hexagon:2611.00572v1

The Tower Side of the DIC Audit

Semistable towers, the barrier $\kappa \le \frac{1}{3}$, and the conductor–discriminant constant

DIC Audit Working Papers

Chunping Zeng

13 November 2026

Abstract

For towers of curves X_n over a number field with a fixed bad-reduction set and $g_n \to \infty$, we study the height–conductor ratio $\kappa_n := h_F(\mathrm{Jac}\, X_n)/\log \mathrm{cond}(\mathrm{Jac}\, X_n)$ and the normalised quantity $R = 12\kappa - 1 - \varepsilon_{\mathrm{ar}} + O(\hat{f}^{-1})$ of the companion preprint *The Height-Conductor Ratio on Curve Towers* (2610.00544v2). Three structural results are established. First, $\sigma_n \to 1$ holds for every semistable tower satisfying the tower axioms, by the combinatorial bound $\mathrm{rank}\, L \le 2g - 3$ together with the Saito–Liu conductor–discriminant identity. Second, the barrier $\kappa \le \frac{1}{3}$ reduces to irreducible isogeny factors; the archimedean branch is closed by the unconditional Theorem 8.7 of 2610.00544v2, and the remaining input $R_\rho \le 3$ for general GL_2 and non-GL_2 towers is shown to be *not derivable* from the currently proved inputs, the obstruction being isolated in an explicit saturating model with $R \equiv 4$, $\kappa \to \frac{5}{12}$. The construction layer is subsequently closed: every tower assembled from the cited existence constructions is bounded by the maximum seed ratio, so the surviving problem reduces to a single scalar – an upper bound for the height–conductor ratio of one absolutely simple non-GL_2 abelian variety. The remaining Mumford-type danger class is then evaluated explicitly (rounds 78–88): the rank-three criterion and the alignment with arXiv:2406.19785 record five exact structures with three retractions; equality-case constructions are rigid; the finite-layer band is continuous with $\gamma = 4 - \frac{3}{2}\mu$; the End^0 MT-group classification leaves genus 4 as the sole survivor, realised by the explicit curve $C_{7,t}$; and all 58 degree-eight CM anchor fields of arXiv:2510.00093v2 are certified to have Galois group identifier 8T13, with 8T24 occurring zero times, closing the 8T24 path ($\Lambda \equiv 0$) and the long-term search items. In the final rounds 89–100 the last-mile chain is run to its end: the genus-four survivor $C_{7,t}$ is realised with explicit MT structure, CM companion and conductor formula, and the characteristic-two analysis strips the remaining finite layers $3, 6, 12, 24$, forcing tower exhaustion; a full literature-cut proof attempt of the characteristic is recorded and fails at an explicit sign barrier, leaving a conditional theorem with three propositions (a)–(c); the sign of γ_A is settled by a three-valued judgement (1274/1274 sample points positive), with $\delta_\gamma = 48c/L = 12\delta_\kappa$; and the CM anchor double-point calibration reduces the surviving input of route (II) to the single sign $\Delta\Phi_{\mathrm{eff}} > 0$, with all 58 endpoints satisfying $d_\Phi = n^2$. Route (II) remains open, conditioned on this single input. Third, the conductor–discriminant constant is not universal: inside the CM class it forms the continuum $(q^2+2q+3)/(q^2-1)$, and the mixed-prime model $u^{p^n} + v^{q^m} = 1$ contributes the new continuum $[-\frac{2}{p-1}, -\frac{2}{q-1}]$ with $\Lambda^{(a)} = 0$ exactly. On the Fermat tower the normalisation $\hat{f}_n = (2n+1)\log p$, the exact layer remainder u_n, $\varepsilon_{\mathrm{ar}} = -2w/(2n+1) + O(p^{-n}/n)$ and the closed form for R_n are derived by finite computation with a two-channel cross-check. The upper half of Problem 10.12 is closed inside the whole Kummer family; the single-curve interior realisation (I'b) and the general bound $R \le 3$ (route (II)) remain open. Nothing in this paper proves *abc*; Theorem A of the source texts is an external input and is not independently verified here.

Contents

1 Introduction **5**

2 The tower framework and its separation from the pointwise line **7**
2.1 Curve towers: the continuous side and its separation from the pointwise side . . . 7

3 The intrinsic tower-side route and the audit of open problems **9**
3.1 Replacing the whole chain: the intrinsic tower-side route . . . 9
3.2 Audit of the open problems on the tower side . . . 12
3.3 Generalising the tower-side results to arbitrary towers . . . 15

4 The Prym lift and weight locking **18**
4.1 The Prym lift: weight locking, tame exclusion, and the exact reduction of Problems 10.8/10.12 . . . 18

