{22}$$

that is, the killed all-history sum collapses at order one onto the free propagation of the first kick. Equivalently: in the killed pairing (P$\star$)*, per unit age the free growth $(1-\sigma_*^2)^{-1}$ of older-kick contributions is exactly offset by the survival cost of one further stage under the descending kill boundary, so the sum over all kicks and sinks telescopes onto its first term.*

Remark 6.12 (numerical status of Conjecture 6.11). *Order one.* Established numerically in exact arithmetic (the pairing at 60 significant digits, against exact contour quadrature of A_m), over three decades of q and the full deck $\hat\mu \in [1.3,\, 5.4]$:

q	$I_m/A_m(b_m)$ over the full deck	flatness
10^4	$1.111 \to 1.153$	$\pm 1.9\%$
10^5	$1.086 \to 1.109$	$\pm 1.0\%$
10^6	$1.071 \to 1.083$	$\pm 0.6\%$

No order-one σ_*-dependence survives: I_m/A_m stays within 2% of a constant and tends to 1 as β_c grows, while the uncancelled accumulation factor $1/\sigma_*^2$ varies across the deck by a factor three to four (≈ 3.2–3.5 at deck entry to ≈ 10–13 at the outer edge).

The remainder. Write $R := \ln(I_m/A_m) - \varsigma(\Delta b_m)$, one function of $\hat\mu$ at each q; writing $R = G\,\beta_c^{-2}$ normalizes it to $G \approx 3.0$–5.0 on the deck, with fit $G \approx 2.50 + 0.45\,\hat\mu$ — a normalization and a fit that Remark 6.13 later supersedes on both counts, and which should not be quoted from here. The β_c-scaling is certified less strongly than the collapse. An independent adversarial reconstruction over *five* decades ($q = 10^3$–10^7, $\beta_c = 4.49$–9.10) finds $R\,\beta_c^2$ drifting upward monotonically by 20–25%; power-law fits give $p \approx 1.69$–1.78 in $R \propto \beta_c^{-p}$ across the deck (fixed-$\hat{\mu}$ decade ratios 1.37–1.40 and 1.30–1.34, against 1.45, 1.37 for exact β_c^{-2} and 1.20, 1.17 for β_c^{-1}), and a free fit $R = c_1/\beta_c + c_2/\beta_c^2$ returns $c_1 \approx +0.13$–$+0.18$. So the remainder decays like β_c^{-p}, $p \approx 1.7$–1.8, at accessible depths; pure $1/\beta_c$ is decisively excluded; the limiting exponent 2 of (22) is not excluded asymptotically (the drift is compatible with a subleading transient) but is *not certified*, so the β_c^{-2} attribution there and in (32) is a hypothesis, not a measurement. (On the LCP archive $R\,\beta_c^2$ had appeared decade-stable to $\pm 3\%$; that archive's $O(q^{-1/2})$ corner-reduction error decays far faster than any β_c^{-p} and can flatten the apparent scaling — the exact ladder is authoritative.) G is measured, not derived; its candidate origin is the $O(\Delta b^2)$ frontier near-cancellation plus boundary-deceleration terms, both formally next-order in the telescoping.

Remark 6.13 (the collapse, reproduced without the archive). The order-one collapse can now be measured on a ladder generated by the closed forms of §6.4 themselves, as follows: $I_0 = 0$, $b_m = \beta_c + \ln(1 + \eta + I_{m-1})$, and I_m from the pairing (P$\star$) with W_j built from the Lagrange kernel (23). No stored archive and no complementarity solver enter at any point, and the ladder satisfies pinning to 4×10^{-37} by construction. Evaluating $I_m/A_m(b_m)$ over the deck reproduces the table above:

q	m-range used	$I_m/A_m(b_m)$	against the archive table above
10^4	8–30	$1.111322 \to 1.153439$ ($\pm 1.86\%$)	$1.111 \to 1.153$ ($\pm 1.9\%$)
10^5	9–30	$1.085847 \to 1.104276$ ($\pm 0.84\%$)	$1.086 \to 1.109$ ($\pm 1.0\%$)
10^6	11–30	$1.070674 \to 1.078746$ ($\pm 0.38\%$)	$1.071 \to 1.083$ ($\pm 0.6\%$)

The lower endpoints agree to every digit quoted; the shortfall at the upper end of the last two rows is truncation of our m-range at 30, short of the $\hat{\mu} = 5.4$ deck edge ($m = 37$ and $m = 43$ respectively), not disagreement.

This bears on provenance. The order-one cancellation is the central empirical claim of this subsection, and Table 3 records the ladder archive as the one [RECONSTRUCTION-ONLY] input on which it rested. It no longer rests on it: the same numbers follow from (23), (P$\star$) and (17) by a route sharing no code, no data and no discretization with the archive — indeed sharing no floating-point PDE solve at all.

On $G(\hat{\mu})$ the agreement is weaker, and we record that too. On this ladder $R\beta_c^2$ runs from 1.92 at $\hat{\mu} = 0.76$ to 4.29 at $\hat{\mu} = 3.27$, concave in $\hat{\mu}$ rather than linear: its size matches the 3.0–5.0 quoted above across the deck proper, but the linear deck fit $G \approx 2.50 + 0.45\hat{\mu}$ is not reproduced, a least-squares fit to these values giving $2.20 + 0.64\hat{\mu}$ with visible curvature. G remains measured rather than derived, and the shape of that measurement should be regarded as open.

The same ladder settles the decay *exponent* more sharply than the archive can. Interpolating $R = \ln(I_m/A_m(b_m)) - \varsigma$ to fixed $\hat{\mu}$ and fitting $R \propto \beta_c^{-p}$ across five decades, $q = 10^3$–10^7 ($\beta_c = 4.49$–9.10), gives

$$p = 1.719, \qquad 1.736, \qquad 1.795 \qquad \text{at} \quad \hat{\mu} = 2,\ 3,\ 4,$$

inside the 1.69–1.78 band reported in Remark 6.12 from an independently constructed exact ladder, and excluding $p = 2$. The consecutive-decade slopes at $\hat{\mu} = 2$ are $1.677,\ 1.745,\ 1.730,\ 1.721$: flat across a doubling of β_c, with no systematic drift toward 2. This bears on the one escape route left open for the β_c^{-2} normalization of (22). Remark 6.12 attributes the measured $p \approx 1.7$–1.8 partly to a substrate effect — the archive's $O(q^{-1/2})$ corner-reduction error decaying faster than any β_c^{-p} and so flattening an apparent scaling — and leaves $p \to 2$ open as a subleading transient. Neither explanation is available here: this ladder satisfies pinning to 4×10^{-37} and carries no corner-reduction error at all, and the exponent shows no drift. We therefore record β_c^{-2} as a form that this measurement does not support; what is measured, on a substrate free of the confound, is β_c^{-p} with $p \approx 1.72$–1.80.

This is why (22) is stated with an unspecified $\mathcal{E}$. An earlier form of the conjecture asserted $\mathcal{E} = G(\hat{\mu})/\beta_c^2$ *with no term of order* $1/\beta_c$, and that clause is contradicted rather than merely uncertified. Forcing $\mathcal{E} = G/\beta_c^2$ requires $R\beta_c^2$ to be constant in q; on this ladder it rises monotonically from 3.12 to 3.82 across the five decades at $\hat{\mu} = 2$, a drift of 22% — independently reproducing the 20–25% reported in Remark 6.12. A two-term fit $\mathcal{E} = c_1/\beta_c + c_2/\beta_c^2$ instead gives $c_1 = 0.162$, $c_2 = 2.41$ at $\hat{\mu} = 2$ and $c_1 = 0.155$, $c_2 = 2.94$ at $\hat{\mu} = 3$, with worst relative residual 1.7%; the c_1 so obtained agrees with the $+0.13$–$+0.18$ found independently in Remark 6.12. Two exact ladders of different construction therefore concur that the $1/\beta_c$ coefficient is nonzero and near 0.16. What survives unambiguously, and is what (22) now asserts, is the order-one collapse itself: $I_m/A_m(b_m) \to 1$ with the deck flatness tabulated above. The decay law of the correction is a measurement, currently β_c^{-p} with $p \approx 1.72$–1.80 or equivalently a $1/\beta_c$ term with coefficient ≈ 0.16, and it is not derived. Correspondingly the $G(\hat{\mu})/\beta_c^2$ appearing in (32) should be read as that same measured correction, not as a certified β_c^{-2} law.

Remark 6.14 (why the cancellation resists direct asymptotics)*.* Two quantified mechanisms defeat termwise saddle or Euler–Maclaurin treatment of the history sum at prefactor order. (i) *Marginal memory*: the per-stage decrement of the sink sum $g(\sigma) = \sigma\Delta b + \ln(1-\sigma^2)$, at the saddle $\sigma = s_*$, has magnitude only 0.014–0.05 near the outer deck edge, so the effective memory is 20–70 stages, comparable to the deck itself; no few-term truncation exists. (ii) *Discreteness poles*: the zeros σ_n of $1 - e^{-g(\sigma)}$ satisfy $\operatorname{Re}\sigma_n \approx -(2/\Delta b)\ln(2\pi n/\Delta b)$, hence sit near the imaginary axis, and the continuum replacement of the stage sum fails exactly at the order at which $\sqrt{2\pi m}$ lives.

The series A and the flow it carries are in turn explicit, through a kernel that is elementary.

Proposition 6.15 (the adjoint kernel in closed form)**.** *Write $v = 1 - \sqrt{1-x}$ and $E_n(y) := [x^n]e^{-vy}$. Multiplying (21) by x^j and summing, polynomiality of the W_j forces $\sum_j W_j(y)x^j = A(x)e^{-vy}$ with $A(x) = \sum_j W_j(0)x^j$ (the e^{-2y} characteristic branch is excluded), so $W_j(y) = \sum_{i=0}^{j} A_i E_{j-i}(y)$ with $A_i = W_i(0)$, and the conditions $W_j(c_j) = 0$ become the triangular recursion $A_j = -\sum_{i<j} A_i\, E_{j-i}(c_j)$, $A_0 = 1$. Then for $n \geq 1$ and $1 \leq r \leq n$,*

$$[x^n]v^r = \frac{r}{n}\,\frac{1}{2^{\,2n-r}}\binom{2n-r-1}{n-1}, \qquad E_n(y) = \sum_{r=1}^{n} \frac{(-y)^r}{(r-1)!\,n}\,\frac{1}{2^{\,2n-r}}\binom{2n-r-1}{n-1}, \tag{23}$$

and, as $n \to \infty$ at fixed y,

$$E_n(y) = -\frac{y\,e^{-y}}{2\sqrt{\pi}}\; n^{-3/2}\bigl(1 + O(1/n)\bigr). \tag{24}$$

Proof. v satisfies $v(2-v) = x$, i.e. $v = x\varphi(v)$ with $\varphi(\lambda) = (2-\lambda)^{-1}$, so Lagrange inversion gives $[x^n]v^r = \frac{r}{n}[\lambda^{n-r}](2-\lambda)^{-n} = \frac{r}{n}2^{-n}\cdot 2^{-(n-r)}\binom{2n-r-1}{n-r}$, which is (23) after $\binom{2n-r-1}{n-r} = \binom{2n-r-1}{n-1}$; the expression for E_n follows from $e^{-vy} = \sum_r (-y)^r v^r/r!$ and $r/r! = 1/(r-1)!$. For (24) write $e^{-vy} = e^{-y}e^{y\sqrt{1-x}}$. The only singularity on $|x| = 1$ is the square-root branch point at $x = 1$, with singular expansion $e^{-y}\{1 + y(1-x)^{1/2} + \frac{1}{2}y^2(1-x) + \frac{1}{6}y^3(1-x)^{3/2} + \cdots\}$; the analytic terms contribute nothing and transfer of the leading singular term, $[x^n](1-x)^{1/2} \sim n^{-3/2}/\Gamma(-\frac{1}{2}) = -1/(2\sqrt{\pi}\, n^{3/2})$, gives the stated form. $\square$

Remark 6.16 (what the factor e^{-y} is). The kernel carries e^{-y}, evaluated in the flow at $y = c_j = b_{m-j+1}$, so every term is weighted by e^{-b_k} — which pinning (17) identifies as $\eta/(1+\eta+I_{k-1})$. The exponential suppression that Conjecture 6.11 asserts is therefore visible in the kernel rather than imposed on it: it is the sink amplitude $\eta = e^{-\beta_c}$, entering through the depth of the boundary. The transfer asymptotic (24) holds for $y = o(\sqrt{n})$ only; on the deck E_n changes sign along $y \approx 0.66\, n$, and its asymptotics uniform across that exchange — two saddles, obtained after a substitution that removes the branch cut entirely — are derived in the companion paper [47] (in preparation), verified to $O(1/n)$ over two decades of y/n and through the sign change itself; no claim of the present paper depends on that uniform form.

