

One Amos bound, three consumer sites: machine-checked Bessel-ratio calculus for lattice gauge expansions

Lluis Eriksson*

July 2026

Abstract

The Amos-type upper bound on the modified-Bessel ratio, $I_{\nu+1}(x)/I_\nu(x) < x/(\nu + \frac{1}{2} + \sqrt{(\nu + \frac{1}{2})^2 + x^2})$, has a distinguished algebraic property: its right-hand side U satisfies the exact calibration identity $1/U - U = (2\nu + 1)/x$. From that identity alone — by ordered-field algebra, with no further analytic input — follow a unit-step inequality $\rho_\nu - \rho_{\nu+1} < 1/x$ for consecutive ratios, the strict increase of the log-derivative $(\log I_\nu)' = \rho_\nu + \nu/x$ across orders, and the strict monotonicity of a φ -sequence arising in a two-dimensional lattice-gauge surface expansion. We formalize this calculus in Lean 4: a single module defines the bound once (`AmosBound`) and proves the calibration engine and four consequence theorems through that one definition, together with two rational satisfiability witnesses whose Amos hypothesis holds by exact Pythagorean arithmetic; all eighteen Lean statements of the development pass the axiom oracle with exactly `[propext, Classical.choice, Quot.sound]` against a pinned Mathlib. A certified companion (256-bit interval arithmetic, self-contained series-plus-tail enclosures, committed transcript) certifies the bound provably strictly at all 1206 points of a pre-registered grid covering the arguments the applications consume. A Bessel interface completes the closure: integer-order I_n is defined by its power series in the same pinned development, with positivity, the three-term recurrence, the termwise-differentiated derivative identity $I'_n = I_{n+1} + (n/x)I_n$, and the logarithmic-derivative identity $(\log I_n)' = \rho_n + n/x$ all proved as theorems, so the consequence theorems — including the unit step read as strict log-derivative monotonicity, in `deriv` form — hold for genuine Bessel ratios with the Amos bound as the single remaining hypothesis. The scope is stated exactly: the Amos bound itself remains a classical cited theorem taken as hypothesis — this paper unifies its three previously scattered uses in our formal development into one named proposition with one oracle and one certified numerical witness, and no downstream result changes its verification class.

1 Introduction

Cluster and character expansions of two-dimensional lattice gauge theory run on ratios of modified Bessel functions: for gauge group $SU(2)$ the mean plaquette is $I_2(\beta)/I_1(\beta)$ in closed

*This manuscript and parts of the formal development were produced with AI assistance under the repository's audit protocol. Every theorem below is machine-checked in Lean 4 against a pinned Mathlib; the grid certification of Section 5 is enclosed by committed interval arithmetic; remaining numerics are labelled paper-level where they occur. Repository: <https://github.com/lluiseiriksson/THE-ERIKSSON-PROGRAMME>.

form, the string tension is $-\log(I_2/I_1)$, and monotonicity structures of expansions in the heat-kernel/von Mises family reduce to inequalities between consecutive ratios $\rho_\nu = I_{\nu+1}/I_\nu$. The workhorse inequality is the Amos-type upper bound [1, 2, 4]

$$\rho_\nu(x) < A(\nu, x) = \frac{x}{\nu + \frac{1}{2} + \sqrt{(\nu + \frac{1}{2})^2 + x^2}} \quad (x > 0, \nu \geq 0), \quad (1)$$

with A as in Definition 2.1 below.

We call a statement that assumes (1) as an explicit hypothesis a *consumer* of the bound, and a *site* a place in the programme where such a statement — or its numerical counterpart — lives. The formal developments happen in a central repository whose Lean toolchain and Mathlib commit are fixed (*pinned*); a file checked against any other toolchain state is *off-pin*. In this programme, (1) entered three sites as a *named hypothesis* — an explicit assumption of a machine-checked theorem, never silently used: (a) the surface track’s φ -monotonicity lemma, whose Lean core carried `hAmos` verbatim and on which the surface theorem’s $E' < 0$ for $0 < \beta \leq 3.5$ is conditional; (b) the Feynman–Hellmann unit-step development for $(\log I_\nu)'$ [7], whose Lean core carried `hAmos` verbatim; (c) the exactly soluble string-tension computation, whose independent numerical program (external repository `ym-lattice-numerics`) computes precisely the ratios of (1) — with no Amos check at all. Copies (a) and (b) were standalone files machine-checked with Lean 4.30.0-rc2 and a Mathlib snapshot *different* from the central repository’s pin — statements shared by textual identity, not by a common definition, with no single artifact certifying both.