5 Exact computations on the Fermat tower **21**
5.1 The weighted average on the Fermat tower: exact layer census and the layer-local $\kappa = \frac{1}{8}$ law . . . 21
5.2 The dual graph Laplacian: locating the $m_p - 1$ competition, checking b_p, and reducing Problem 10.1 . . . 24
5.3 Stable-graph bounds: unconditional $\sigma \to 1$ on the Fermat tower, and a correction to 3.11 . . . 27
5.4 The three remaining items: 3.11 reconciled, $\mu(\Phi)$ localised, and the stable model located in the literature . . . 30
5.4.1 Item (i): the coefficient 3.9150 is correct . . . 30
5.4.2 Item (ii): $\mu(\Phi)$ cannot produce the $+1$ . . . 31
5.4.3 Item (iii): the stable model exists in the literature . . . 32
5.5 The three remaining items: CM discriminant multiplicity, the archimedean exponent, and the dual graph . . . 33
5.5.1 Item (ii): the multiplicity vector of d_Φ . . . 34
5.5.2 Item (i): the archimedean exponent e . . . 35
5.5.3 Item (iii): dual graph assembly . . . 36
5.5.4 Combining the three items . . . 36
5.6 The archimedean term computed directly: $e = -w$ without an equality hypothesis 37
5.6.1 The normalisation is fixed by two anchors . . . 37
5.6.2 δ_{Fal} by orbit-layer summation . . . 38
5.6.3 Wilms inversion and the closed form for R_n . . . 39
5.6.4 Ledger update . . . 40
5.7 The layer remainder computed by a finite sum: the Kronecker asymptotic is deleted 40
5.7.1 The CM type is the set of odd Dirichlet characters . . . 41
5.7.2 Exact evaluation of $\mu(\Phi_c)$ . . . 41
5.7.3 The layer remainder becomes exact . . . 42
5.7.4 The total remainder u_n . . . 42
5.7.5 Tightened error terms and unchanged conclusions . . . 43
5.8 The conductor measured: exact tower conductor mass and the stable-model cross-check . . . 44
5.8.1 Channel A: the exact tower conductor mass . . . 44
5.8.2 Quantifying the negation of 5.27 . . . 45
5.8.3 Channel B: the stable-model cross-check . . . 46
5.8.4 Ledger update . . . 46

6 General towers and the mixed-prime model **47**
6.1 General towers: what generalises and what does not . . . 47
6.1.1 Setup . . . 47
6.1.2 $\sigma \to 1$ is universal . . . 48
6.1.3 The barrier reduces to irreducible isogeny factors . . . 49

6.1.4 The Λ constant is not universal . . . 49
6.1.5 What does not generalise . . . 50
6.1.6 Ledger update . . . 50
6.2 The next tower-side step: single-curve realisation and the mixed-prime reduction . . . 51
6.2.1 A density theorem for single curves . . . 51
6.2.2 What remains of Problem 10.12 . . . 52
6.2.3 Torelli dimension count: the shortcut is closed . . . 52
6.2.4 The mixed-prime reduction . . . 53
6.3 The mixed-prime model: decoupling wild inertia . . . 54
6.3.1 The construction . . . 54
6.3.2 Wild-inertia decoupling . . . 54
6.3.3 Density decomposition and the tower axioms . . . 55
6.3.4 Weights: a single curve is not an AV product . . . 56
6.3.5 $\sigma \to 1$ and the stable-model question . . . 56
6.4 Exact constants of the mixed-prime model . . . 57
6.4.1 The exact identity . . . 57
6.4.2 The discriminant layer vanishes . . . 58
6.4.3 Relation between the two constants . . . 58
6.4.4 Endpoint correction to 6.20 . . . 58
6.4.5 Exact form of the σ bound . . . 59
6.5 Geometric description of the mixed semistable model . . . 60
6.5.1 Fibre-product structure . . . 60
6.5.2 The cluster dendrogram . . . 60
6.5.3 Tree-likeness transfers without Liu . . . 61
6.5.4 Symmetry and exact σ . . . 61
6.5.5 Why the Fermat tower is not a tree . . . 61
6.6 Finite decomposition of T and the exact constant $c(p)$ . . . 62
6.6.1 Layer conductor sums without Colmez . . . 62
6.6.2 $T = D + \frac{1}{2}R$ is an algebraic identity . . . 62
6.6.3 Exact finite-n closed form for Λ . . . 63
6.6.4 $c(p)$ upgraded, with exact error coefficient . . . 63
6.6.5 Exact three-term decomposition . . . 64
6.6.6 Independent recovery of the 5.59 asymptotic . . . 64

7 The upper half of Problem 10.12 and mixed CM types 65
7.1 The upper half of Problem 10.12: exact CM census and the $R > 1/2$ threshold . . . 65
7.1.1 Exact mixed CM type and isogeny decomposition . . . 65
7.1.2 Conductor mass, CM discriminant, and a constant excess . . . 66
7.1.3 Exact archimedean terms and κ-neutrality . . . 66
7.1.4 The exact threshold for the upper half . . . 67
7.1.5 Under the additive Colmez law the mixed tower stays below 1/8 . . . 67
7.1.6 The remaining gap is one scalar function . . . 67
7.2 Exploring hypothesis (A): refutation and withdrawal of 7.5 . . . 68
7.2.1 The four verbatim formulas . . . 68
7.2.2 Even factors make the factorwise additive law meaningless . . . 69
7.2.3 The layer mean is one universal constant, not a sum . . . 69
7.2.4 The slope survives: the upper-half threshold is not met . . . 70
7.2.5 Calibration failure and the honest boundary . . . 70
7.3 Textual alignment: the tower-level normalisation . . . 71
7.3.1 Three verbatim normalisations . . . 71
7.3.2 Prop 2.5 verified digit by digit . . . 71
7.3.3 The calibration does not fail: $c_0 = O(n/p^n)$ . . . 72