Remark 6.17 (discharge routes). Two routes to a proof are on record; both are developed in the companion paper [47] (in preparation), and we summarize their state here — for orientation only, since the one companion result this paper re-uses, the archive-free reproduction of the collapse, is carried out independently in Remark 6.13. *Adjoint reciprocity* is now explicit and exact: two integrations by parts against the flow equation close the pairing (P$\star$) on the *past* pinning identities, converting I_m into a linear recursion in its own history $\{I_{m-j}\}$ whose kernel is the boundary-flux sequence $W_j'(c_j)$, with forcing supported on the fresh strips — verified there in exact arithmetic, end to end against the archive, and shown to be *generative*: iterated with pinning it reproduces the ladder, and the collapse table above, with no solver and no archive (whence Remark 6.13). The flux kernel is elementary — reverse-Bessel coefficients through the kernel $E_n(y) = [x^n]e^{-(1-\sqrt{1-x})y}$ of Proposition 6.15 — and admits two-saddle asymptotics uniform across its sign change, verified to $O(1/n)$. What remains is a single uniform estimate on that recursion; a shortcut that fails (treating boundary motion as a perturbation of the frozen flux $-u_j$) is documented there. *Dual summation*: Poisson summation over the discreteness poles σ_n converts the marginal stage sum into a rapidly convergent dual sum. Any proof must reproduce the measured facts: cancellation of $1/\sigma_*^2$ at order one, and a remainder of shape $-\frac{1}{2}\delta_{\mathrm{pref}}$ (Section 6.5) plus a flat part with the size and decay of Remark 6.13.

6.4 The collapsed kernel in closed form: reverse Bessel polynomials

Under Conjecture 6.11 the analytic content sits in A_m, which is elementary.

Proposition 6.18 (closed form for K_m and A_m). *Let $\theta_n(x) = \sum_{i=0}^{n} \dfrac{(n+i)!}{i!\,(n-i)!\,2^i} x^{n-i}$ be the reverse Bessel polynomial of degree n. Then $K_m(y) = e^{-|y|}\theta_{m-1}(|y|)/\big((m-1)!\,2^m\big)$, so $e^y K_m(y)$ is a polynomial on $y > 0$ and, for $b \geq 0$ (always the case on the deck), A_m is an elementary polynomial of degree m:*

$$A_m(b) = \frac{1}{4^m}\binom{2m-1}{m} + \sum_{i=0}^{m-1} \frac{(m-1+i)!\, b^{m-i}}{i!\,(m-1-i)!\,(m-1)!\,(m-i)\,2^{m+i}}, \tag{25}$$

with leading term $b^m/(m!2^m) = \beta^m/m!$ in the half-height variable $\beta := b/2$ and constant term $A_m(0) = \int_{-\infty}^{0} e^y K_m = \binom{2m-1}{m}/4^m$. (For $b < 0$ the $y < 0$ branch of K_m applies and $A_m(b)$ is e^{2b} times a polynomial; (25) is the $b \geq 0$ evaluation.) Moreover, exactly and pointwise for all $m \geq 2$ and all b,

$$A_{m-1}(b) = e^b\big[K_m(b) - K_m'(b)\big]. \tag{26}$$

Proof sketch. Invert $\widehat{K}_m(s) = (1-s^2)^{-m}$ for $y > 0$ by the order-m residue at $s = 1$ (the pole branch) and use symmetry in y; the residue produces exactly the reverse Bessel coefficients, and integrating $e^y K_m$ term by term gives (25), its constant term by the hockey-stick identity. For (26), apply $(1-\partial_y^2)K_m = K_{m-1}$ and integrate once by parts in the definition of A_{m-1}. $\square$

Remark 6.19 (verification). These checks test algebra, not modelling, so carry no corner-order error. Against exact contour quadrature at archived (m, b), $q = 10^6$, the worst relative difference for (25) is -6×10^{-12} (at $m = 10$; -2×10^{-19} at $m = 24$, -4×10^{-28} at $m = 43$) — quadrature artefacts, the identity itself holding below 10^{-60} for $m \leq 24$, $b \in \{0, 0.7, 6.1, 19.4\}$. Identity (26) is exact to 50 digits for $m = 2, \ldots, 25$, both signs of b.

Proposition 6.20 (frozen-b generating function)**.** *For fixed b and $w = \sqrt{1-x}$,*

$$\sum_{m \geq 1} A_m(b)\, x^m = \frac{(1+w)e^{(1-w)b} - 2w}{2w} =: \mathfrak{A}(x, b), \tag{27}$$

and, consistently with Lemma 6.7, $\partial_b \mathfrak{A} = x e^{(1-w)b}/(2w) = e^b \sum_m x^m K_m(b)$: the kill-free adjoint flow reproduces A_j term by term.

Remark 6.21 (what does *not* close). (27) closes only for frozen argument b. The tower itself — b_m generated by $\eta e^{b_m} = 1 + \eta + A_{m-1}(b_{m-1})e^{\mathcal{E}+\varsigma}$ — mixes the index into the argument, and no closed equation for $\sum_m e^{b_m} x^m$ results. The obstruction is the pinned-trace feedback of Section 6.1; its consequences for closed-form solvability and difference-Galois criteria are taken up in §8, and in the continuum limit of Section 6.7 it degenerates to an algebraic constraint and dissolves.

Remark 6.22 (the substitution $u = e^b$). The natural linearizing substitution $u_m = e^{b_m}$ makes pinning *affine*, $\eta u_m = 1 + \eta + I_{m-1}$, but the logarithm reappears at once in the flow: every pairing term is an exponential polynomial in (b_k, e^{b_k}) jointly, so the substitution moves the transcendence between the two halves of the substrate — pinning natural in e^b, the flow in b — rather than removing it. Any linearizing substitution must act on both halves simultaneously, which the exponential does not. Details, with the worked $j = 1$ case, are in the companion paper [47].

6.5 The increment law

The saddle asymptotics of A_m are classical and unconditional; we record them with the exact width identity that fixes the prefactor class.

Definition 6.23 (Legendre rate and saddle)**.** For $b > 0$, $m \geq 1$ define the Legendre rate of $(1-\partial_y^2)^{-m}$ and the *endpoint–width defect*

$$S(b, m) = \max_{s \in (0,1)} \big\{ sb + m \ln(1-s^2) \big\} = s_* b + m \ln(1 - s_*^2), \qquad \frac{2ms_*}{1-s_*^2} = b, \tag{28}$$

$$\delta_{\mathrm{pref}}(s) = \tfrac{1}{2} \ln \frac{2(1+s^2)}{(1+s)^2} \geq 0, \qquad \delta_{\mathrm{pref}}(1) = 0, \tag{29}$$

with saddle $s_* = s_*(b, m) \in (0, 1)$, explicitly $s_* = \big[-\hat{\mu} + \sqrt{\hat{\mu}^2 + 4\hat{B}^2}\big]/(2\hat{B})$ in the scaled variables of Section 6.7; nonnegativity in (29) is $2(1+s^2) - (1+s)^2 = (1-s)^2 \geq 0$.

Proposition 6.24 (saddle asymptotics of A_m; exact width identity)**.** *As $m \to \infty$ with s_* in a compact subset of $(0,1]$, $A_m(b) = e^{\,b-S(b,m)}\big(1+O(1/m)\big)/\big(\sqrt{2\pi m}\, e^{\,\delta_{\rm pref}(s_*)}\big)$, the prefactor arising from the exact identity*

$$(1-s_*)\sqrt{2\pi E''(s_*)} = \sqrt{2\pi m}\,\frac{\sqrt{2(1+s_*^2)}}{1+s_*} = \sqrt{2\pi m}\; e^{\,\delta_{\rm pref}(s_*)}, \qquad E''(s) = \frac{2m(1+s^2)}{(1-s^2)^2}: \quad (30)$$

the endpoint factor $1/(1-\sigma)$ produced by the e^y-projection cancels the $(1-s_*^2)^{-2}$ divergence of the spatial kernel width E'', and what survives is the arrival-stage (Stirling) width m — the fluctuation of which stage's *kick reaches the boundary, not the spatial spread. On the pole branch $s_* \to 1$ the prefactor is exactly $\sqrt{2\pi m}$, since $\delta_{\rm pref}(1) = 0$.*

Both equalities in (30) hold to 10^{-61}, and the asymptotic is verified against exact contour quadrature to 0.1–0.9% on the deck, its relative error times m being constant to three digits over $m = 10$–400. Combining Proposition 6.2 (which gives $\Delta b_m = \beta_c + \ln I_m - b_m + O(1/I_m)$, the error being $O(\eta)$ on the deck, where $I_m \gg 1$), Conjecture 6.11 and Proposition 6.24 yields the law.

Theorem 6.25 (uniform increment law; conditional on Conjecture 6.11)**.** *In the corner layer $\hat\mu = O(1)$, $\beta_c \to \infty$,*

$$\Delta b_m = \beta_c - S(b_m, m) - \tfrac{1}{2}\ln(2\pi m) + c_0(\hat\mu; q), \qquad (31)$$

with the explicit, parameter-free correction

$$c_0 = -\,\delta_{\rm pref}(s_*) + \varsigma(\Delta b_m) + \mathcal{E}(\hat\mu; \beta_c). \qquad (32)$$

Remark 6.26 (measured status of the law). Unconditionally, independent of Conjecture 6.11: over 89 deck stages at $q = 10^4$–10^6 the raw constant $c_0 := \Delta b_m - [\beta_c - S] + \frac{1}{2}\ln(2\pi m)$ satisfies $|c_0| \le 0.08$; a deficit-model comparison selects the prefactor $\frac{1}{2}\ln m + 0.9165$ at RMS 0.023, every tested alternative ($\frac{1}{2}\ln(m\beta_c)$, spatial widths, etc.) being at least 2.3× worse, and 0.9165 agrees with $\frac{1}{2}\ln 2\pi = 0.9189$ to 0.3%. With (32) the law closes *pointwise* to ± 0.01 across the deck and three decades of q. The prediction that the correction decays faster than $1/\beta_c$ holds; the stronger claim that the $1/\beta_c$ coefficient vanishes exactly is *not* certified, the measured exponent being $p \approx 1.7$–1.8 (Remark 6.12), so (32) fixes the amplitude while its exponent stays hypothesized. Five pre-registered falsification specifications at $q = 10^7$ ($\beta_c \approx 9.10$) — mean $R = \mathrm{mean}[c_0 + \delta_{\rm pref} - \varsigma] = 0.048 \pm 0.006$ over $\hat\mu \in [1.3, 5]$; slope $+0.0067 \pm 0.0015$ per unit $\hat\mu$; raw $|c_0| \le 0.1$; the 10^7 curve below the 10^6 curve by 0.005–0.012 on $\hat\mu \in [2,5]$; decade ratio $(\beta_c(10^6)/\beta_c(10^7))^2 = 0.76$ — were run in the independent verification on the exact reconstructed ladder. The first four pass (measured 0.0498, +0.0070, 0.053, $-0.008\ldots-0.010$); the fifth, the only one probing the decay exponent, comes out at 0.79, outside the β_c^{-2} specification and exactly where $p \approx 1.75$ puts it. Size and shape of the law are thus confirmed at 10^7; the exponent attribution is revised as in Remark 6.12.

On the inner (pole) branch the law must reproduce, with no spare constant, the singulant normalization of the inner region established in §5. It does, and exactly: the reduction does not use Conjecture 6.11, because as $s_* \to 1$ the kill goes inactive (the boundary is quasi-static, the Dirichlet constants in (21) stop biting), the W-flow values revert to the Poisson terms $W_j(0) \to \beta^j/j!$, and the all-age truncated exponential sum is restored in place of the first-kick collapse.

Proposition 6.27 (pole-branch reduction to the inner singulant; exact)**.** *Let $\beta = b/2$, $Q(m,\beta) = \Gamma(m,\beta)/\Gamma(m)$ the regularized upper incomplete gamma function, $\varphi(m,\beta) = m\ln(m/\beta) - (m-\beta)$ the Stirling (singulant) exponent. On the inner side $I_m = e^{\beta_m} Q(m,\beta_m)\big(1+O(1/\beta_c)\big)$, $\beta_m = b_m/2$, and: (i) pinning gives $b_{m+1} = \beta_c + \beta_m + \ln Q(m,\beta_m)$, whose stationary point ($\Delta b \to 0$, $b = 2\beta$) is exactly the inner closure $\beta = \beta_c + \ln Q(m,\beta)$; (ii) with $\lambda = \beta/m$, the root $\beta_{\rm root}$ of the closure approaches the ceiling β_c according to*

$$\big(\beta_c - \beta_{\rm root}\big)(1-\lambda)\sqrt{2\pi m}\, e^{\varphi} \;\longrightarrow\; 1 + O(1/m); \tag{33}$$

(iii) on the pole branch $S(b,m) = \varphi(m,b/2) + b/2 + o(1)$ and $\delta_{\rm pref}(s_*) \to 0$, so exponent and prefactor of (31) continue analytically into the $(\sqrt{2\pi m}, e^{\varphi})$ normalization of the inner region term by term, with no spare constant.