This paper contributes:

1. a single Lean module `AmosClosure` in the pinned central core defining (1) once (Definition 2.1) and proving the calibration engine (Lemmas 3.1, 3.2) and all four consumer theorems (Theorems 3.3–3.6) — Lemma 3.1 generic by design, the other five through the single definition — the toolchain-fragmentation gap closed, with a clean axiom oracle on all eighteen statements including the Bessel interface and derivative layer of Section 7;
2. rational non-vacuity witnesses (Section 6) in which the Amos hypothesis holds by *exact* arithmetic (a Pythagorean calibration point makes the square root rational);
3. a certified companion (Section 5) certifying (1) provably strictly at all 1206 points of a pre-registered grid covering the domain the consumers actually consume.

Statuses are kept separate throughout: **exact** means a Lean theorem in the pinned central core with axiom oracle [`propext`, `Classical.choice`, `Quot.sound`]; **certified** means a committed interval-arithmetic transcript with ball-plus-boolean output; **paper-level** means ordinary prose, always flagged. The formal development lives in a repository whose top-level namespace (`YangMills/`) reflects its original motivation in lattice gauge theory; no reduction to any continuum or gauge-theoretic problem is claimed or used, and — stated plainly — *no downstream theorem changes its verification class by this paper*: consumers of (1) were conditional on a classical bound before and remain so; what improves is that the condition is now one audited object instead of two copies plus an unchecked numerical site.

2 Setting

Throughout, I_ν is the modified Bessel function of the first kind; for $x > 0$ and $\nu \geq 0$ it is positive, and consecutive orders satisfy the three-term recurrence $I_{\nu-1}(x) - I_{\nu+1}(x) = \frac{2\nu}{x} I_\nu(x)$

[6, 10.29.1]. Write $\rho_\nu = I_{\nu+1}/I_\nu$. These facts are classical background (paper-level here); the Lean theorems below take the recurrence and the positivity of the relevant values as explicit hypotheses and are, as formal objects, statements of ordered-field algebra plus one square root.

Definition 2.1 (the bound, defined once). *For $\nu, x, r \in \mathbb{R}$,*

$$A(\nu, x) = \frac{x}{\nu + \frac{1}{2} + \sqrt{(\nu + \frac{1}{2})^2 + x^2}} \quad (\text{Lean: } \text{amosRHS}), \quad \text{AmosBound } \nu \ x \ r : \iff r < A(\nu, x).$$

*Every consumer theorem below (Theorems 3.3–3.6, and Lemma 3.2) hypothesizes **AmosBound** — never a textual copy of the right-hand side; the calibration engine (Lemma 3.1) is deliberately generic in its first argument and does not mention the definition. The truth of **AmosBound** $\nu \ x \ \rho_\nu(x)$ for actual Bessel ratios is the classical theorem of Amos [1] (sharp forms: [2, 4]); it is not formalized here and remains the single analytic input of every consumer.*

3 Main results

All statements are in module **AmosClosure** (file **AmosClosure/Core.lean**) of the central repository, compiled against its pinned toolchain; the recorded build and oracle facts are in Section 10.