7.3.4 $\{u < N/2\}$ is a legitimate three-point orbit . . . 72
7.3.5 Two notions of $\bar{\rho}$; $C_{\mathrm{mix}} = O(1)$ promoted . . . 73
7.3.6 The mixed model is a different curve . . . 73
7.4 The mixed CM type as a disjoint union of complete orbits . . . 74
7.5 The exact leading term of μ_{ratio} and the joint-limit qualification . . . 77

8 Collapse of the barrier branches and non-semistable fibres **81**
8.1 The CM–Parseval bound and the collapse to two open branches . . . 81
8.2 The non-semistable (L1): exact scope and the thickness criterion . . . 83

9 Isogeny factors: the two-gap verdict **85**
9.1 Non-automorphic simple isogeny factors: category correction and the two-gap verdict 85

10 The construction layer: the weight ruler and the cover **88**
10.1 The construction layer of gap (I): the weight ruler and the curve-level reduction . . 88
10.2 Doing (I′): the reindexing ruler, the product obstruction, and the two halves of the curve realisation . . . 90
10.3 Attempting the cover: three explicit constructions, multiplicity–conductor cancellation, and the Prym dilution bound . . . 93
10.4 Computing κ_P: corrected multiplicities, Prym dominance, and a conditional solution of (I′b) . . . 96

11 Routes toward the general bound $R \leq 3$ **99**
11.1 Verification of $\kappa_\chi = 1/4$ and channel additivity . . . 99
11.2 Breakthrough routes for the long-term item (II) . . . 101
11.3 Pushing P1 and P2: one-sided Yuan bound, degeneration asymptotics relocated, and a single criterion for (II) . . . 102
11.4 Pushing P3, P5 and P6: three routes converge on the constant four . . . 105
11.5 Completing the P3 alignment: the weak layer on all wild Kummer towers . . . 109
11.6 P5: the normalization criterion is settled by the source, and (II) is not a normalization issue . . . 112
11.7 General-tower upper bounds: the ladder, the optimal closure, and the rigidity gap 114
11.8 Torelli non-GL2 factors: three existence layers and a literature ledger . . . 117
11.9 Verbatim verification of the prescribed-factor preprint and its effect on route (II) . 120
11.10 Final reduction of the remaining open problem: dilution, limited decoupling, and the positive-weight gap . . . 123
11.11 The tower search is closed: seed-spectrum conservation, the (T3) constraint, and reduction to a single factor . . . 125

12 The Mumford-type danger class: explicit evaluation **128**
12.1 Five external routes from four back to three: a literature check and a verdict . . . 128
12.2 Alignment of the twisted-Fermat canonical height with the tower audit . . . 131
12.3 Equality Orbits, Adelic Metrics and a Rigidity-Gap Threshold . . . 134
12.4 The Rigidity-Gap Sample Interval via the Mu-Ratio Dictionary . . . 136
12.5 Explicit Endomorphism and Mumford–Tate Determination of the Narrow Gap . . 138
12.6 Arithmetic Invariants of the Explicit Mumford-Type Family $C_{7,t}$ . . . 140
12.7 CM Anchor Parameters and the Point-Value Obstruction for γ_A . . . 143
12.8 Recovering CM Types from the Repositories and the Colmez Disjunction . . . 146
12.9 The 8T24 Specialisation of Λ . . . 149
12.10 Closing the Text-Level Items on the 8T24 Source Predicates . . . 152
12.11 Pointwise Group Identifiers of the Fifty-Eight CM Points . . . 154
12.12: New tools and new theorems for item (II) . . . 156
12.13: the first Archimedean machine value . . . 159

12.14: completion of the Archimedean pair . . . 161
12.15: the first exact finite-place conductor exponent . . . 163
12.16: $p = 3$, a p^2-level wild place and the odd conductor spectrum . . . 165
12.17: $p = 2$ and the collapse of the two long-term directions . . . 167
12.18: an attempted proof of the full dossier cut . . . 170
12.19: the sign of γ_A is a three-value judgement . . . 173
12.20: the normalisation anchor equation . . . 175
12.21: what can be computed about $\Delta\Phi$ . . . 177
12.22: two-point calibration does not shrink the band . . . 178