Proof sketch. (i) restates (17), with $e^{\beta} Q(m,\beta) = \sum_{j<m} \beta^j/j!$ the truncated exponential (Poisson) sum. (ii) inserts into the closure the geometric tail estimate $1 - Q = \big(\beta^m e^{-\beta}/m!\big)\big(1 + O(1/m)\big)/(1-\lambda)$ and Stirling's formula $\beta^m e^{-\beta}/m! = e^{-\varphi}\big(1+O(1/m)\big)/\sqrt{2\pi m}$. (iii): at the saddle (28) with $s_* = 1 - \varepsilon_*$ the condition gives $\varepsilon_* \simeq m/b$, whence $S = b - m + m\ln(2m/b) + o(1) = \varphi(m,b/2) + b/2 + o(1)$, and $\delta_{\rm pref}(1) = 0$. □

Identity (33) is the hard gate of the derivation: it is the inner-region normalization obtained independently in the error-law analysis of §5, and (31) reproduces it exactly. Its left-hand side evaluates to 0.9514 at $m = 20$ and 0.9823 at $m = 30$, converging to 1 within the $O(1/m)$ bound (observed somewhat faster, since $\lambda = \beta/m$ shrinks simultaneously). The $m = 1, 2$ base cases are reproduced by the exact machinery (Remark 6.9); the collapsed leading recursion, being a deck asymptotic, is *not* valid at $m = 1$, and seeding it there shifts b_2 by $\approx +0.155$ (1.6% at $q = 10^6$) — an $O(1/\beta_c)$ error of exactly the expected size for a deck formula used outside the deck.

6.6 Welding to the continuum: the 8π constant

Deep in the bulk ($\hat\mu \to \infty$ at fixed τ) the scheme converges to the continuum corner law $\xi^2 = \Lambda + 2/\xi^2$, i.e. the refined law (8) of §4 [12], whose constant is the 8π inside $\Lambda = \ln\big(\sigma^2/(8\pi r^2\tau)\big)$. Ladder runs at $(q,m) = (10^7, 10^4)$ and $(10^{10}, 10^4)$ — i.e. $\tau = 10^{-3}, 10^{-6}$ — converge to (8) to $|\xi^2 - \ln(1/\tau) - \varepsilon_{\rm pred}| \le 0.025$ across the five decades of τ certified in §4, verifying the constant $\sigma^2/(8\pi r^2)$ to $\pm 2.5\%$ and the coefficient 2 of the $1/\xi^2$ term to roughly $\pm 10\%$. The increment law welds onto (8) through its constant: in the far bulk

$$c_0(\text{bulk}) = \frac{1}{2}\Big[-\ln 2 + \frac{2}{\xi^2} + 2\frac{db}{dm} - \frac{\xi^4}{4m}\Big] + o(1) \;\longrightarrow\; -\tfrac{1}{2}\ln 2, \tag{34}$$

where $\xi^4 := (\xi^2)^2$ and db/dm is the per-stage increment of b in the bulk. This predicts $c_0 = -0.181/-0.249/-0.280$ at the three runs against measured $-0.17\ldots-0.18$, -0.247, -0.291 (agreement ± 0.01). Hence the scheme prefactor welds to the continuum one, $-\tfrac{1}{2}\ln(2\pi m) + c_0 \to -\tfrac{1}{2}\ln(4\pi m)$, and since $\Lambda = 2\beta_c - \ln(4\pi m)$ exactly, the ladder reproduces the continuum 8π. The bookkeeping: the arrival-stage Stirling width supplies the 2π; the welded constant $e^{-2c_0} = 2$ doubles it to the 4π inside $\ln(4\pi m)$; the remaining factor 2 in $8\pi = 4\pi \cdot 2$ enters through the $\sqrt{2}$ in $C = \ln\big(\sigma/(\sqrt{2}r)\big)$, via $2\beta_c = \ln\big(q\sigma^2/(2r^2)\big)$. The deep-bulk limit $-\tfrac{1}{2}\ln 2$ is thus the entire content of the constant separating the scheme's natural prefactor from the continuum's — and the reason the fixed-τ $\ln n$ anomaly of §5 is a constant reallocation, not a change of order.

Remark 6.28 (open: the ξ^{-4} completion). Formula (34) is verified numerically to ±0.01 across the three runs, but its ξ^{-4} completion — that the bracket, including the $-\xi^4/(4m)$ term and the implied $O(\xi^{-4})$ tail, is the full next-order expansion — is *unproven*; the derivation is open. The residual ±0.02 in the convergence to (8) is likewise consistent with the $O(\xi^{-4})$ truncation of the reference formula plus the discretization variable $z = \ln^3(1/\tau)/m \le 0.3$ and the scheme-layer term $\beta_c^2/m \le 0.016$ of these runs, but that apportionment is not derived.

6.7 The corner-layer curve at $q \to \infty$

At fixed $\hat{\mu}$ the recursion rides the quasi-static manifold obtained by setting the increment exponent to zero at leading order, and that manifold is an elementary curve. Scale $\hat{B} = b/(2\beta_c)$, so the ceiling $b = 2\beta_c$ (the stationary rung) is $\hat{B} = 1$.

Proposition 6.29 (exact $q \to \infty$ corner-layer curve; leading order, conditional on Conjecture 6.11 at order one)**.** *As $\beta_c \to \infty$ at fixed $\hat{\mu}$, the ladder concentrates on the parametric curve*

$$\hat{\mu}(s) = \left[\frac{2s^2}{1-s^2} + \ln(1-s^2)\right]^{-1}, \qquad \hat{B}(s) = \hat{\mu}(s)\frac{s}{1-s^2}, \qquad s \in (0,1), \tag{35}$$

with $b \simeq 2\beta_c\hat{B}$, the parameter s being the Legendre saddle s_* itself. Along the curve $d\hat{B}/d\hat{\mu} = -\ln(1-s^2)/(2s)$ with $s = \left[-\hat{\mu} + \sqrt{\hat{\mu}^2 + 4\hat{B}^2}\right]/(2\hat{B})$, so s is algebraic in $(\hat{B}, \hat{\mu})$. The quasi-static identity $2\hat{B}s + \hat{\mu}\ln(1-s^2) = 1$, which holds identically along (35)*, removes the logarithm already at first order: it gives $s = 1/\big(2(\hat{B} - \hat{\mu}\hat{B}')\big)$, and the saddle constraint $\hat{B}s^2 + \hat{\mu}s - \hat{B} = 0$ then yields the degree-five* first-order *algebraic ODE*

$$4\hat{B}\big(\hat{B} - \hat{\mu}\hat{B}'\big)^2 - 2\hat{\mu}\big(\hat{B} - \hat{\mu}\hat{B}'\big) - \hat{B} = 0. \tag{36}$$

Alternatively, one further differentiation removes the logarithm through the explicit second-order algebraic ODE

$$\big(\hat{\mu}\hat{B}' - \hat{B}\big)^4 = \hat{\mu}^2\hat{B}^2\big(\hat{B}''\big)^2\big(\hat{\mu}^2 + 4\hat{B}^2\big), \qquad ' = d/d\hat{\mu}. \tag{37}$$

The curve is therefore D-algebraic of order one (a fortiori of order two). As an explicit function of $\hat{\mu}$ it is implicit-elementary (the inverse of an elementary map), not expressible in the named special functions attempted. Endpoints: $s \to 0$ gives $\hat{B} = \sqrt{\hat{\mu}}$ (the outer matching), and $\hat{B} = 1$ is reached at $s = 0.7422$, $\hat{\mu} = 0.6052$ — bare deck entry; the measured *deck entry at finite q is shifted by the prefactor and is q-dependent, $\hat{\mu}_* = 1.50$, 1.32, 1.20 at $q = 10^4, 10^5, 10^6$.*

Remark 6.30 (validation at finite β_c). Against the archive at $q = 10^6$, the offsets $\hat{B}_{\text{curve}} - \hat{B}_{\text{data}} = +0.343, +0.416, +0.483, +0.568$ at $\hat{\mu} = 2.0, 3.0, 4.0, 5.4$ stand to the predicted finite-β_c linearized offset $(\beta_c - S)/(2s_*\beta_c)$ in ratios $0.925, 0.912, 0.903, 0.893$ — deviation from 1 of second order in the linearization, well inside the 15% acceptance threshold, and reproduced to the last quoted digit by the independent verification. That verification also produced (37), checked it along the curve to residual 4.5×10^{-60} ($s = 0.12$–0.96), and excluded by an SVD rank test over 60 curve points any first-order algebraic relation $P(\hat{\mu}, \hat{B}, \hat{B}') = 0$ of total degree ≤ 4 — consistent with (36), which has degree five. The $e^{\mathcal{E}}$ correction of (22) is benign here: $\mathcal{E} \to 0$ on the deck, it enters below the $-\frac{1}{2}\ln(2\pi m)$ prefactor (itself $O(\ln\beta_c/\beta_c)$ relative to the exponent), and it vanishes in the limit — so (35) is the exact $q \to \infty$ corner object whatever the decay rate of $\mathcal{E}$ turns out to be. The curve result is in this respect independent of the exponent question of Remark 6.13.

Remark 6.31 (scope: what the curve is, and is not). Proposition 6.29 is *scheme asymptotics*: the exact $q \to \infty$ shape of the Carr boundary's traversal of the layer $m = O(\beta_c)$, for use in discretization analysis. It is *not* a closed form for the continuum boundary $B(\tau)$, for a structural reason: fixed physical τ has $\hat{\mu} = q\tau/\beta_c \to \infty$, so the entire physical corner maps to the single endpoint $s \to 0$, under $s \simeq \sqrt{\beta_c/(q\tau)} = \xi/\sqrt{2m}$ — a scheme quantity vanishing identically at any fixed physical location as $q \to \infty$. There $\hat{\mu} = s^{-2} - \frac{3}{2} + O(s^2)$ and $\hat{B}^2/\hat{\mu} = 1 + s^2/2 + O(s^4)$ give

$$\xi^2_{\text{curve}} = 2\beta_c + \frac{\beta_c}{\hat{\mu}} + O\Big(\frac{\beta_c}{\hat{\mu}^2}\Big) = 2\beta_c + \frac{\beta_c^2}{m}, \tag{38}$$

precisely the leading continuum cusp $\ln\big(K/B(\tau)\big) \simeq \sigma\sqrt{\tau \ln(1/\tau)}$ in the matched limit [12] — a validation of the machinery — plus scheme-layer bookkeeping ($2C + \ln m$, β_c^2/m). The curve's *interior* (s bounded away from 0) never touches fixed-τ physics: it encodes Erlang stage discreteness, the resolvent-tail-to-Gaussian crossover of $(1-s^2)^{-m}$ in the layer $m = O(\beta_c)$, and carries no continuum content beyond the leading cusp. Quantitatively, at the three bulk runs of Section 6.6 the bare curve's endpoint overshoots the continuum ξ^2 by $+\ln(4\pi m) + \beta_c^2/m$, measured $+11.75$, $+11.76$ and $+12.85$: all the refined structure of (8) lives at the ladder's prefactor level, which the curve, leading order by construction, lacks. On the measured side the cusp ratio $\xi^2/\ln(1/\tau)$ runs $0.9776 \to 0.9811 \to 0.9818$ over $\tau = 2.5 \times 10^{-4} \to 10^{-9}$ (the first from an auxiliary run at $\tau = 2.5 \times 10^{-4}$, $m = 2.5 \times 10^3$, $q = 10^7$, whose own overshoot is $+10.39$), the deviation tracking $\big(\ln(2/\pi) + 2/\xi^2\big)/\ln(1/\tau)$ to one or two digits, within the ± 0.025 band of Section 6.6 — the scheme converging to the continuum law, not the curve describing it. Finally, the low-order D-algebraicity above belongs to the scheme curve alone: the coordinate map degenerates at the corner ($s \equiv 0$ there), so no D-algebraicity statement about $B(\tau)$ transfers in either direction; see §8.

Remark 6.32 (summary of statuses). Proved exactly: Propositions 6.8, 6.18, 6.20, 6.15, Lemmas 6.6, 6.7. In the companion paper [47] (in preparation, hence quoted for orientation and load-bearing for nothing here): the reciprocity telescoping, the summed linear recursion, and the kernel's uniform two-saddle asymptotics (the last verified to $O(1/n)$ over more than two decades of y/n). Proved at corner order: Propositions 6.2, 6.3, 6.24, 6.27, Lemma 6.4.

Conjectural with quantified numerical support: the first-kick collapse (Conjecture 6.11), whose order-one content is confirmed to 0.4–1.9% across the deck and, in Remark 6.13, independently of the ladder archive; on it Theorem 6.25 and the leading order of Proposition 6.29 are conditional. Also conjectural: the ξ^{-4} completion of the bulk weld (Remark 6.28; verified ± 0.01, derivation open), and the deck-exponent statement that sink and kill contributions do not shift the deck exponent (supported at ± 0.08 by the raw c_0 measurement, subsumed in Conjecture 6.11).

Measured, not derived — and now known *not* to follow a β_c^{-2} law: the correction $\mathcal{E}(\hat{\mu}; \beta_c)$ of (22). At accessible depths the data give β_c^{-p} with $p \approx 1.72$–1.80, equivalently a $1/\beta_c$ coefficient near 0.16, on two exact ladders of independent construction (Remarks 6.12, 6.13). The earlier deck fit $G \approx 2.50 + 0.45\hat{\mu}$ is superseded: on the archive-free ladder a fit gives $2.20 + 0.64\hat{\mu}$ with visible curvature, so the *shape* of $\mathcal{E}$, and not merely its amplitude, should be treated as open.