Lemma 3.1 (calibration engine; **amos_calibration**). *Let $a, b > 0$ and $0 < t < b/(a + \sqrt{a^2 + b^2})$. Then*

$$\frac{2a}{b} < \frac{1}{t} - t.$$

Lemma 3.2 (Amos smallness; **amos_small**). *If $\nu \geq 0$, $x > 0$, $r > 0$ and **AmosBound** $\nu \ x \ r$, then $r < x/(2\nu + 1)$.*

Theorem 3.3 (φ unit step; **phi_unit_step**). *Let $m \geq 1$, $b > 0$, $u > 0$ with **AmosBound** $m \ b \ u$. Then*

$$\left(\frac{1}{u^2} - u^2\right) - \frac{2}{b} \left((m-1)u + \frac{m+2}{u}\right) + \frac{4(2m+1)}{b^2} > 0.$$

Theorem 3.4 (φ -monotonicity step; **phi_step_of_recurrences**). *Let $m \geq 1$, $b > 0$, and let $R_0 > 0$, $R_1 > 0$, R_2 satisfy the two recurrence relations $1/R_0 = R_1 + 2m/b$ and $R_2 = 1/R_1 - 2(m+1)/b$ (the shapes taken by $\rho_{m-1}, \rho_m, \rho_{m+1}$ under [6, 10.29.1]) together with **AmosBound** $m \ b \ R_1$. Then*

$$\varphi_m < \varphi_{m+1}, \quad \text{where} \quad \varphi_m = \frac{(m-1)/R_0^2 + (m+1)R_1^2}{m}.$$

Theorem 3.5 (unit step; **unit_step_of_recurrence_and_amos**). *Let $x > 0$, $\nu \geq 0$, $I_0, I_1 > 0$ and I_2 satisfy $I_0 - I_2 = \frac{2(\nu+1)}{x}I_1$ and **AmosBound** $\nu \ x \ (I_1/I_0)$. Then*

$$\frac{I_1}{I_0} - \frac{I_2}{I_1} < \frac{1}{x}$$

(positivity of I_2 is not needed).

Theorem 3.6 (log-derivative unit step; **logderiv_unit_step_increase**). *Under the hypotheses of Theorem 3.5, in the closed form $(\log I_\nu)' = \rho_\nu + \nu/x$,*

$$\left(\frac{I_2}{I_1} + \frac{\nu+1}{x}\right) - \left(\frac{I_1}{I_0} + \frac{\nu}{x}\right) > 0.$$

4 Proof architecture

The formal proofs are the authority; the mechanism is short enough to give almost in full.

One identity drives everything. Let $s = \sqrt{a^2 + b^2}$ and $U = b/(a + s)$. Rationalizing, $U = (s - a)/b$ (from $(s - a)(s + a) = b^2$), and therefore

$$\frac{1}{U} - U = \frac{s + a}{b} - \frac{s - a}{b} = \frac{2a}{b}$$

exactly. Since $t \mapsto 1/t - t$ is strictly decreasing on $(0, \infty)$, any $t < U$ satisfies $1/t - t > 2a/b$: that is Lemma 3.1, and with $a = \nu + \frac{1}{2}$, $b = x$ it converts the Amos bound into the calibrated inequality $1/\rho - \rho > (2\nu + 1)/x$. The unit step (Theorem 3.5) is then one line: the recurrence gives $I_0/I_1 - I_2/I_1 = 2(\nu + 1)/x$ exactly, and subtracting the calibrated inequality leaves exactly $1/x$ of room. Theorem 3.6 is the same statement re-indexed.

The φ step. Writing $u = R_1$ for the candidate ρ_m , the difference $\varphi_{m+1} - \varphi_m$ reduces, after substituting both recurrences and clearing denominators, to the exact algebraic identity

$$(m/R_1^2 + (m + 2)R_2^2)m - ((m - 1)/R_0^2 + (m + 1)R_1^2)(m + 1) = 2m(m + 1)\delta,$$

$$\delta = \left(\frac{1}{u^2} - u^2\right) - \frac{2}{b} \left((m - 1)u + \frac{m + 2}{u}\right) + \frac{4(2m + 1)}{b^2} = \left(u + \frac{1}{u} - \frac{3}{b}\right) \left(\left(\frac{1}{u} - u\right) - \frac{2m + 1}{b}\right) + \frac{2m + 1}{b^2}.$$

The second factor of the product is positive by the calibrated inequality; the first by the smallness $u < b/(2m + 1) \leq b/3$ (Lemma 3.2, itself one comparison of denominators); the additive term is positive outright. That is Theorems 3.3 and 3.4. Both exact identities are discharged in Lean by `field_simp`; `ring`; the inequalities by `nlinarith`/`linarith` over the ordered field. No analytic fact about I_ν enters anywhere: the analytic content is exactly the hypothesis `AmosBound` plus the recurrence shapes.