1 Introduction

This paper organises the *tower side* of the DIC audit of the two companion preprints of Zeng Chunping: *The Height-Conductor Ratio on Curve Towers* (2610.00544v2, the "master" text) and *Asymptotic closed forms for the Faltings height of Fermat Jacobians and the spectrum of logarithmic derivatives of L-functions* (2610.00712v2, the Fermat-tower text). Throughout, a *tower* $(X_n)_{n\geq 1}$ is a sequence of smooth projective curves over a number field K, unramified outside a fixed finite set S of primes, with $g_n \to \infty$ and conductor mass $\hat{f}_n \to \infty$ (axioms (T1)–(T3)). For such a tower the ratio

$$\kappa_n \;=\; \frac{h_F(\mathrm{Jac}\, X_n)}{\log \mathrm{cond}(\mathrm{Jac}\, X_n)}$$

and its normalisation $R = 12\kappa - 1 - \varepsilon_{\mathrm{ar}} + O(\hat{f}_n^{-1})$ are the central invariants: by the dictionary of 2610.00544v2, upper bounds on κ are equivalent to bounds on R, and the target barrier is $\kappa \leq \frac{1}{3}$ (equivalently $R \leq 3$ at $\sigma = 1$).

Statement of results

Our first group of results is *universal*. For every semistable tower, $\sigma_n \to 1$ (Section 6, Theorem 6.1): the only geometric input is the graph-theoretic bound $\mathrm{rank}\, L(\Gamma) \leq 2g - 3$ on the dual graph of the special fibre, combined with the Saito–Liu identity $\Delta_{\mathrm{Del},p} - f_p = |V(\Gamma_p)| - 1$; the two remainder terms are of different orders and decouple. On the Fermat tower $F_{p^n}\colon x^{p^n} + y^{p^n} = 1$ everything is made exact (Section 5): the normalisation $\hat{f}_n = (2n+1)\log p$, the conductor mass $\mu(\Phi_c) = \log d_{\Phi_c} + \frac{1}{2}\log p$, the archimedean exponent $e = -w$, the layer remainder $\varepsilon_c = \log p\,[p^{c-1}(p-1)]^{-1}$, and the closed form

$$R_n = \frac{1}{2} + \frac{12C_p + 2\log(2\pi^4)}{(2n+1)\log p} + O(p^{-n}/n),$$

all certified by exact finite computation and a two-channel cross-check.

The mixed-prime model $C_{p^n,q^m}\colon u^{p^n} + v^{q^m} = 1$ is treated in Section 6: wild inertia decouples ($\hat{f}_p = 2F_p^{(a)}\log p/(a-1)$ without any stable model), the dual graphs are explicit depth-n cluster trees with $\mathrm{rank}\, L = 0$, the pure-cover constant vanishes identically ($\Lambda^{(a)} = 0$), and the shape constant has the closed form $c(p) = (p^2+2p+3)/(p^2-1)$, exact to $O(p^{-n})$ with explicit first error term. The constant Λ/g_N is thereby shown to be *non-universal*: it already forms a continuum inside the CM class, extended by the mixed model to $[-\frac{2}{p-1}, -\frac{2}{q-1}]$. In Section 7 the upper half of Problem 10.12 of 2610.00544v2 is reduced to the single scalar $C_{\mathrm{mix}} > 0$, closed unconditionally inside the whole Kummer family via the exact mixed-to-three-point orbit correspondence, and the limit $\mu_{\mathrm{ratio}} \to \frac{2}{3}$ is proved path-dependently: the Heisenberg endpoint $\kappa \to \frac{1}{12}$ requires the joint limit $n, m \to \infty$, while a one-sided tower gives $(1 + q^{-m})/12$.

The barrier $\kappa \le \frac{1}{3}$ is analysed in Sections 8–9. The archimedean branch is closed by the unconditional Theorem 8.7 of 2610.00544v2; the curvature branch is identified with the upper half of Problem 10.12 via a CM–Parseval bound; the surviving item is the factor-level bound $R_\rho \le 3$ for general GL_2 and non-GL_2 towers (route (II)). This route is proved *not derivable* from the currently proved inputs: the three thresholds $R \le 3$, $R \le 3 + 2/c$, $R \le 4$ are matched by an explicit saturating model satisfying the Wilms–Noether identity, the archimedean upper bound, arithmetic Bogomolov–Gieseker with equality, isogeny convexity and $\sigma \to 1$, yet having $R \equiv 4$ and $\kappa \to \frac{5}{12}$. The constant 4 reappears three times independently (Parshin's $3+B_n \to 4$, rank-two Bogomolov $4\hat{c}_2$, and the normalisation $C_B = 4$), and the only inputs reaching 3 are the proved modular classes. Non-GL_2 simple factors are classified into three existence layers, and the 2026 preprint of Qiu on Jacobians with prescribed factors is verified verbatim, promoting one layer to a theorem relative to that preprint.