7 Practitioner deliverables

Four artifacts fall out of the analysis; all ship with the released code.

7.1 Log-basis Richardson extrapolation

The measured error structure (13) makes the three-point extrapolant that annihilates $\{1/n, \ln n/n\}$ the natural upgrade of classical Richardson for Canadized prices and boundaries. On the campaign data it beats the same-cost classical three-point extrapolant (annihilating $\{1/n, 1/n^2\}$) exactly where practitioners live:

$$\begin{array}{llll} \tau = 2^{-8},\ n = 32: & \text{raw } 5.1 \times 10^{-5} & \to\ \text{classic } 9.8 \times 10^{-7} & \to\ \text{log-basis } 5.3 \times 10^{-8}\ (18\times) \\ \tau = 1,\ n = 32: & \text{raw } 2.1 \times 10^{-4} & \to\ 9.6 \times 10^{-6} & \to\ 7.9 \times 10^{-7}\ (12\times) \end{array}$$

with the advantage decaying into the extraction-noise floor by $n \gtrsim 512$. Honest summary: up to an order of magnitude at $n \le 256$ stages, free of charge. This comparison also carries analytic weight: it is the basis-discrimination test of §5.2 — a genuine $1/n^2$ second component would be annihilated exactly by the classical extrapolant, which instead loses by the margin a residual logarithm predicts.

Scope. Richardson extrapolation of Carr randomization is not new: Carr's own Eq. (36) extrapolates the critical stock price, and Levendorskiĭ [32] justifies extrapolation to arbitrary order for prices and sensitivities of barrier options with an exogenous barrier. What was missing is the expansion the boundary extrapolant presupposes. §5 supplies it, and the basis correction here follows from it: the $\ln n$ term the classical $1/n$-basis ignores is small but not zero, and annihilating it is what buys the order of magnitude above.

7.2 Parameter-free corner approximant

$B \approx K \exp\big(-\sigma\sqrt{\tau\, \xi_u^2(\Lambda)}\big)$ with ξ_u^2 from (8): zero constants, all maturities, exact through $2/\Lambda$ (Table 1). Accuracy in B: $\sim 6 \times 10^{-3}$ at $\tau = 1$, improving to 10^{-5}–10^{-6} grade by $\tau \le 2^{-10}$ ($\tilde{r} = 1.25$). In the deep corner, prefer Chen–Chadam's fourth-order resummed approximant; ours is the crossover-uniform option.

7.3 Desk approximant

Adding one calibrated rational correction,

$$B_{\text{desk}} = K \exp\Big(-\sigma\sqrt{\tau\,[\xi_u^2(\Lambda) + \delta(\Lambda)]}\Big), \qquad \delta = \frac{p_0 + p_1\Lambda + p_2\Lambda^2}{1 + p_3\Lambda + p_4\Lambda^2 + p_5\Lambda^3 + p_6\Lambda^4},$$

with the seven constants released alongside the code, achieves

$$\max_{\tau \in [2^{-28},\, 1]} \big|B_{\text{desk}} - B\big| = 1.7 \times 10^{-7} \qquad (\tilde{r} = 1.25).$$

The calibration is per-$\tilde{r}$: applied unmodified at $\tilde{r} = 0.625$ or $\tilde{r} = 2.5$ the error degrades to 3–4 $\times 10^{-3}$ (the universality floor of §4); recalibration is a single least-squares fit against ~ 25 solver calls.

7.4 Certification harness

A method-independent test suite for any American put pricer, built from exact identities only: (C1) the contact curvature (3) in log-coordinates, $w_{xx} = 2rK/\sigma^2$, a pointwise check requiring no reference solution — and, on a dividend-paying underlying, its yield form $w_{xx} = 2(rK - \delta B)/\sigma^2$

of (5), still reference-free but now a consistency check between the pricer's own boundary and its own curvature (Remark 3.4); (C2) the velocity identity (4), with (6) under a yield; (C3) the delta-matching residual (1) evaluated on the *candidate's own boundary curve* — a functional certificate. Calibration of discriminating power: the CN-LCP solver passes C1 at $O(\Delta x^2)$ grade ($4.6 \times 10^{-4} \to 3.8 \times 10^{-5}$); the collocation boundary passes C3 at 10^{-7}–10^{-11}; a deliberately 0.2%-shifted boundary is flagged by C3 with defects 1.6×10^{-2}–2.3×10^{-1}, three-plus orders above the pass floor.

8 What corner asymptotics cannot decide

The conjecture left standing by the preceding sections is that $B(\tau)$ is *hypertranscendental* over $\mathbb{C}(\tau)$: that it satisfies no nonzero algebraic differential equation with rational-function coefficients. This section is a bounded negative result plus a signpost. A classical meta-theorem — everything built from $\mathbb{C}(\tau)$ by finitely many algebraic operations, exponentials, logarithms and compositions is differentially algebraic (Theorem 8.4) — transported through the Chen–Chadam dictionary of §4.2 yields the new content here: *no finite-order corner analysis can witness hypertranscendence* (Corollary 8.5), which tells the field where not to spend effort. Theorem 8.4 claims no novelty; it repackages standard facts about logarithmico-exponential towers together with composition-closure [2, 24, 33, 41]. We then record briefly the obstruction that blocks the difference-Galois machinery of Hardouin and Singer [15] and Adamczewski et al. [1] at finite randomization intensity — the *pinned-trace feedback* of Remark 8.7 — and note that it dissolves in the continuum without leaving a decision behind. Finally, in one compact subsection labeled throughout as program and conjecture, we point at the route that survives both no-go results: the branch structure of the maximal analytic continuation of the Laplace transform of the boundary gap in Hedenmalm's balayage equation [17]. Throughout, ξ^2 and Λ are the corner variables of (7); note the elementary but load-bearing relation

$$\tau\,\Lambda'(\tau) \;=\; -1, \tag{39}$$

an algebraic ODE of order one: the logarithmic clock is itself differentially algebraic.

8.1 Log towers are differentially algebraic

We work with germs at $\tau = 0^+$. A germ f is *differentially algebraic* (D-algebraic) over a differential field K if $P(f, f', \ldots, f^{(n)}) = 0$ for some nonzero polynomial P over K; otherwise it is *hypertranscendental* over K; the classical theory is surveyed by Rubel [43]. The lemmas below are pure differential algebra, stated over an arbitrary differential field of characteristic zero; Hardy fields — fields of real-valued germs at 0^+, closed under differentiation, in which every nonzero element is eventually nonvanishing and hence of constant sign — enter only in Definition 8.1, where they are load-bearing: by the extension theorems [42] (in essence Hardy's logarithmico-exponential orders of infinity [16]) such a field may be enlarged by any element algebraic over it, by e^f, and by $\log f$ for f eventually positive, the resulting germs again being nonoscillating.

Definition 8.1 (corner tower class)**.** A *tower over* $\mathbb{C}(\tau)$ is a finite chain of Hardy fields of germs at $\tau = 0^+$, $\mathbb{R}(\tau) = F_0 \subset \cdots \subset F_M$ with $F_{i+1} = F_i(g_i)$, where each g_i is either *(a)* algebraic over F_i, or *(b)* $g_i = e^f$ with $f \in F_i$, or *(c)* $g_i = \log f$ with $f \in F_i$ eventually positive. The *corner tower class* $\mathcal{T}$ is the union of all tops F_M; concatenating towers shows $\mathcal{T}$ is a field. (Complex constants are admitted via $F_i \otimes_{\mathbb{R}} \mathbb{C}$, a field since -1 is not a square in a Hardy field, though the complexification is only a differential field. In the composition statements below the inner germ may be based at $+\infty$ rather than 0^+ — interchanged by $w \mapsto 1/w$ — since natural inner germs such as Λ diverge.)

The class $\mathcal{T}$ contains every corner approximant of the American put boundary we are aware of: Λ, the iterated logarithms $\ln \Lambda, \ln\ln \Lambda, \dots$, all algebraic functions of these — in particular the uniform law (8) — and the exponentials converting log-moneyness back to a price; the Kuske–Keller, Bunch–Johnson and Stamicar–Ševčovič–Chadam approximants [8, 13, 25, 44] are all of this form. Membership is strictly stronger than D-algebraicity and is not automatic: $\mathcal{T}$ is closed under algebraic operations, exp and log, but *not* under solving an implicit transcendental equation, so an approximant defined implicitly (Lambert-W type, say) need not lie in $\mathcal{T}$ — though it will still be D-algebraic if it satisfies an algebraic ODE, which is all Corollary 8.5 requires. Below, $K\langle f\rangle = K(f, f', f'', \dots)$ for f in a differential extension of K.

Lemma 8.2. *f is D-algebraic over K if and only if $\operatorname{trdeg}_K K\langle f\rangle < \infty$. More generally, every element of a differential subfield M with $\operatorname{trdeg}_K M < \infty$ is D-algebraic over K.*

Proof. If $\operatorname{trdeg}_K K\langle f\rangle \le n$ then $f, \dots, f^{(n)}$ are algebraically dependent over K, a D-algebraic relation. Conversely take $P(f, \dots, f^{(n)}) = 0$ with $P \ne 0$, n minimal, and among order-n relations take P of minimal degree in the top variable Y_n. Minimality of n forces $\deg_{Y_n} P \ge 1$, so the separant $S = (\partial P/\partial Y_n)(f, \dots, f^{(n)})$ evaluates a nonzero polynomial of smaller Y_n-degree and therefore, by minimality of the degree, $S \ne 0$ — hence invertible, being a nonzero element of a field. Differentiating, $0 = P_\tau + \sum_{i\le n} (\partial P/\partial Y_i) f^{(i+1)}$, so $f^{(n+1)} \in K(f, \dots, f^{(n)})$; by induction that field is differential and contains every $f^{(k)}$, so $\operatorname{trdeg}_K K\langle f\rangle \le n+1$. Finally $f \in M$ gives $K\langle f\rangle \subseteq M$, of finite transcendence degree. $\square$

Lemma 8.3 (closure properties; classical)**.** *Let Ω be a differential field of characteristic zero containing $\mathbb{C}(\tau)$. The elements of Ω that are D-algebraic over $\mathbb{C}(\tau)$ form a differential subfield of Ω: they are closed under the field operations and under* (i) *differentiation. Moreover:* (ii) *every element algebraic over a differential subfield of finite transcendence degree over $\mathbb{C}(\tau)$ is D-algebraic;* (iii) *if $f \in \Omega$ is D-algebraic and $g \in \Omega$ satisfies $g' = f'g$ (an exponential of f) or $g' = f'/f$ with $f \ne 0$ (a logarithm of f), then g is D-algebraic; and* (iv) (composition) *if g is D-algebraic over $\mathbb{C}(w)$ as a function of one variable w, h is D-algebraic over $\mathbb{C}(\tau)$ and nonconstant, and the composite $g \circ h$ is defined, then $g \circ h$ is D-algebraic over $\mathbb{C}(\tau)$.*

These are classical. Part (iv), composition-closure, follows from the transcendence-degree criterion of differential algebra [24, 41]; a modern statement with proof and an order bound is Ait El Manssour et al. [2, Prop. 3.7], and it is recorded as a remark in Lipshitz and Rubel [33].[2] We prove them because the transcendence-degree bookkeeping is what Corollary 8.5 consumes.

Proof. Write $K = \mathbb{C}(\tau)$, $d_i = \operatorname{trdeg}_K K\langle f_i\rangle < \infty$. Items (i)–(iii) are Lemma 8.2 applied to a differential field of finite transcendence degree built in each case: the compositum $M = K\langle f_1\rangle K\langle f_2\rangle$, with $\operatorname{trdeg}_K M \le d_1 + d_2$, for the field operations and derivatives; $M(g)$ for (ii), differential because separability in characteristic zero turns $m(g) = 0$ into $g' = -m^\partial(g)/m_Y(g) \in M(g)$ without raising the transcendence degree; and $K\langle f\rangle(g)$, of degree $\le d_1 + 1$, for (iii), since $g' = f'g$ puts every $g^{(k)}$ in it and $g' = f'/f$ puts g' in $K\langle f\rangle$.

[2] The analogous statement is *false* for D-finite (holonomic) functions, which are *not* closed under composition — $\sec = (1/x) \circ \cos$ is not holonomic — a common source of confusion.