5 The certified companion

Script `scripts/c2_amos_arb.py` (SHA-256

0528dadf563d9a51eb80a95e34fb813830a5968983b5820861a3a044e5ab1def

`python-flint` 0.9.0, precision 256 bits; transcript committed alongside it with a dual-hash witness, exit 0) certifies (1) for *true* integer-order Bessel ratios, pointwise and provably, on the grid pre-registered in the charter (`docs/C2-CHARTER.md`, committed before the companion script and its transcript):

$$(\nu, x) \in \{0, \dots, 20\} \times \{0.25, 0.50, \dots, 14.00\} \text{ (1176 points),}$$

$$(\nu, x) \in \{0, \dots, 5\} \times \{20, 50, 100, 200, 300\} \text{ (30 points).}$$

The bulk box covers the surface track's *argument* range ($b = 4\beta c \leq 14$ for $\beta \leq 3.5$, where $c \in (0, 1]$ is the surface track's chart parameter); in the *order* direction the grid reaches $\nu = 20$, which is a grid choice, not a consumer-derived bound — a consumer whose ladder exceeds order 20 would need the grid extended. The wide columns probe the large- x regime. Enclosures are self-contained: $I_n(x)$ is the truncated power series plus an explicit certified remainder majorant $t_K/(1 - q)$, where t_K is the first omitted term and the term ratio is provably $< \frac{1}{2}$ from the truncation index on (checked per point; the ratio decreases in the index, so one check covers

the whole remainder; all grid abscissae are dyadic, so the series starts from exact inputs), with outward rounding throughout and tri-state verdicts (PASS only when the strict inequality holds for every point of the enclosing balls; FAIL only when the reverse does; UNDECIDED otherwise).

Result (committed transcript, run 2): **all 1206 points pass; zero fail; zero undecided**. The measured slack floor — the smallest certified lower bound of $A(\nu, x) - \rho_\nu(x)$ over the grid — is

$$2.782430724 \times 10^{-6} \pm 1.91 \times 10^{-16} \quad \text{at } (\nu, x) = (0, 300),$$

consistent with the slack being eventually decreasing in x (on the grid it first rises — e.g. at $\nu = 0$: 0.112 at $x = 0.25$, 0.184 at $x = 0.75$ — and decays only in the large- x regime, which is where the floor sits). The programme’s independent numerical program for the same ratios (`exact2d.py`, external repository `ym-lattice-numeric`s, not archived here; outward-rounded decimal arithmetic) recorded — per a reconstructed figure, marked as such in the transcript — a conservative floor lower bound of 2.5×10^{-7} on $\nu \in [1, 100]$, $x \in [0.01, 300]$; being a loose lower bound it is compatible with, and weaker than, the certified floor above. Per the charter, we registered *no* prediction of the slack floor — the figure above is measured, and the comparison is diagnostic-only. Scope: the certification is pointwise on the registered grid, for integer orders; it is a numerical witness that the classical hypothesis is comfortably true where the consumers consume it, not a proof of (1).

Remark 5.1 (a measured enclosure defect, found by audit). *The first companion run used the remainder majorant $t_K \cdot q/(1 - q)$, which bounds only the terms after the first omitted one: t_K itself was in no bucket, so the I_n enclosures provably failed containment of the true values (by $\approx t_K$, at most 6.5×10^{-34} relative on the grid). The independent numeric audit caught the defect and measured its direction: at every grid point where it is visible, both enclosures undershoot with the denominator undershooting relatively more, so the computed ratio lies above the true ratio and every PASS verdict of run 1 was a true statement reached by an unsound documented method. The majorant was corrected to $t_K/(1 - q)$, the run repeated (identical printed values: the correction is below the 10^{-16} print precision), and both transcripts are committed — run 1 marked superseded, never deleted. We record the failure as evidence for the audit protocol: enclosure methods are audited for containment, not only for verdicts.*