The construction layer of the remaining gap is developed in Section 10: the weight ruler $w = (24\kappa - 3)/5$ with the convexity identity $R_{\max} \equiv 3$; the single-curve GL_2 realisation (I'a) is *closed* by the Eichler–Shimura / Atkin–Lehner newform decomposition of $J(X_0(p^{2k}))$, while (I'b) is reduced to the Prym tensor-factor height κ_P, for which a closed form $\kappa_P = (\kappa_2 + r\kappa_\chi)/(1+r)$ with $\kappa_2 = \frac{1}{3}$, $\kappa_\chi = \frac{1}{4}$ is established, conditionally on the twist-tower height anchor, giving a conditional dense realisation of $(\frac{1}{4}, \frac{1}{3})$. Section 11 records the breakthrough routes toward route (II): the semistable-degeneration and Yuan-constant routes are closed or relocated, and Problem 10.2 is closed on all wild Kummer and Belyi-type towers. Two final reductions then close the construction layer itself: the exact weight law $\kappa(J_C) = (\lambda\kappa_A + \kappa_B)/(1+\lambda)$ shows that the Qiu insertion dilutes ($\lambda \to 0$, so no saturated counterexample arises); axiom (T3) forces the seed spectrum to refresh at every level; seed-spectrum conservation bounds every tower built from the cited constructions by $\max_\rho \kappa_\rho$; the K_0 quotient of the Qiu preprint carries no arithmetic data; and the boundary special subvarieties satisfy $\kappa \le (g+5)/(8g) < \frac{1}{3}$. Route (II) itself remains the single open mathematical item, now reduced to an upper bound for $h_F/\log \mathrm{cond}$ of a single absolutely simple non-GL_2 abelian variety.

The Mumford-type danger class

The final group of results evaluates the remaining danger class explicitly (Section 12, rounds 78–88). The rank-three criterion and the alignment with arXiv:2406.19785 record five exact structures with three retractions; the twist-channel constant is deliberately not promoted to U. Equality-case constructions are rigid: $\mathrm{PU}(2,1)$, Albanese triviality, and the exact threshold $\gamma \ge \lambda$. The finite-layer band is continuous, $\gamma = 4 - \frac{3}{2}\mu$ with μ a unit criterion and band $[8/3, 3)$. The explicit End^0 MT-group classification leaves genus 4 as the sole survivor, realised by the curve $C_{7,t}$ whose discriminant and conductor structure are computed, with axioms (T2) and (T3) mutually exclusive along the family. The CM anchor is certified by the 58 algebraic specialisations of arXiv:2510.00093v2, every one of Galois group identifier 8T13; the 8T24 path closes ($\Lambda \equiv 0$, group order reversal $1344 \to 48$). In rounds 89–100 the last-mile chain is run to its end: the genus-four survivor $C_{7,t}$ is realised with explicit MT structure, CM companion and conductor formula $N \cdot \prod p^4$; the characteristic-two analysis strips the remaining finite layers and forces tower exhaustion; a full literature-cut proof attempt of the characteristic is recorded and fails at an explicit sign barrier, leaving Theorem 12.121 conditional on three propositions; the sign of γ_A is settled by a three-valued judgement with 1274/1274 sample points positive and $\delta_\gamma = 48c/L = 12\delta_\kappa$; and the CM anchor double-point calibration reduces the surviving input to the single sign $\Delta\Phi_{\mathrm{eff}} > 0$. The long-term search items are thereby closed, and route (II) remains open, conditioned on this single scalar sign.

Status conventions and scope

Statements are tagged U when derived here from finite computation or standard identities, and C when cited at text level and not independently re-proved; U is always *relative to standard*

cited input, not absolute unconditionality. Theorem A of the source texts, the averaged Colmez formula, Wilms' identity and the Obus stable model are external inputs; *no part of this paper constitutes an independent verification of Theorem A*. Nothing here proves *abc*: the fourteen No-Go reasons and the pointwise conclusions of the companion pointwise file are unchanged. References to Appendices A–AA are to that pointwise companion file, which is not reproduced here.

Organisation

The ten sections group the forty underlying audit appendices thematically; within each section the subsections retain the chronological order of the analysis. Section 2 sets up the tower framework and the disjointness theorem separating the two lines; Section 3 audits the intrinsic tower-side route and the open problems of 2610.00544v2 §10; Section 4 develops the Prym lift; Sections 5 and 6 carry out the exact Fermat-tower and mixed-model computations; Section 7 treats the upper half of Problem 10.12; Sections 8–9 collapse the barrier branches and classify the isogeny factors; Section 10 develops the construction layer; Section 11 records the routes toward $R \le 3$; Section 12 evaluates the remaining Mumford-type danger class and certifies the CM anchor, and runs the last-mile chain to its end (rounds 78–100).