(iv). Let $d = \operatorname{trdeg}_{\mathbb{C}(w)} \mathbb{C}(w)\langle g\rangle$, derivatives in w. The key point is that substitution of h is a *field embedding*, not a specialization: h nonconstant takes infinitely many values, so for $r \in \mathbb{C}(w)$, $r \neq 0$, the composite $r \circ h$ cannot vanish identically (that would force h into the finite zero set of r); hence $r \mapsto r \circ h$ is an embedding $\mathbb{C}(w) \hookrightarrow \mathbb{C}(\tau)(h)$ and *every* algebraic relation over $\mathbb{C}(w)$ among $g, \ldots, g^{(m)}$ transfers verbatim to the same relation over $\mathbb{C}(h)$ among $g \circ h, \ldots, g^{(m)} \circ h$. Therefore, for *every* m, $\operatorname{trdeg}_{\mathbb{C}(h)} \mathbb{C}(h)(g \circ h, \ldots, g^{(m)} \circ h) \leq \operatorname{trdeg}_{\mathbb{C}(w)} \mathbb{C}(w)(g, \ldots, g^{(m)}) \leq d$, so the whole family $\{g^{(j)} \circ h\}_{j \geq 0}$ has transcendence degree $\leq d$ over $\mathbb{C}(h)$, hence over the larger base $\mathbb{C}(\tau)\langle h\rangle$, and $N = \mathbb{C}(\tau)\langle h\rangle(g^{(j)} \circ h : j \geq 0)$ has finite transcendence degree over $\mathbb{C}(\tau)$. By the chain rule N is a differential field, and by Faà di Bruno every derivative of $g \circ h$ lies in N; Lemma 8.2 concludes. □

Theorem 8.4 (meta-theorem: towers are D-algebraic; classical)**.** *Every $f \in \mathcal{T}$ is differentially algebraic over $\mathbb{C}(\tau)$. Moreover $\mathcal{T}$ is closed under composition (when defined), so the same holds for arbitrary finite superpositions of tower functions. In particular Λ, every finite iterated-log tower $\Lambda,\ \ln\Lambda,\ \ln\ln\Lambda,\ \ldots$, and every algebraic function of finitely many of these — including the uniform law* (8) *— is D-algebraic over $\mathbb{C}(\tau)$.*

Proof. Induct along Definition 8.1; $F_0 = \mathbb{R}(\tau)$ is trivially D-algebraic. Given F_i D-algebraic and $F_{i+1} = F_i(g)$: in case (a) g is algebraic over the differential field generated by the finitely many elements of F_i in its minimal polynomial, of finite transcendence degree by Lemmas 8.2 and 8.3(i), so 8.3(ii) applies; in cases (b), (c), 8.3(iii) applies; and every element of $F_i(g)$ is then a rational combination of D-algebraic germs. Composition-closure is Lemma 8.3(iv), and $\mathcal{T}$ itself is closed under composition because $(P/Q) \circ h \in \mathbb{C}(h)$, $e^F \circ h = e^{F \circ h}$, $(\log F) \circ h = \log(F \circ h)$, and algebraic dependence composes. For the examples, $\Lambda = \ln(\sigma^2/8\pi r^2\tau) \in \mathcal{T}$ by one application of (c), with the explicit order-one relation (39); iterated logarithms follow by repeating (c), algebraic functions of these by (a). □

That log-exp towers over $\mathbb{C}(\tau)$ are D-algebraic is classical, as is the one nontrivial ingredient, composition-closure; what is new is the transport to the free boundary, via the dictionary (10) of §4.2. In the Chen–Chadam variables $\alpha = \frac{1}{2}\xi^2$, $\xi_{\mathrm{CC}} = -\frac{1}{2}\Lambda$, Chen and Chadam [12, Theorem 6.5] prove the full series (9): there are constants $c_1 = -\frac{1}{2}, c_2, \ldots$ with $\alpha \sim -\xi_{\mathrm{CC}} + \sum_{i \geq 1} c_i \xi_{\mathrm{CC}}^{-i}$ as $\tau \to 0^+$, every truncation carrying a rigorous remainder bound — their induction produces, for every n, constants M_{n+1}, ξ_{n+1} with $|\alpha - \alpha_{n+1}| \leq M_{n+1}|\xi_{\mathrm{CC}}|^{-n-2}$ for $\xi_{\mathrm{CC}} \leq \xi_{n+1}$.

Corollary 8.5 (corner rigidity: no finite-order witness)**.** *For $N \geq 1$ define the truncated boundary*

$$B_N(\tau) \;=\; K \exp\Big(-\sigma\sqrt{2\tau\,\alpha_N(\tau)}\Big), \qquad \alpha_N \;=\; \frac{\Lambda}{2} + \sum_{i=1}^{N} c_i \Big(-\frac{2}{\Lambda}\Big)^i.$$

Then (i) *each B_N lies in $\mathcal{T}$ and is therefore D-algebraic over $\mathbb{C}(\tau)$; and* (ii) $\ln(K/B(\tau)) = \ln(K/B_N(\tau)) + O(\sqrt{\tau}\,\Lambda^{-N-1/2})$ *as $\tau \to 0^+$, so the corner expansion of B agrees with that of a D-algebraic function to order N, for every N. Consequently the near-expiry corner expansion of $B(\tau)$ is D-algebraic order by order, and* no finite-order corner analysis can witness hypertranscendence of $B(\tau)$*: any functional of finitely many corner coefficients, at any order, is realized identically by a D-algebraic function. A proof or disproof of Conjecture 8.11 must therefore use strictly more than any finite jet of the corner expansion.*

Proof. (i) $\Lambda = \ln(\sigma^2/8\pi r^2\tau)$ is operation (c) applied to an eventually positive element of $\mathbb{R}(\tau)$; $\alpha_N \in \mathbb{R}(\tau)(\Lambda)$; since $\alpha_N \sim \Lambda/2 \to +\infty$ the germ $2\tau\alpha_N$ is eventually positive and its square root is operation (a); the outer exponential is (b); multiplication by K stays in the field. Apply Theorem 8.4. (ii) By Chen and Chadam [12, Theorem 6.5] and (10), $\alpha - \alpha_N = O(\Lambda^{-N-1})$, so

$$\ln(K/B) - \ln(K/B_N) \;=\; \sigma\sqrt{2\tau}\,\big(\sqrt{\alpha} - \sqrt{\alpha_N}\big) \;=\; \sigma\sqrt{2\tau}\,\frac{\alpha - \alpha_N}{\sqrt{\alpha} + \sqrt{\alpha_N}} \;=\; O\big(\sqrt{\tau}\,\Lambda^{-N-3/2}\big),$$

which is the stated bound (indeed slightly better). For the consequence: a "finite-order corner analysis" is by definition a computation whose input is the N-jet of the expansion in the corner scales for some finite N; by (i)–(ii) that input is identical for B and for the D-algebraic B_N, so no such computation can output a certificate of hypertranscendence — nor, symmetrically, can agreement with finitely many coefficients of a D-algebraic candidate certify that B satisfies its ODE exactly. □

Remark 8.6 (what the corollary does *not* block)*.* Two limitations should be stated plainly. First, the quantifier order is $\forall N\,\exists B_N$: for each order there is a D-algebraic function with that jet. We do *not* prove $\exists g\,\forall N$ — a single D-algebraic function carrying the entire expansion — and nothing below should be read as asserting it. Second, and more importantly, the corollary blocks arguments whose input is a finite jet; it does *not* block arguments whose input is the *growth* of the coefficient sequence $(c_i)_{i\geq 1}$. Gevrey- and resurgence-type arguments, the standard route to hypertranscendence statements about asymptotic series, use series data rather than jet data and are untouched by Corollary 8.5; making the growth of (c_i) a live line of attack — for which the Chen–Chadam constants would have to be generated to high order — is the most concrete corner-side route left open. Finally, the corollary is close to tautological once "finite-order corner analysis" is *defined* as a computation on an N-jet; its mathematical content is item (i) together with the Chen–Chadam input.

Order-by-order D-algebraicity likewise says nothing about the sum of (9): an infinite sum of D-algebraic functions need not be D-algebraic, and conversely a hypertranscendental function can have every truncation D-algebraic — as would be the case here if Conjecture 8.11 holds. Whether the Chen–Chadam series converges appears to be open; in either case the information capable of deciding Conjecture 8.11 is not contained in any finite jet: it must come from the growth of (c_i), from the function–series gap ($O(\Lambda^{-\infty})$ data at the corner), or from away from the corner altogether (§8.3). The identifiability barrier of §4 is the numerical shadow of the same fact: the corner is a logarithmic clock (39), and logarithmic clocks are D-algebraically inert.

8.2 The discrete ladder at finite intensity

Remark 8.7 (pinned-trace feedback, and its dissolution)*.* One may instead attack the *discrete* object: the Carr-randomized ladder $b_m = b_m(q)$ of §2 and §6 [9], whose m-dependence is a difference equation, the natural habitat of the criteria of Hardouin and Singer [15] and Adamczewski et al. [1]. It does not work, for a reason worth naming. Say that a difference system $Y_{m+1} = A\,Y_m + B_m$ exhibits *pinned-trace feedback* if the inhomogeneity B_m is a function of a fixed-argument trace (a pointwise evaluation) of the solution Y_m itself, so that no difference–differential base field independent of the solution contains the forcing. The ladder's exact substrate is of this type: the moment recursion (18) is rank-one linear in J_m with constant coefficient $(1 - s^2)^{-1}$ and telescopes once the boundary sequence is given, but its forcing contains e^{sb_m} with

$$b_m \;=\; \ln\frac{1 + \eta + J_{m-1}(1)}{\eta} \tag{40}$$

by the pinning identity (Proposition 6.2), the $s = 1$ trace of the solution re-entering its own inhomogeneity. Consequently the system *as presented* is not σ-linear over any base field we can specify independently of the solution, and the hypotheses of Hardouin and Singer [15] and Adamczewski et al. [1] — a linear system $Y_{m+1} = AY_m + B$ with A, B in a ∂-difference base field — are not met. Two qualifications are essential. Non-linearity of a *presentation* is not non-linearizability of the *object*: Riccati recursions become linear after a Cole–Hopf-type substitution; we know of no substitution putting the substrate in the required form, and we do not claim that none exists. Nor is the feedback an obstruction to computation, being causal and triangular (b_m depends on J_{m-1}, not on J_m). Freezing $\{b_k\}$ does not help either: the frozen equation is rank-one with difference Galois group inside $\mathbb{G}_m$, its inhomogeneous content exhausted by whether the forcing telescopes in the base field, so every conclusion is relative to the differential structure of the field generated by $\{b_k\}$, which is precisely the unknown. Hypertranscendence of $m \mapsto b_m$ is therefore *not* established by this route; it remains conjectural. A decision would require either an extension of difference Galois theory to such skew systems — the Malgrange–Casale groupoid formalism is the natural candidate, though we are not aware of a classification theorem covering this class — or a different normal form altogether; we offer the above not as a proof of anything about b_m, but as a structural explanation consistent with the failure of every closed-form class tested in §6. Finally, the obstruction is a finite-q phenomenon and dissolves in the continuum in a way that decides nothing. As $\beta_c \to \infty$ at fixed $\hat{\mu}$ the pinned loop collapses onto the quasi-static manifold of §6.7, an elementary curve satisfying the first-order degree-five algebraic ODE (36), $4\hat{B}(\hat{B} - \hat{\mu}\hat{B}')^2 - 2\hat{\mu}(\hat{B} - \hat{\mu}\hat{B}') - \hat{B} = 0$, hence D-algebraic of order one — conditional, as there, on the kill-cancellation collapse of §6, whose final telescoping step is an explicit conjecture with 0.6–1.9% numerical support. Neither direction transfers: the finite-q hardness does not survive the limit, the feedback becoming an algebraic constraint; and the curve's tameness is a statement about the *scheme's* corner layer, not about $B(\tau)$, since the coordinate map degenerates at the physical corner (§6.7).

8.3 The surviving route: the balayage continuation

Where, then, could a decision live? In data beyond any finite corner jet. Exact, nonperturbative, global characterizations of the continuum boundary do exist — foremost the early-exercise-premium / free-boundary integral equation of Kim [21], Jacka [19] and Carr et al. [11], which is proven to determine B — but their analytic content has been developed in directions orthogonal to transcendence. One is uniquely suited to the question at hand: Hedenmalm's balayage equation [17, eq. (2.9)], an exact global identity forcing one diagonal slice of the two-variable Laplace transform of the graph of the boundary gap φ to be an *elementary* function, and whose *analytic-continuation* structure is unexplored. (He remarks, without proof, that under reasonable restrictions the equation characterizes the free boundary [17, §2.1]; we use only the direction that the boundary satisfies it, which is what he proves.) For $F = \mathcal{L}[\varphi]$ he proves holomorphy in $\operatorname{Re} p > -1$ and continuation past a square-root branch point at $p = -1$ into $\{\operatorname{Re} p > -2\}$ minus the slit $(-2, -1]$, reports that the scheme "seems to run into difficulty after a few steps", and leaves the maximal continuation open [17, §2.4, §3]. Since that continuation encodes the perpetual-end trans-series of φ — of which only the leading $k = 1$ sector is established [17, eq. (2.24)] — it is the natural beyond-all-orders object. What follows is program and conjecture; nothing beyond the cited results of Hedenmalm [17] is claimed as proved.