6 Non-vacuity, machine-checked

Both consumer families are instantiated in Lean at concrete rational values (file `NonVacuity.lean` of the module) — instance terms, not existentials. The calibration point is chosen Pythagorean: at $a = \nu + \frac{1}{2} = \frac{3}{2}$ (order $\nu = 1$) and argument $x = 2$, $\sqrt{(\frac{3}{2})^2 + 2^2} = \sqrt{\frac{25}{4}} = \frac{5}{2}$ *exactly*, so $A = \frac{1}{2}$ exactly and the Amos hypothesis for the witness ratio $\frac{2}{5}$ is the rational comparison $\frac{2}{5} < \frac{1}{2}$:

- φ family (`nonvacuous_phi_step`): $m = 1$, $b = 2$, $(R_0, R_1, R_2) = (\frac{5}{7}, \frac{2}{5}, \frac{1}{2})$; both recurrences hold by rational arithmetic, and the conclusion $\varphi_1 = \frac{8}{25} < \frac{7}{2} = \varphi_2$ is strict.
- unit-step family (`nonvacuous_unit_step`): $\nu = 1$, $x = 2$, $(I_0, I_1, I_2) = (1, \frac{2}{5}, \frac{1}{5})$; the recurrence is $1 - \frac{1}{5} = \frac{4}{5} = \frac{2 \cdot 2}{2} \cdot \frac{2}{5}$, and the conclusion $\frac{2}{5} - \frac{1}{2} < \frac{1}{2}$ is strict.

What this proves: the hypothesis sets of all four consumer theorems are satisfiable with all-positive values and a strict Amos margin, so the theorems have non-trivial instances. What it does not prove: the witnesses are not true Bessel values (the certified companion is what ties the hypotheses to actual ratios), and satisfiability is not sharpness.

7 The Bessel interface

The preceding sections treat the recurrence shapes as hypotheses. This section removes that layer: we define integer-order modified Bessel functions by their power series *in the pinned core* and prove the analytic facts the consumers need, so that the recurrence hypotheses are discharged and **AmosBound** — now a statement about genuine Bessel ratios — remains the sole named hypothesis. (File **BesselInterface.lean** of the module.)

Definition 7.1 (**besselI**). For $n \in \mathbb{N}$ and $x \in \mathbb{R}$,

$$\mathbf{besselI} \, n \, x = \sum_{k=0}^{\infty} \frac{(x/2)^{n+2k}}{k! (n+k)!},$$

as a *tsum* over the summable term family (summability for all $x \geq 0$ by comparison with the exponential series, via $(n+k)! \geq n! k!$'s consequence $(n+k)! \geq n!$; Lean: **summable_besselTerm**).

Theorem 7.2 (positivity; **besselI_pos**). For every n and $x > 0$, $\mathbf{besselI} \, n \, x > 0$.

Theorem 7.3 (the recurrence as a theorem; **besselI_recurrence**). For every n and $x > 0$,

$$\mathbf{besselI} \, n \, x - \mathbf{besselI} \, (n+2) \, x = \frac{2(n+1)}{x} \mathbf{besselI} \, (n+1) \, x$$

— DLMF 10.29.1 at order $n+1$, proved from the series by the termwise identity $\frac{2(n+1)}{x} t_{n+1,k} = t_{n,k} - t_{n+2,k-1}$ (with $t_{n+2,-1} := 0$) and an index shift of the summed family.

Theorem 7.4 (true-Bessel consumers). Let $x > 0$ and write $\rho_n = \mathbf{besselI} \, (n+1) \, x / \mathbf{besselI} \, n \, x$. If **AmosBound** holds at (n, x, ρ_n) , then $\rho_n - \rho_{n+1} < 1/x$ and the log-derivative unit step of Theorem 3.6 holds; if it holds at $(m+1, x, \rho_{m+1})$, then $\varphi_{m+1} < \varphi_{m+2}$ in the true-ratio φ -sequence. No recurrence hypothesis appears: both recurrence shapes are instances of Theorem 7.3. (Lean names: **besselI_unit_step**, **besselI_logderiv_step**, **besselI_phi_step**.)