---

2 The tower framework and its separation from the pointwise line

2.1 Curve towers: the continuous side and its separation from the pointwise side

We keep Appendices F–AA of the pointwise companion file. The companion note's master identity (Wilms' normalised Noether formula), verbatim-checked in Appendix R, reads

$$12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}, \qquad \kappa = \frac{h_F}{\log \mathrm{cond}}, \quad \sigma = \frac{\Delta_{\mathrm{Del}}}{\log \mathrm{cond}}, \quad R = \frac{\bar{\omega}^2}{\Delta_{\mathrm{Del}}}. \tag{2.1}$$

Theorem 2.1 (The linear dictionary)**.** *Dividing* (2.1) *by* $\log \mathrm{cond}$ *gives* $12\kappa = \sigma(1+R) + \delta_{\mathrm{Fal}}/\log \mathrm{cond}$. *In the tower regime* $\sigma \to 1$ *and* $\delta_{\mathrm{Fal}}/\log \mathrm{cond} \to 0$ *(the archimedean channel absorbed by the product formula, §4.4 of the companion note),*

$$\boxed{R = 12\kappa - 1}.$$

Consequently the attainable closure $\kappa \in [1/12, 1/3]$ *is exactly* $R \in [0,3]$*, and the arithmetic BMY inequality* $\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g)$ *is precisely the right endpoint. The four landmarks are exact:*

	Heisenberg	*Fermat*	κ_{ff}	*modular/BMY*
κ	$1/12$	$1/8$	$1/6$	$1/3$
$R = 12\kappa - 1$	0	$1/2$	1	3

Proof. Divide (2.1) by $\log \mathrm{cond}$: $12\kappa = \bar{\omega}^2/\log \mathrm{cond} + \sigma + \delta_{\mathrm{Fal}}/\log \mathrm{cond} = \sigma R + \sigma + \delta_{\mathrm{Fal}}/\log \mathrm{cond} = \sigma(1+R) + \delta_{\mathrm{Fal}}/\log \mathrm{cond}$. With $\sigma \to 1$ and the archimedean term $\to 0$ this is $12\kappa = 1 + R$. The four substitutions are $12 \cdot \frac{1}{12} - 1 = 0$, $12 \cdot \frac{1}{8} - 1 = \frac{1}{2}$, $12 \cdot \frac{1}{6} - 1 = 1$, $12 \cdot \frac{1}{3} - 1 = 3$. Since $\kappa \le 1/3$ is equivalent to $R \le 3$, the "one-third barrier" of the companion note is the same statement as the BMY bound, and the companion note's implication $\kappa \le 1/3 \Rightarrow R \le 4/\sigma - 1 \le 3$ specialises to $R \le 3$ at $\sigma = 1$. Appendix AA gives the third independent reproduction of $C_{\mathrm{BMY}} = 3 \iff E = 2 \iff C_{\mathrm{Szp}} = 12$. □

Theorem 2.2 (Separating the two "3"s)**.** *Let $\lambda := \delta(F)/\Delta_{\mathrm{Del}}$ bridge the two discriminant normalisations, so that $R = \lambda\,(\bar{\omega}^2/\delta(F))_{\mathrm{actual}}$. Here $C_{\mathrm{BMY}} = 36/N$ of Theorem 5.4 is an* upper-bound coefficient*, whereas R is the* actual *ratio; $C_{\mathrm{BMY}} > 3$ does not imply $R > 3$. Conversely, if the arithmetic BMY inequality $R \le 3$ holds and $\lambda = 1$ (tame, fixed bad set), then the Theorem-A-type upper bound is slack by at least*

$$C_{\mathrm{BMY}} - 3 \in \left\{33,\ \tfrac{15}{7}\right\} \quad \text{for } \Theta \in \{8, 10\}\ (\textit{Theorem Z.4}),$$

and compatibility requires $\lambda \le 3/C_{\mathrm{BMY}}$, i.e. $\lambda \le 1/12$ for $\Theta = 8$ and $\lambda \le 7/12$ for $\Theta = 10$.

Proof. By Theorem Z.4, $\Theta \in \{8, 10\}$ after the unconditional closure, so $N = 3\Theta - 23 \in \{1, 7\}$ and $C_{\mathrm{BMY}} = 36/N \in \{36, 36/7\}$. If $R \le 3$ and $\lambda = 1$, then $(\bar{\omega}^2/\delta(F))_{\mathrm{actual}} = R/\lambda \le 3$, while the proved upper bound has coefficient C_{BMY}; the difference is an unavoidable slack of $C_{\mathrm{BMY}} - 3$, equal to 33 and $36/7 - 3 = 15/7$. Solving $\lambda\, C_{\mathrm{BMY}} \le 3$ gives $\lambda \le 3/C_{\mathrm{BMY}} = 1/12$ and $7/12$. The rows $\Theta = 12, 14$ are excluded by Theorems Z.4 and U.3 and are therefore not listed. □

Remark 2.3 (No contradiction with Appendix F). Theorem F.1 states that $12 = \frac{1}{3} \cdot \frac{m}{\rho - mc}$ is an algebraic invariant, not an estimation remainder. Theorem 2.2 is compatible: 12 is the coefficient *of this particular proof*, whereas 3 is the conjectural geometric optimum reached by a different channel (Miyaoka's rank-2 extension, the arithmetic analogue of Ω^1_X). The factor 12 between them explains why successive-minima equality criteria cannot move 12: they act on the estimation layer, not on c_{eff}.