Conjecture 8.8 (singularity lattice with bounded log-depth)**.** *The maximal analytic continuation of $F = \mathcal{L}[\varphi]$ has singularities exactly at the negative integers $p = -k$, $k \geq 1$, the germ at*

$p = -k$ being generated by the k-fold convolution content of φ, with local type

$$(p+k)^{(3k-2)/2} \quad (k \text{ odd}), \qquad (p+k)^{(3k-2)/2}\log(p+k) \quad (k \text{ even}), \tag{41}$$

and no logarithm of order two or higher at any k. The alternation is forced rather than incidental: $(3k-2)/2$ is a half-integer for k odd and an integer for k even, and a pure integer power is analytic, so at even k the logarithm is the only available carrier of the singularity. The case $k = 1$ is proven [17], and its leading coefficient λ_0 is strictly positive — immediate from his eq. (2.25), though he does not state the inequality.

Remark 8.9 (bounded depth: the transfer argument, and a forecast withdrawn). An earlier form of this program predicted logarithmic depth growing like j at $p = -2j$, and proposed to run route (H) below through that growth. The prediction is withdrawn. Germ and trans-series are related by the Hankel integral

$$\frac{1}{2\pi i}\int_H p^s \log^j p \, e^{pt}\, dp \;=\; \Big(\frac{d}{ds}\Big)^j \frac{t^{-s-1}}{\Gamma(-s)},$$

and $1/\Gamma(-s)$ vanishes at every non-negative integer s: a *single* germ logarithm sitting on an integer power therefore transfers to a pure power t^{-s-1}, the would-be $\log t$ term being killed by $1/\Gamma(-s) = 0$, and only a second germ logarithm would produce one. Feeding this back through the convolution hierarchy — level k of the trans-series being assembled from products of lower levels — the induction closes with every level purely power-like in t and every germ of depth at most one, in the alternating pattern (41). One input of that induction is unproved and we flag it as such: the apparently deeper singularities carried by individual pieces of the convolution decomposition must cancel, as they are obliged to if $\mathcal{L}[\varphi^k]$ is to be independent of the decomposition used to compute it.

Remark 8.10 (the $k = 2$ rung, measured). The first falsifiable instance is the coefficient $c_2^\flat$ of $(p+2)^2\log(p+2)$, and the balayage identity at level two predicts it outright in terms of the level-one constant $\beta_1 = \Gamma(3/2)\lambda_0$ of §5: the germ should be $-\frac{1}{4}\beta_1^2\sqrt{p+1}\,(p+2)^2\log(p+2)$, equivalently a slit-density contribution of amplitude $c_L = \beta_1^2/4\pi = 1.2415\times10^{-3}$. A companion high-precision solution of the balayage integral equation returns $\lambda_0 = 0.140938227753$ to ten digits (three independent grids, mutually concordant) and measures $c_L = 1.2478(1)\times10^{-3}$ — a ratio 1.0051, stable across three datasets, the half-percent excess being the finite-window bias of the neighbouring $(p+2)^3$ terms. Hence $c_2^\flat \neq 0$: the depth-one logarithm at $k = 2$ is present and carries the predicted λ_0^2 coefficient, and the exact-cancellation alternative $c_2^\flat = 0$, which would have been evidence for closure, is excluded. These are numerical results of a companion program [46], quoted for the status of the lattice and not re-derived here; no claim of §§5–6 depends on them.

The dichotomy to be decided is this. Under *(H)* the lattice is infinite and F is not D-algebraic, whence — modulo the transfer lemma below — hypertranscendence propagates to φ and to $B(\tau)$; under *(C)* the continuation closes on a finite-sheeted Riemann surface, compatible with D-algebraicity of F, and the problem becomes a certificate search against the balayage identity. Conjecture 8.8 makes route (H) harder rather than easier, and that should be said plainly. Counting cannot separate the two branches: a D-algebraic function may have an arbitrarily wild singular set (j and the Jacobi theta functions satisfy algebraic ODEs [34, 37] yet have the unit circle as a natural boundary in the nome), and while *proving* the lattice infinite would refute D-*finiteness* of F — though not D-algebraicity — verifying finitely many lattice points refutes nothing. Depth cannot separate them either, and that is what has changed: a bounded-log-depth theorem for nonlinear algebraic ODEs, which an earlier version of this program proposed to prove and then contradict, would now be useless here even if true, since the depth it would bound is bounded already by Remark 8.9. What route (H) has left is arithmetic rather than shape — the level-k germ coefficients are predicted to lie in the tower generated by λ_0, of which the $k = 2$ rung is verified (Remark 8.10) — together with the infinitude of the lattice itself. Both are open, and one gap is named as before. *(L2) A transfer step* must carry non-D-algebraicity of F back to φ — the Laplace transform preserves D-algebraicity in neither direction in general, so this step must use the balayage identity itself — and thence to B, which is routine, the two being related by tower operations (Theorem 8.4).

8.4 The open problem

Conjecture 8.11 (hypertranscendence of the free boundary)**.** *The American put boundary $\tau \mapsto B(\tau)$ satisfies no nonzero algebraic differential equation over $\mathbb{C}(\tau)$.*

The conjecture stands open, and this section sharpens its burden rather than its plausibility. By Theorem 8.4 and Corollary 8.5 the expansion of $B(\tau)$ at expiry is D-algebraic order by order, at every order of the Chen–Chadam series, so *no finite-order asymptotic-expansion argument at the corner can settle the conjecture, in either direction.* The qualifier is essential: what is proved is that each finite jet is shared with a D-algebraic function, not that one D-algebraic function carries the whole expansion, and arguments driven by the *growth* of the coefficients (c_i) — Gevrey or resurgence type — remain available at the corner (Remark 8.6). The discrete route is blocked only at the level of available technique, and dissolves in the continuum without leaving a decision behind (Remark 8.7). What survives is the branch-structure dichotomy (H)/(C) for the balayage continuation, with one anchor proven, its second rung now measured against an exact prediction (Remark 8.10), a sharpened local structure that removes logarithmic depth from the list of possible witnesses (Remark 8.9), and the transfer lemma (L2) still separating even a confirmed lattice from a theorem. If $B(\tau)$ is hypertranscendental, the witness lives in the growth of the corner coefficients, beyond all orders at expiry, or at the perpetual end of the time axis.

9 Reproducibility, audit trail, and erratum ledger

Every load-bearing claim was classified when first asserted as (i) an exact identity verified in machine arithmetic, (ii) a measurement with stated error bars and an independent re-implementation, (iii) a gated hypothesis, asserted only as a falsification target and killed on failure, or (iv) an assumption carried under an explicit [UNPROVEN] flag. Failures are reported alongside successes: the ledger below is itself a methods result, a small but fully documented sample of how retraction risk distributes across those classes.

9.1 The verifier stack

All continuum boundary values are certified against two independently implemented solvers for $B(\tau)$ at $\sigma = 0.2$, $r = 0.05$, $K = 1$ ($\tilde{r} = 1.25$), no dividends.

Method A (primary): Kim integral equation. From the smooth-pasting (delta-matching) form $B\,\Phi(d_+(B/K,\tau)) = rK\,I_2[B](\tau)$ [21]: Chebyshev–Lobatto collocation in $z = \sqrt{\tau}$, boundary carried as $H(z) = \ln^2(K/B)$, quadrature substitution $u = (\sqrt{\tau}\sin\theta)^2$ (which removes both endpoint square-root singularities), and a global Newton iteration (Powell hybrid [39]) on the log-residual (`kim_solver.py`). The result $B(1) = 0.80875056967553 \pm 5 \times 10^{-14}$ (self-convergence) is monotone in collocation order over $M = 12$–96 with $|B_{M=96} - B_{M=64}| = 3.7 \times 10^{-14}$, and quadrature-independent to all fourteen digits across $N_q = 96$–256; it confirms the campaign fingerprint $B(1) = 0.8087505696 \pm 3 \times 10^{-10}$ from a fresh code lineage at full quoted precision. Two naive approaches fail, and are documented as such: the obvious fixed-point iteration on the boundary equation is *expansive* at small τ, each sweep amplifying the boundary error by $\sim 1/\Phi(d_+)$, which diverges as the boundary deepens; and per-node scalar root-finding on the log-residual admits no sign bracket, because the residual's leading terms cancel as the boundary becomes deep, leaving it non-monotone in the unknown. The global Newton solve on the fully coupled collocated system is a necessity, not a convenience.

Method B (independent cross-check): CN–LCP. A Crank–Nicolson scheme with Rannacher startup [40] in $x = \ln(K/S)$, a mirrored Brennan–Schwartz projection [7] per step (exact for single-interval contact on an M-matrix), boundary extraction by the $\sqrt{w}$ window fit below, $dt = dx$, Richardson extrapolation over dx (`pde_check.py`). At the price level the two methods agree to 1.3×10^{-7}: $P_A(K, 1) = 0.0609037059$ (Kim-implied) versus 0.060903576 (PDE-Richardson), with a clean $O(dx^2)$ error sequence ($-7.2 \times 10^{-6} \to -3.0 \times 10^{-7}$ over $dx = 4 \times 10^{-3}$ to 5×10^{-4}). At the boundary level agreement is 4–6×10^{-6}, limited by the PDE scheme's own contact-staircase noise floor (§9.5) rather than by any systematic offset; the Kim side is 10^{-13}-converged. Separately, `certify.py` supplies method-independent moment-identity checks (contact curvature $w_{xx} = 2rK/\sigma^2$, a boundary-velocity identity, and the Kim residual as a functional certificate on any candidate boundary), applicable to any pricer's output.

Two-lineage cross-validation of the lead results. The error law of §5 is cross-validated by two independently written LCP solvers — the single-sweep Brennan–Schwartz ladder and a Howard policy-iteration ladder (`crosscheck.py`, `policy_ladder.py`) — which agree on $n \cdot D$ to between 9×10^{-9} and 6×10^{-6} on quantities of size ≈ 0.03 and ≈ 0.007; at stage one the two lineages agree with each other, and with Carr's one-stage closed form (Proposition 5.1), to ten digits, and the Brennan–Schwartz solution satisfies complementarity to 2.4×10^{-16}. The exact identities of §6 (exact pairing, reverse-Bessel form, closed-form A_m, width identity, generating function) were reproduced to 88–99 digits by an independent party working only from the statements as printed here, by two internal routes; singulant, deck offsets, and decade ratios were reproduced digit-for-digit. One caveat is recorded rather than hidden: artifacts from an earlier phase of this project (three legacy solver generations, a ball-arithmetic audit, a claimed cross-family agreement of 7×10^{-12}) could not be re-executed and are excluded. Every numerical statement here is made against the present stack; the reference value itself survives, with improved precision.

Provenance of the ladder archive. The rung archive underlying §6 was produced by the single-sweep Brennan–Schwartz lineage, and is carried under a [RECONSTRUCTION-ONLY] flag: it reconstructs an earlier generation of the campaign whose solver could not be re-executed, so bit-identity with that generation is unproven and remains so. What can be established is the property that actually carries the section's claims — that the archived numbers are correct solutions of the stated stage problem — and it now is. Forty archived rungs spanning every decade of the archive ($q = 32$ to 10^6) and the full stage range ($m = 1$ to 450) were re-solved by the Howard policy-iteration lineage (`policy_ladder.py`), blind: the solver's own heuristic sizes its grid, and no archived value is supplied to it. At a matched spacing $dx = 10^{-2}$ the worst relative difference over the forty is 6.4×10^{-4}, attained at the shallowest rung ($q = 96$, $m = 1$, where $b \approx 3.4$ is smallest and the extraction window is coarsest relative to it); through the deck interior the agreement is 10^{-5} to 10^{-6}, and the worst absolute difference anywhere is 2.2×10^{-3} in corner units. The residual is discretization, not disagreement: refining dx from 2×10^{-2} to 2.5×10^{-3} moves the policy-iteration value monotonically *toward* the archived one — at $(q, m) = (10^4, 50)$ the gap runs $+7.4 \times 10^{-4}$, $+2.6 \times 10^{-4}$, $+1.1 \times 10^{-4}$, $+2.2 \times 10^{-5}$, and at $(10^6, 112)$ it runs $+9.6 \times 10^{-4}$ to $+2.5 \times 10^{-4}$. The archive therefore sits where an independently written complementarity solver converges, across eight decades of intensity. This does not recover the lost generation and we do not claim it does; it removes the question that mattered, which was whether the reconstruction is *right*. Script: `prov_check.py`, `prov_refine.py`.

9.2 The erratum ledger: retraction pattern by verification class

The released archive includes a compiled erratum ledger (`ERRATUM_LEDGER.md`) over the six audited campaign reports (Gate 0, Gate $0'$, derivation, kill-cancellation, closed-form, continuum translation), logging every retraction, correction, and supersession event, each cited to the report and section that made the correction, with a classification of the original claim: *was the specific proposition, as stated, backed by a direct numerical or exact-arithmetic check when first asserted?* Nine events are logged, the ninth (L9) added while this paper was being prepared. Eight concern asserted claims; L8 is an infrastructure near-miss caught before any claim was asserted on the corrupted numbers (§9.5). Among the eight, the number retracted or corrected is **1** for claims numerically verified as stated at assertion and **7** for claims asserted from inference, hypothesis, uncontrolled reading, unverified identification, or stale bookkeeping. Over the six audited reports alone, as the ledger was first compiled, the tally was **0** versus **7**; the single verified-class retraction is L9, and the fourth qualification below is given over to it. The seven: L1, an uncontrolled level-set reading of under-resolved data; L2, a definitional artifact never measured under a stated criterion; L3 and L4, unmeasured scaling hypotheses (L4 being the θ-deck ansatz itself, killed at its gate before any derivation was built on it); L5, a mis-targeted gate criterion inferred from a diagnostic; L6, an unverified algebraic identification bridging two verified facts; L7, a stale convention carry-over.