The derivative. Termwise differentiation of the series — dominated derivatives on the ball $|y| < x+1$, through Mathlib's series-differentiation machinery — yields the derivative identity as a theorem (file **BesselDeriv.lean**):

Theorem 7.5 (derivative identity; **besselI_hasDerivAt**). For every n and $x > 0$, the function $\mathbf{besselI} \, n$ has at x the derivative $\mathbf{besselI} \, (n+1) \, x + \frac{n}{x} \mathbf{besselI} \, n \, x$ (as a **HasDerivAt**). Consequently, by positivity,

$$\frac{d}{dx} \log I_n(x) = \rho_n(x) + \frac{n}{x}$$

holds as a formal **HasDerivAt** (**besselI_log_hasDerivAt**) — the logarithmic-derivative identity is no longer classical background. Under the Amos bound at order n , the strict inequality

$$\mathbf{deriv}(\log \circ \mathbf{besselI} \, n)(x) < \mathbf{deriv}(\log \circ \mathbf{besselI} \, (n+1))(x)$$

is a theorem (**besselI_logDeriv_lt**): the unit step is log-derivative monotonicity, formally.

What remains hypothetical after this section is exactly one statement: **AmosBound** $n \, x \, \rho_n(x)$ for the true ratios — the classical bound (1) itself. Its formalization over **besselI** is the natural next arc; everything downstream of it in this paper — including the reading of the unit step as monotonicity of the logarithmic derivative — is now theorem.

8 What the closure buys, site by site

Site	Before	After
Surface track φ -lemma (<code>surface-theorem/</code> <code>PhiMonotone.lean</code>)	standalone file, Lean 4.30.0-rc2 / Mathlib <code>cd3b69ba</code> , textual <code>hAmos</code>	statements re-proved in the central pin through <code>AmosBound</code> (Theorems 3.3, 3.4); original artifact untouched
Feynman–Hellmann core [7] (archived here under <code>papers/bessel-amos-fh/</code> ; origin: external <code>lean-transfer-matrix</code> tree)	standalone file, same off-pin toolchain, textual <code>hAmos</code> , not wired into any lakefile	statements re-proved in the central pin through <code>AmosBound</code> (Theorems 3.5, 3.6)
String-tension computa- tion (<code>exact2d.py</code> , external <code>ym-lattice-numerics</code> reposi- tory, not archived here)	certified ratio enclosures, no Amos check	the same ratios certified against (1) on the registered grid (Section 5)

The surface theorem’s $E' < 0$ for $\beta \leq 3.5$ was conditional on the classical bound (1) before this paper and remains so; likewise the Feynman–Hellmann application [7] and the closure map of [9]. The gain is architectural and auditable: one definition, one oracle, one pin, one certified numerical witness — in place of two textual copies on an off-pin toolchain and a numerical program with no certified Amos witness.

9 Related work

Bounds of the form (1) originate with Amos [1], who also gave the computational scheme for the ratios; Segura [2] derived sharp two-sided bounds and the associated Turán-type inequalities; Hornik and Grün [3] studied the Amos-type family systematically; Ruiz-Antolín and Segura [4] gave the sharp family at both ends of the parameter range; and Hornik and Grün [5] treat the generalized Amos-type family $t/(\alpha + \sqrt{\beta t^2 + \gamma^2})$, showing in particular that Amos’s *upper* bound cannot be improved within that family’s parameter freedom. Why this programme carries (1) rather than a sharper bound: the consumers need the *exact* calibration $1/U - U = (2\nu + 1)/x$ of Section 4, an identity available precisely for the Amos form; sharper bounds [2, 4] trade that algebraic exactness for tightness, and by [5] the Amos upper bound is already optimal within its own family — so for consequence-mining of this kind it is the natural choice, not a convenience. The recurrence is [6, 10.29.1]. On the programme side, the prose companions of the two Lean consumer families are the author’s preprints [7] (whose “machine-checked core” is precisely the file ported here) and [9] (the surface-expansion closure map consuming the φ -lemma); a weighted Turán-type variant via the calibrated Amos bound is [8]. The formalization builds on Mathlib [10]. We are not aware of prior *formalized* treatments of Amos-type ratio bounds or their expansion-theoretic consumers; readers aware of one are invited to point us to it.