Theorem 2.4 (Disjointness of the two limit directions)**.** *Extrapolate the general coefficient chain of Appendix K along the rank-2 locus $b = 2g_0 - 2$, $r = 2$, $\lambda = 3$. With $W = 6r(g_0 - 1)$, $\mu = 8g_0 - 8 + r(1 - b)$, $\rho = (b - 1)(b + 3)/3$, $c_{\mathrm{GRR}} = (2b^2 + 2b - 1)/12$ and $B_0 = W(\rho - \mu c_{\mathrm{GRR}})$, positivity of t requires*

$$\Theta > \Theta^*(g_0) := -\frac{B_0}{\mu} = W\Big(c_{\mathrm{GRR}} - \frac{\rho}{\mu}\Big) = \frac{(g_0 - 1)(16g_0^3 - 40g_0^2 + 26g_0 + 3)}{2g_0 - 1} \sim 8g_0^3.$$

The tower regime requires $g_n \to \infty$ with a fixed bad set, whereas the first-order degeneration model gives $\Theta = 2\sum_{j=1}^{4} n_j \le 15$, and Theorem Z.4 gives $\Theta \le 10$. Since $\Theta^*(2) = 23/3 < 8$ but $\Theta^*(3) = 306/5 = 61.2 > 10$, the chain starts only *at $g_0 = 2$ and fails for every $g_0 \ge 3$. Hence no family satisfies both "coefficient chain starts" and "$g_n \to \infty$".*

Proof. Substituting $b = 2g_0 - 2$, $r = 2$ into the invariants of Appendix K of the pointwise companion file gives $W = 12(g_0 - 1)$, $\mu = 8g_0 - 8 + 2(3 - 2g_0) = 4g_0 - 2$, $\rho = (2g_0 - 3)(2g_0 + 1)/3$, and $c_{\mathrm{GRR}} = (8g_0^2 - 12g_0 + 3)/12$; a direct rational simplification of $12(g_0 - 1)\big(c_{\mathrm{GRR}} - \rho/\mu\big)$ yields the displayed polynomial quotient, whose leading term is $8g_0^3$. Exact values: $\Theta^*(2) = 23/3$, $\Theta^*(3) = 306/5$, $\Theta^*(4) = 210.43\ldots$, $\Theta^*(13) = 13791.84\ldots$, matching Appendix K verbatim. Theorem Y.1/S.1 give $\Theta = 2\sum n_j$ with $n_j \ge 1$; the exponent-1 optimality bound of Theorem 5.5 gives $\Theta \le 15$, improved to $\Theta \le 10$ by Theorems U.3 and Z.4. Since $\Theta^*(3) > 10$, no $g_0 \ge 3$ is admissible, and the tower limit $g_n \to \infty$ is incompatible with the chain. □

Corollary 2.5 (Thirteenth No-Go reason)**.** *The curve-tower side cannot inherit the polynomial-type bound of Theorem A by taking a limit: doing so would require replacing the whole DRR–Minkowski chain, i.e. the second or third branch of the trichotomy of Theorem N.4 (an arithmetic BMV hypothesis or abc itself). This reason is unconditional and independent of the truth of Theorem A; with Appendices N, O, P of the pointwise companion file it brings the No-Go ledger to thirteen.*

Proof. Theorem 2.4 excludes the limit; Theorem N.4 classifies all possible replacements of the external inequality; the two admissible replacements are exactly the arithmetic BMV hypothesis and *abc* itself, neither of which removes the conditionality of Theorems 7.5/9.2 (Theorem H.2). No step uses the truth of Theorem A. □

Theorem 2.6 (Continuous versus discrete)**.** *The tower-side attainable set is the* continuous *interval* $\kappa \in [1/12, 1/3] \simeq R \in [0, 3]$*, densely filled by the product convexity construction of the companion note, whereas the pointwise side is the* discrete *lattice* $\{144/N : N \equiv 1 \pmod 3\}$ *of Theorem AA.1. Continuity comes from products of towers; discreteness comes from the integrality of torsion lengths* $n_j \in \mathbb{Z}$ *(Theorem Y.1). Consequently the tower side has strictly more room to push constants down, but its output is asymptotic (*$\kappa \leq 1/3$*) rather than pointwise (*$c \leq N^{E+\varepsilon}$*).*

Proof. Theorem 2.1 identifies the tower closure with $R \in [0, 3]$; the companion note's product construction gives countable density. Theorem AA.1 gives $C_{\mathrm{Szp}} = 144/N$ with $N = 3\Theta - 23 \equiv 1 \pmod 3$ and $\Theta \in 2\mathbb{Z}$ (Theorem Y.1), a discrete set. The two outputs are of different type (Theorems H.2 and M.11): no pointwise polynomial bound follows from an asymptotic κ-bound. This reproduces, from the tower side, the meta-conclusion of Theorem P.1 that dropping the pointwise requirement C1 is the only unconditionally successful exit of this paradigm. □