The denominator on the survivor side: approximately fifteen load-bearing claims that *were* numerically verified at assertion were re-tested by later reports — among them the projection identity (2×10^{-6}), the pinning identity and corner-reduction anchors, the moment recursion ($\leq 5 \times 10^{-4}$), the exact pairing, the closed-form base cases b_1, b_2 (10^{-12}), the measured increment law (RMS 0.023 across two independent implementations, later closed pointwise to ± 0.01), the width identity re-verified by contour quadrature, closed-form A_m (10^{-12}–10^{-28}), the order-one $G(\hat{\mu})$ collapse (deck flatness ± 0.6–1.9% over three decades of q), and the continuum dictionary and refined-law convergence. Zero retractions occurred in this class across the six audited reports. Four qualifications accompany the count, the last of them the single verified-class retraction and the most instructive of the nine events.

1. *L7 is a boundary case.* The number ($+0.14$, a residual peak at deck entry) was a genuine, verified computation — under the superseded $\frac{1}{2}\ln(m\beta_c)$ convention. What was retracted

is the claim as recorded: the number attached to the final $\frac{1}{2}\ln(2\pi m)$ convention, together with a physical "unresolved sub-structure" interpretation; recomputed under its stated convention the value is $+0.02$. Read the other way — as the retraction of a verified number — the tally becomes 1 versus 6. We count it as unverified-as-stated because the assertion was never checked under the convention it was stated in, and we disclose the judgment call here.

2. *L6's components were verified; the equivalence was not.* The closed forms for b_1, b_2 (10^{-12}) and the pinning identity (2×10^{-6}) were individually verified and never retracted. What failed was the unverified identification equating the $m = 1$ corner pinning with the $m = 2$ closure identity: compared directly, the two predictions for b_2 differed at the few-percent level, and the exact pairing decided between them to 8×10^{-4}.

3. *Selection caveat.* The ledger is compiled from the same reports that made the corrections; a claim wrongly asserted and never re-tested would be invisible to it. The mitigation is structural rather than statistical: under the campaign protocol every load-bearing claim carries either a falsification test or an [UNPROVEN] flag, and the Gate-0 verdicts were independently re-derived by a fresh-code implementation.

4. *L9, the one verified claim that fell, and the most instructive event here.* The kill-cancellation report established the remainder $R\beta_c^2$ to be decade-stable to $\pm 3\%$, with the $1/\beta_c$ coefficient 0 ± 0.03, and the ledger as first compiled listed that measurement among the survivors. It is superseded, and the supersession surfaced only when the numbers were re-derived for this paper. Re-measured on the exact reconstructed ladder over five decades rather than on the LCP archive, $R\beta_c^2$ drifts monotonically by 20–25% and a free fit returns a $1/\beta_c$ coefficient of $+0.13$ to $+0.18$, so the decay exponent is $p \approx 1.7$–1.8 and the β_c^{-2} normalization is a hypothesis rather than a measurement (Remark 6.12). The mechanism deserves the attention: the LCP archive carries an $O(q^{-1/2})$ corner-reduction error that decays *faster* than any β_c^{-p}, so it flattens an apparent scaling law instead of perturbing it. A measurement can be sound, reproducible and correctly executed, and still be reporting a property of its own substrate. This is the event behind the **1** versus **7** above, and it is logged as L9 in the released ledger. It has since been confirmed a third time, on a ladder generated entirely from the closed forms of §6 with no archive and no PDE solve: $p = 1.72$–1.80 across five decades of q, flat, with no drift toward 2 (Remark 6.13). It is also the sharpest vindication of the protocol's premise available here: the event was found by re-deriving a published number against a second substrate, which is the only check that could have found it. What survives untouched is the order-one collapse, which the exact ladder confirms.

We claim no statistical significance from nine events. What the ledger does support is a statement sharper than its tally, and we think it is the more useful of the two.

Verification class at assertion predicted retraction for eight of the nine. The ninth, L9, was verified as stated and fell anyway — and it fell for a reason its own check was structurally incapable of detecting. The archive's $O(q^{-1/2})$ reduction error decays *faster* than any β_c^{-p}, so it flattens the scaling law it was being used to measure instead of perturbing it; no amount of care taken on that substrate could have revealed the drift, because the substrate supplied the flatness. The same shape recurs in the deep-ladder anomaly of §9.5. There a global rescale of the solution was read as refuting a floating-point explanation, but rescaling moves the signal and the roundoff floor together, so the test returns bit-identity whether or not the mechanism is present. In both cases the check was sound, reproducible, correctly executed, and blind by construction.

The operative principle is therefore not that verified claims survive. It is that claims survive verification *against a substrate that does not share their error mode.* Every retraction in this ledger that had previously passed a check had passed one taken on the very substrate whose bias it was measuring; and every one was caught the same way, by re-deriving the same quantity along a route with different failure modes — a second solver family (§9.1), a second discretization (§9.5), or, in §6.3, closed forms carrying no discretization at all. That last route is the strongest available here precisely because it shares nothing with the others: no grid, no projection, no floating-point PDE solve. The discipline this suggests is cheap, and we recommend it in preference to the tally: record the verification class of a claim at the moment of assertion, and record alongside it what that verification was incapable of seeing.

9.3 Open-liabilities register

Table 3 lists the flags open at submission. The principal analytic liability is the kill-cancellation telescoping (the single [UNPROVEN] load-bearing statement of the increment-law derivation); the second is the ξ^{-4} completion of the c_0 continuum matching. Five falsification tests at $q = 10^7$ were pre-registered in the released reports with concrete numerical predictions and have since been run in the independent verification (Remark 6.26). Four pass. The fifth — the decade ratio ≈ 0.76, the only one of the five probing the decay *exponent* rather than the size or shape of the law — returns 0.79, and that failure is what demoted the β_c^{-2} attribution to a hypothesis in the table below. We record the falsified prediction here, and not only in the reports, because it is the one instance in this campaign of a pre-registered prediction being killed by its own gate.

9.4 Code and data manifest

The complete engine, data, and campaign reports are archived at `github.com/charlieyanhx/carr-boundary-er` Released engine: `kim_solver.py` (Method A); `pde_check.py` (Method B); `error_law.py` and `error_campaign.py` (the §5 campaign: $\tau = 1$ to 2^{-18}, $n = 32$–4096 half-octave, dx-Richardson, window-fit extraction, image-charge-safe domains); `policy_ladder.py`, the independent Howard policy-iteration ladder, with `crosscheck.py` for the two-lineage comparison; `certify.py`, the method-independent identity harness above; and `div_check.py`, the dividend-yield CN–LCP solver and contact-relation test behind Proposition 3.3; and `prov_check.py`, `prov_refine.py`, the blind archive re-solve of §9.1. Data: `reference_boundary.json`, the reference boundary table ($M = 96$, $N_q = 256$, dyadic τ-grid); the ladder archive (`archive.csv`, 1350 rungs (q, m, b) with per-rung diagnostics); campaign outputs `error_campaign.json` and `deep_tau.json`; and the window-sensitivity systematics log. All campaign reports are released verbatim — the six audited reports of §9.2, the verifier report, and the erratum ledger itself — so every claim in this section can be checked against its source text rather than a summary.

9.5 Numerical traps documented

The $x_{\min}$ domain wall (image-charge drain). The legacy solver domain $x \in [-40, 110]$, adequate for the original archive ($m \leq 450$), is silently inadequate at large m: the absorbing wall at $x_{\min}$ drains mass like $e^{-|x_{\min}|(b+|x_{\min}|)/m}$, reaching 20% at the continuum cross-check scales $m \sim 10^3$–3×10^4 ($q = 10^7$–3×10^{13}). The scaling rule is that $|x_{\min}|(b + |x_{\min}|)$ must grow linearly with m against the target tolerance; domains inherited from smaller runs violate it with no visible symptom in the solver output. The first large-m cross-check run was corrupted exactly this way; it was caught because the verification step (comparison against the reference continuum law) ran *before* any claim was asserted, only rescaled-domain results were published, and the derived drain bound retroactively certifies the original archive ($\leq 2.5\times10^{-5}$ in the deck, $\lesssim 0.005$ in the deepest tail). This is the ledger's L8: an infrastructure erratum, not a retraction.

Deep-ladder anti-convergence at $q = 3\times10^{13}$ (found by the gate itself; resolved — a cancellation floor). Re-running the deepest table rows, the gate reproduced the archived values exactly at archived resolution ($dx = 0.01$) — and then found that refinement moves the boundary *away*: at $m = 3\times10^4$, successive dx-halvings change b by -2.19 and -8.05 corner units ($b = 1104.909,\ 1102.724,\ 1094.674$ at $dx = 0.01,\ 0.005,\ 0.0025$; growth $\times3.7$ per halving). The effect is bit-reproducible and invariant under domain enlargement and extraction-window choice. Onset is progressive: at $q = 10^7$ the driver converges cleanly (fan $\leq 4\times10^{-4}$ in b), at $q = 10^{10}$ a small non-monotone fan appears (≤ 0.005 in ε), and at 3×10^{13} the fan reaches $O(0.5)$, with successive refinements changing b by factors of about four — a divergence, not a slow convergence.

The mechanism, identified. Holding dx fixed and varying m separates a per-stage accumulation from a static discretization error, and the answer is unambiguous: over $m = 250$ to 7500 at $q = 3\times10^{13}$ the fan $b(dx) - b(dx/2)$ grows like $m^{0.98}$ and $m^{0.97}$ for the two spacing pairs, while the implied per-stage bias is constant to 8% ($8.09\times10^{-5} \to 7.47\times10^{-5}$ over a 30-fold range in m). The error is therefore one constant bias per stage, accumulated linearly, and growing like dx^{-2} — a total $\propto m\,dx^{-2}$. This *refutes* the mechanism we previously suspected, that the per-stage bias scales like the squared cells-per-stage contact motion $(\Delta b_{\text{stage}}/dx)^2$: across the same range $(\Delta b_{\text{stage}}/dx)^2$ falls by a factor 16.7 while the measured bias is flat, so their ratio drifts monotonically by a factor 15.

The true cause is a cancellation floor in the extraction, and it is quantitative. Near contact the excess $w = v - g$ vanishes quadratically, $w \simeq \frac{1}{2}\kappa\,(b-x)^2$, so the first off-contact node carries $w_1 \sim \frac{1}{2}\kappa\,dx^2$ — which shrinks like dx^2 under refinement — while w is formed by subtracting two numbers of size $|v|$, whose representable resolution is the fixed quantity $\mathrm{ulp}(|v|)$. At $q = 3\times10^{13}$, $m = 7500$: $|v| \approx 1.47\times10^{-5}$ and $\kappa \approx 1.7\times10^{-15}$, giving

dx	w_1	$\mathrm{ulp}(\lvert v\rvert)$	w_1/ulp	b
0.01	1.39×10^{-19}	1.69×10^{-21}	82.0	571.183
0.005	1.19×10^{-20}	1.69×10^{-21}	7.0	570.624
0.0025	1.69×10^{-21}	1.69×10^{-21}	1.0	568.464

At the finest spacing the extraction signal is *exactly one unit in the last place*: there is no information left in w_1 to locate the boundary with, and the solver is reading rounding error. The signal-to-floor margin degrades like dx^{-2}, which is precisely the observed growth rate, and a one-ULP perturbation of w_1 displaces the extracted boundary by $\mathrm{ulp}/\sqrt{2\kappa w_1} = 7.8\times10^{-5}$ at $dx = 0.01$, against a measured per-stage bias of 7.5×10^{-5} — the two agreeing within a few percent, and the predicted totals matching the observed fans to better than a factor of two across all three spacings. Correspondingly, a `float32` run of the identical algorithm does not degrade gracefully but fails outright, the boundary running to the domain edge.

Two corrections follow. First, we previously reported that a global rescale of the solution left the result bit-identical and read that as refuting a floating-point explanation. The inference was wrong: rescaling v rescales w and $\mathrm{ulp}(|v|)$ together, leaving their ratio — the only quantity that matters here — invariant, so that test could not have detected this mechanism whether or not it was present. Second, the anomaly is not a defect of the corner discretization at extreme intensity, as we had supposed, but a limit of double precision in an obstacle problem whose contact curvature κ is driven to 10^{-15} by the intensity. The practical rule is a margin condition, $w_1/\mathrm{ulp}(|v|) \gtrsim 10^2$, i.e. $\kappa\, dx^2 \gtrsim 10^2\, \varepsilon_{\mathrm{mach}}|v|$, which for given precision caps the usable q from *below* in dx — the reverse of the usual refinement instinct, and the reason refining looked like divergence. It also means the archived $dx = 0.01$ values, with a margin of 82, are the most trustworthy of the three, which is how they are quoted.