10 Reproducibility

Toolchain: Lean `leanprover/lean4:v4.29.0-rc6`; Mathlib pinned to commit

07642720480157414db592fa85b626dafb71355b

(lakefile and manifest agree); `python-flint` 0.9.0 on CPython 3.12.6; `tectonic` 0.15.0.

Recorded build facts (transcripts committed the day they were produced, named `.txt` so that the repository’s `*.log` ignore rule cannot silently exclude them): `lake build AmosClosure` completed successfully (**8161 jobs**) at the Lean-integration commit

`b0aa4f7017e5aeb5a7e923fb13d9b0e8839632d9,`

with the port from the off-pin sources compiling with zero API drift. The committed build-summary transcript (final jobs plus success line) is

`scripts/c2_lake_build_run1_GREEN_transcript.txt.`

At tag `c2-v1.2`, the extended module including `BesselInterface.lean` and `BesselDeriv.lean` rebuilt successfully in **8163 jobs**; transcript:

`scripts/c22_lake_build_run3_GREEN_transcript.txt.`

The release tag `c2-v1.0` was placed on commit `8fffa16f` after the three pre-tag audit desks reported; the audit-expedient patch tags `c2-v1.0.1` and `c2-v1.0.2` followed (five-role audit closed, charter Amendments 1–3); the Bessel interface shipped as `c2-v1.1` and the derivative layer as `c2-v1.2` (Amendments 5–7). Tags never move. The axiom oracle (the `Oracle.lean` step of the command sequence below; log committed under `scripts/` at the same commit) reports, for all eighteen declarations listed in the theorem-to-artifact map below (logs `c2_oracle_output.txt`, `_v2.txt`, and `_v3.txt` for the successive layers), exactly

`[propext, Classical.choice, Quot.sound].`

Theorem-to-artifact map (all Lean files in module `AmosClosure`):

Result	Lean name	File	Status
Lem. 3.1	<code>amos_calibration</code>	<code>Core.lean</code>	exact
Lem. 3.2	<code>amos_small</code>	<code>Core.lean</code>	exact
Thm. 3.3	<code>phi_unit_step</code>	<code>Core.lean</code>	exact
Thm. 3.4	<code>phi_step_of_recurrences</code>	<code>Core.lean</code>	exact
Thm. 3.5	<code>unit_step_of_recurrence_and_amos</code>	<code>Core.lean</code>	exact
Thm. 3.6	<code>logderiv_unit_step_increase</code>	<code>Core.lean</code>	exact
§6 witnesses	<code>nonvacuous_*</code>	<code>NonVacuity.lean</code>	exact
Def. 7.1 summability	<code>summable_besselTerm</code>	<code>BesselInterface.lean</code>	exact
Thm. 7.2	<code>besselI_pos</code>	<code>BesselInterface.lean</code>	exact
Thm. 7.3	<code>besselI_recurrence</code>	<code>BesselInterface.lean</code>	exact
Thm. 7.3 (ratio form)	<code>besselI_ratio_recurrence</code>	<code>BesselInterface.lean</code>	exact
Thm. 7.4 (i)	<code>besselI_unit_step</code>	<code>BesselInterface.lean</code>	exact
Thm. 7.4 (ii)	<code>besselI_logderiv_step</code>	<code>BesselInterface.lean</code>	exact
Thm. 7.4 (iii)	<code>besselI_phi_step</code>	<code>BesselInterface.lean</code>	exact
Thm. 7.5	<code>besselI_hasDerivAt</code>	<code>BesselDeriv.lean</code>	exact
Thm. 7.5 (log)	<code>besselI_log_hasDerivAt</code>	<code>BesselDeriv.lean</code>	exact
Thm. 7.5 (capstone)	<code>besselI_logDeriv_lt</code>	<code>BesselDeriv.lean</code>	exact
§5 grid	—	<code>scripts/c2_amos_*</code>	certified
(1) for true I_ν	—	[1, 2, 4]	classical, cited
§2 Bessel background	—	[6]	paper-level

From a clean clone at the release revision:

```
lake exe cache get
lake build AmosClosure
lake env lean AmosClosure/Oracle.lean
python scripts/c2_amos_arb.py
tectonic -X compile papers/c2-amos-closure/c2_amos_closure.tex
```

Dual hashes (LF blob and CRLF worktree) for the companion transcript are committed alongside it as

```
scripts/c2_amos_transcript_HASHES.txt;
```

canonical artifact hashes are recorded, tag-scoped, in `RELEASE-MANIFEST.md` in this paper's directory (each manuscript revision's own hashes are stamped by the addendum of the tag that ships it). Compare against `git show <rev>:path`, never a worktree checkout.

11 Conclusion

The machine-checked development routes both Lean consumer families of the Amos bound through one audited definition in one pinned core, with rational satisfiability witnesses; a certified 1206-point numerical witness covers the numerical site on the consumed domain. Limitations, stated where they arise: the bound (1) itself is classical and not formalized — with the interface of Section 7 in place it is the *only* analytic input left, now a statement about the in-core `besselI`, and proving it there is the natural next formal target; the grid certification is pointwise, not uniform; no downstream theorem changes its verification class. Reasonable next uses: consuming `AmosBound` and the `besselI` interface directly in any future surface-track or transfer-matrix formalization, and upstreaming the modified-Bessel API to Mathlib.

References

- [1] D. E. Amos, *Computation of modified Bessel functions and their ratios*, Math. Comp. **28** (1974), 239–251.
- [2] J. Segura, *Bounds for ratios of modified Bessel functions and associated Turán-type inequalities*, J. Math. Anal. Appl. **374** (2011), 516–528. doi:10.1016/j.jmaa.2010.09.030.
- [3] K. Hornik and B. Grün, *Amos-type bounds for modified Bessel function ratios*, J. Math. Anal. Appl. **408** (2013), 91–101. doi:10.1016/j.jmaa.2013.05.070.
- [4] D. Ruiz-Antolín and J. Segura, *A new type of sharp bounds for ratios of modified Bessel functions*, J. Math. Anal. Appl. **443** (2016), 1232–1246. arXiv:1606.02008.
- [5] K. Hornik and B. Grün, *Generalized Amos-type bounds for modified Bessel function ratios*, Math. Inequal. Appl. **27** (2024), 775–787. doi:10.7153/mia-2024-27-55.
- [6] NIST Digital Library of Mathematical Functions, <https://dlmf.nist.gov/>; F. W. J. Olver et al., eds.

- [7] L. Eriksson, *Recurrence–Amos proof of the unit-step order-monotonicity of $(\log I_\nu)'$, with a Feynman–Hellmann application to two-dimensional lattice gauge theory*, preprint, [ai.viXra.org:2607.0020](https://arxiv.org/abs/2607.0020), July 2026.
- [8] L. Eriksson, *A weighted Turán-type monotonicity lemma for modified Bessel functions via the calibrated Amos bound, with an application to the ordering of a surface expansion in two-dimensional lattice gauge theory*, preprint, [ai.viXra.org:2607.0017](https://arxiv.org/abs/2607.0017), July 2026.
- [9] L. Eriksson, *Ratio monotonicity for a killed von Mises bridge: exact bridge structure, certified negative results, a single-Bessel reduction, and a two-scale closure map for a surface expansion in two-dimensional lattice gauge theory*, preprint, [ai.viXra.org:2607.0023](https://arxiv.org/abs/2607.0023), July 2026.
- [10] The mathlib Community, *The Lean mathematical library*, in: Proc. 9th ACM SIGPLAN Int. Conf. on Certified Programs and Proofs (CPP 2020), ACM, 2020, pp. 367–381. doi:10.1145/3372885.3373824.