The natural suspect — that the single-sweep Brennan–Schwartz projection mishandles a contact set moving several cells per stage — is *refuted*. We re-ran the deepest row with a second, independently written solver replacing the single sweep by Howard policy iteration, solving each stage's linear complementarity problem to convergence (`policy_ladder.py`); it agrees with the single-sweep solver to ten digits at stage one against the closed form of Proposition 5.1, reproduces the $q \le 10^{10}$ rows exactly, and at $q = 3 \times 10^{13}$ exhibits the *same* refinement fan, slightly larger — at $m = 7500$, over the same three spacings, $570.99 \to 569.92 \to 566.62$ against $571.18 \to 570.62 \to 568.46$. The Brennan–Schwartz solution moreover satisfies complementarity to 2.4×10^{-16}: it solves the LCP it is meant to solve. The single-sweep projection is therefore exonerated: both lineages read the same corrupted extraction, which is what the paragraph above identifies. That both anti-converge, and by similar amounts, is exactly what a shared cancellation floor predicts — it is a property of the arithmetic, not of either complementarity solve. The relevant small parameter is the corner payoff curvature $\eta = \sigma/\sqrt{2q} \approx 2.6 \times 10^{-8}$, which drives the contact curvature κ that the extraction must resolve against a fixed floating-point floor. No dx level at $q = 3 \times 10^{13}$ finer than the archived one is usable in double precision, so the two deepest rows are quoted at archived resolution — the spacing with the largest margin — and the six-decade claim is stated as five certified decades plus one provenance-only decade. The diagnostic scripts are released (`anom.py`, `anom4.py`).

Contact-staircase noise and the window fit. In the CN–LCP solver the discrete contact point advances through the grid in unit jumps, so the contact node carries an $O(dx)$ sawtooth in time and any boundary fit anchored there inherits it — non-smooth in both n and dx, which destroys Richardson extrapolation. The replacement extraction: with $w = v - g$ vanishing quadratically at the free boundary by smooth pasting, $\sqrt{w}$ vanishes linearly, and the boundary is read off as the zero crossing of a quadratic fit to $\sqrt{w}$ over a window of nodes offset from the contact set (nodes adjacent to contact, which carry the staircase noise, are excluded; the window scales with dx). This yields the clean $O(dx^2)$ error sequences above, with measured window systematics $\le 1.5 \times 10^{-4}$ in $n \cdot \Delta B$ — subdominant to every effect quoted in §5. The same extraction is used for all error-law campaign runs.

References

[1] B. Adamczewski, T. Dreyfus, and C. Hardouin. Hypertranscendence and linear difference equations. *Journal of the American Mathematical Society*, 34:475–503, 2021.

[2] R. Ait El Manssour, A.-L. Sattelberger, and B. Teguia Tabuguia. D-algebraic functions. *Journal of Symbolic Computation*, 128:102377, 2025. arXiv:2301.02512; DOI 10.1016/j.jsc.2024.102377.

[3] L. Andersen, M. Lake, and D. Offengenden. High-performance American option pricing. *Journal of Computational Finance*, 20(1):39–87, 2016.

[4] F. Avram, A. E. Kyprianou, and M. R. Pistorius. Exit problems for spectrally negative Lévy processes and applications to (Canadized) Russian options. *Annals of Applied Probability*, 14(1):215–238, 2004.

[5] G. Barles, J. Burdeau, M. Romano, and N. Samsoen. Critical stock price near expiration. *Mathematical Finance*, 5(2):77–95, 1995.

[6] B. Bouchard, N. El Karoui, and N. Touzi. Maturity randomization for stochastic control problems. *Annals of Applied Probability*, 15(4):2575–2605, 2005. arXiv:math/0602462.

[7] M. J. Brennan and E. S. Schwartz. The valuation of American put options. *Journal of Finance*, 32(2):449–462, 1977.

[8] D. S. Bunch and H. Johnson. The American put option and its critical stock price. *Journal of Finance*, 55(5):2333–2356, 2000.

[9] P. Carr. Randomization and the American put. *Review of Financial Studies*, 11(3):597–626, 1998.

[10] P. Carr and D. Faguet. Fast accurate valuation of American options. Working paper, Cornell University, 1996.

[11] P. Carr, R. Jarrow, and R. Myneni. Alternative characterizations of American put options. *Mathematical Finance*, 2(2):87–106, 1992.

[12] X. Chen and J. Chadam. A mathematical analysis of the optimal exercise boundary for American put options. *SIAM Journal on Mathematical Analysis*, 38(5):1613–1641, 2007. DOI 10.1137/S0036141003437708.

[13] J. D. Evans, R. Kuske, and J. B. Keller. American options on assets with dividends near expiry. *Mathematical Finance*, 12(3):219–237, 2002.

[14] J. Goodman and D. N. Ostrov. On the early exercise boundary of the American put option. *SIAM Journal on Applied Mathematics*, 62(5):1823–1835, 2002.

[15] C. Hardouin and M. F. Singer. Differential Galois theory of linear difference equations. *Mathematische Annalen*, 342(2):333–377, 2008.

[16] G. H. Hardy. *Orders of Infinity: The 'Infinitärcalcül' of Paul du Bois-Reymond*. Number 12 in Cambridge Tracts in Mathematics. Cambridge University Press, Cambridge, 1910.

[17] Håkan Hedenmalm. On the asymptotic free boundary for the American put option problem. *Journal of Mathematical Analysis and Applications*, 314(1):345–362, 2006. arXiv:math/0412332; DOI 10.1016/j.jmaa.2005.03.082.

[18] S. Howison. Matched asymptotic expansions in financial engineering. *Journal of Engineering Mathematics*, 53(3–4):385–406, 2005.

[19] S. D. Jacka. Optimal stopping and the American put. *Mathematical Finance*, 1(2):1–14, 1991.

[20] P. Jaillet, D. Lamberton, and B. Lapeyre. Variational inequalities and the pricing of American options. *Acta Applicandae Mathematicae*, 21:263–289, 1990.

[21] I. J. Kim. The analytic valuation of American options. *Review of Financial Studies*, 3(4): 547–572, 1990.

[22] Yerkin Kitapbayev. Closed form optimal exercise boundary of the American put option. *International Journal of Theoretical and Applied Finance*, 24(1):2150004, 2021. arXiv:1912.05438; DOI 10.1142/S0219024921500047.

[23] F. Kleinert and K. van Schaik. A variation of the Canadisation algorithm for the pricing of American options driven by Lévy processes. *Stochastic Processes and their Applications*, 125(8):3234–3254, 2015. arXiv:1304.4534.

[24] E. R. Kolchin. *Differential Algebra and Algebraic Groups*. Academic Press, 1973.

[25] R. A. Kuske and J. B. Keller. Optimal exercise boundary for an American put option. *Applied Mathematical Finance*, 5(2):107–116, 1998.

[26] A. E. Kyprianou and M. R. Pistorius. Perpetual options and Canadization through fluctuation theory. *Annals of Applied Probability*, 13(3):1077–1098, 2003.

[27] D. Lamberton. Convergence of the critical price in the approximation of American options. *Mathematical Finance*, 3(2):179–190, 1993.

[28] D. Lamberton. Error estimates for the binomial approximation of American put options. *Annals of Applied Probability*, 8(1):206–233, 1998.

[29] D. Lamberton. Brownian optimal stopping and random walks. *Applied Mathematics and Optimization*, 45:283–324, 2002.

[30] Damien Lamberton and Mohammed Adam Mikou. Exercise boundary of the American put near maturity in an exponential Lévy model. *Finance and Stochastics*, 17(2):355–394, 2013. DOI 10.1007/s00780-012-0194-z.

[31] Guillaume Leduc. Convergence speed of Bermudan, randomized Bermudan, and Canadian options. *Mathematics*, 13(2):213, 2025. DOI 10.3390/math13020213.

[32] Sergei Levendorskiĭ. Convergence of price and sensitivities in Carr's randomization approximation globally and near barrier. *SIAM Journal on Financial Mathematics*, 2(1):79–111, 2011. DOI 10.1137/100788331.

[33] L. Lipshitz and L. A. Rubel. A gap theorem for power series solutions of algebraic differential equations. *American Journal of Mathematics*, 108(5):1193–1213, 1986. composition-closure appears as Remark 4.3.

[34] K. Mahler. On algebraic differential equations satisfied by automorphic functions. *Journal of the Australian Mathematical Society*, 10:445–450, 1969.

[35] H. P. McKean, Jr. Appendix: A free boundary problem for the heat equation arising from a problem in mathematical economics. *Industrial Management Review*, 6:32–39, 1965.

[36] G. H. Meyer and J. van der Hoek. The valuation of American options with the method of lines. *Advances in Futures and Options Research*, 9:265–285, 1997.

[37] Y. Ohyama. Differential relations of theta functions. *Osaka Journal of Mathematics*, 32(2): 431–450, 1995.

[38] G. Peskir. On the American option problem. *Mathematical Finance*, 15(1):169–181, 2005.

[39] M. J. D. Powell. A hybrid method for nonlinear equations. In P. Rabinowitz, editor, *Numerical Methods for Nonlinear Algebraic Equations*, pages 87–114. Gordon and Breach, London, 1970.

[40] R. Rannacher. Finite element solution of diffusion problems with irregular data. *Numerische Mathematik*, 43:309–327, 1984.

[41] J. F. Ritt. *Differential Algebra*, volume 33 of *AMS Colloquium Publications*. American Mathematical Society, 1950.

[42] M. Rosenlicht. Hardy fields. *Journal of Mathematical Analysis and Applications*, 93(2): 297–311, 1983.

[43] L. A. Rubel. A survey of transcendentally transcendental functions. *American Mathematical Monthly*, 96(9):777–788, 1989.

[44] R. Stamicar, D. Ševčovič, and J. Chadam. The early exercise boundary for the American put near expiry: numerical approximation. *Canadian Applied Mathematics Quarterly*, 7 (4):427–444, 1999.

[45] Yasumasa Sugita and Kenta Kobayashi. High-precision numerical computation for the optimal exercise boundary of American put option. *Japan Journal of Industrial and Applied Mathematics*, 43(2):Article 21, 2026. DOI 10.1007/s13160-026-00777-y.

[46] Charlie Yan. The American put boundary as a fixed point of inverse first passage: exact laws, coefficient arithmetic, and the case against a closed form. Companion paper, in preparation, 2026.

[47] Charlie Yan. The Carr-randomization ladder as a closed system: exact kernel, telescoping, and the first-kick recursion. Companion paper, in preparation, 2026.

Flag	Statement (abridged)	Numerical status
Kill-cancellation telescoping [UNPROVEN]	Exact pairing asymptotically equals $A_m e^{O(1/\beta_c^2)}$; proof route stated (telescoping onto the pinning identities, or Poisson summation over discreteness poles).	Order-one collapse established across three decades of q (deck flatness ± 0.6–1.9%); remainder $\propto \beta_c^{-p}$, $p \approx 1.7$–1.8; pure $1/\beta_c$ excluded, but free-fit coefficient $+0.13$–$+0.18$, so exponent 2 is hypothesized (Rem. 6.12).
ξ^{-4} completion of the c_0 matching [UNPROVEN]	Bulk weld $c_0 \to -\frac{1}{2}\ln 2$, i.e. scheme prefactor $-\frac{1}{2}\ln(2\pi m) + c_0 \to -\frac{1}{2}\ln(4\pi m)$, beyond the verified order.	± 0.01 over three τ-decades.
Constant in $G(\hat{\mu})$ [UNPROVEN]	$G(\hat{\mu})$ not derived; candidate origin identified.	Correction scales like β_c^{-p}, $p \approx 1.7$–1.8 at accessible depths (independent remeasurement; exponent 2 not excluded asymptotically, *not certified*).
Kill-remainder shape [UNPROVEN]	Kill remainder carries a compensating $\approx -\frac{1}{2}\delta_{\mathrm{pref}}$-like shape.	Characterized numerically only.
Corner-order truncation (benign)	$O(\nu^{-1})(1+b)$ corrections in the corner reduction controlled but not carried; the projection identity is the exact version.	Consistently observed at $\sim 1/\sqrt{q}$.
Hypertranscendence of $m \mapsto b_m$ (conjecture)	Explicitly conjectural; what is proved is inapplicability of the linear criteria and the pinned-trace feedback obstruction (§8).	Not decidable from corner asymptotics at any finite order.
Archive provenance [RECONSTRUCTION-ONLY]	Ladder archive is a reconstruction; bit-identity with the lost generation is unproven and stays so. Correctness, which is what the claims rest on, is now independently established.	40 rungs over $q = 32$–10^6, $m = 1$–450 re-solved blind by the policy-iteration lineage: worst relative difference 6.4×10^{-4}, 10^{-5}–10^{-6} through the deck, and the gap closes monotonically under refinement (§9.1). Independently, the order-one collapse table is reproduced from the closed forms with no archive and no PDE solve of any kind (Rem. 6.13); the shape of $G(\hat{\mu})$ is the residual reconstruction-sensitive object.

Table 3: Open liabilities at submission. Each entry cites, in the released reports, the section where the flag was raised and the checks that constrain it.