Remark 2.7 (What remains open on this side)*.* Five items, all recorded in the audit ledger: (1) an upper bound for $\delta_{\mathrm{Fal}}/\log \mathrm{cond}$ — Soulé's IHÉS notes state that $\hat{c}_2(\bar{E})$ "contains an archimedean summand . . . not easy to bound from above" and that one is "unable to prove Parshin's conjecture by Miyaoka's argument"; (2) the Bergman-kernel $\to \bar{\omega}^2$ step, which §4.4 of the companion note leaves "not supplied"; (3) quantitative wild ramification, since Lemma 2.4 there gives only necessity (purely tame / semistable towers cannot enter); (4) the equality case $\sigma \to 1$, i.e. a tower-side quality plane; (5) realisation by a single curve rather than by products. The general Problem 10.1 remains open. Yuan's uniform Bogomolov result (Annals 203 (2026), Prop. 2.6) is in the *opposite* direction and must not be cited as the arithmetic BMY theorem; the proved case is the modular tower $X_0(N)$ with $\bar{\omega}^2 \sim 3g \log N$.

Remark 2.8 (Scope)*.* This appendix is an algebraic recomputation and exact-rational verification built on the verbatim-checked master identity of Appendix R. It does not independently verify the proofs of Wilms' formula, the companion note's Lemma 2.4, the modular-tower slope inequalities, Soulé's Theorems A/3, Moret-Bailly 6.4 or Robin. Here C_{BMY} is an upper-bound coefficient and R an actual ratio; $\lambda = 1$ is the tame assumption (tag C). All conclusions are complementary to, not in conflict with, the pointwise closure of Theorem Z.4. The status of *abc* is unchanged; all previous conclusions, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5 - 18c)$, are unchanged.

3 The intrinsic tower-side route and the audit of open problems

3.1 Replacing the whole chain: the intrinsic tower-side route

Theorem 2.3 shows that no family satisfies both "the DRR–Minkowski coefficient chain starts" and "$g_n \to \infty$". We therefore abandon any attempt to make the tower inherit Theorem A and restart from the tower's own invariants $(\kappa, \sigma, R, \mathrm{cond}, g_n)$ of Section 2.1.

Theorem 3.1 (Three pointwise closures become inapplicable)**.** *After the replacement, the following three obstructions of the pointwise framework do not apply:*

1. *bridge breaking (Theorem 10.3): the jump of* $\mathrm{supp}(\mathfrak{d}_F)$ *to new integers* $c^5 + 4a^5$*;*
2. *the logical impossibility of a fixed finite S (Theorem O.4, S-unit finiteness);*
3. *the mod-*3 */ torsion-length lattice* $\Theta \in 2\mathbb{Z}$*,* $N \equiv 1 \pmod 3$*.*

The single mechanism is Lemma 2.4 of the companion note: a tower with a fixed bad set grows through wild ramification exponents *of the cover, not through new rational points. S-unit finiteness controls solutions of* $a/c + b/c = 1$*, not ramification data; product convexity makes the attainable set continuous (Theorem 2.5), so no integrality lattice survives. Retired with these are the Kummer field* $\mathbb{Q}(\alpha^{1/q})$*,* $\eta(q) = (q-1)/q$*, the* M_5*-discount,* $\gamma = 2/5$*,* $t = 12/N$ *and the bridge* $\delta(F)/d \asymp \log N$*.*

Proof. The tower's input is cond, whose growth at a fixed finite bad set is carried by wild inertia indices; Theorem 10.3 concerns the support of a number-field discriminant produced from a single triple and is therefore off-type. Theorem O.4 derives its contradiction from Evertse–Schlickewei–Schmidt finiteness of the S-unit equation $a/c + b/c = 1$; a tower varies the covering, not the rational solution, hence the hypothesis of that theorem is not met. Theorem 2.5 shows the tower attainable set is the continuous interval $R \in [0, 3]$, whereas the lattice of Theorem 5.4 comes from $n_j \in \mathbb{Z}$; the latter has no counterpart. Lemma 2.4 states that purely tame / semistable towers cannot enter, so wild ramification is exactly the growth mechanism in all three cases. □

Remark 3.2 (The cost). The output changes type: the tower gives an asymptotic bound $\kappa \le 1/3$, not a pointwise $c \le A_\varepsilon N^{E+\varepsilon}$. By Theorems H.2 and 2.5 no pointwise polynomial-type weak *abc* bound follows. This reproduces Theorem P.1 from the tower side: dropping the pointwise requirement C1 is the only unconditionally successful exit.

Theorem 3.3 (The intrinsic upper-bound chain)**.** *Dividing the master identity* (2.1) *by* $\log \mathrm{cond}$ *gives* $12\kappa = \sigma(1 + R) + \varepsilon_{\mathrm{ar}}$ *with* $\varepsilon_{\mathrm{ar}} := \delta_{\mathrm{Fal}}/\log \mathrm{cond}$*. If* $R \le C_{\mathrm{tower}}$ *then*

$$\boxed{\kappa \le \frac{\sigma(1 + C_{\mathrm{tower}}) + \varepsilon_{\mathrm{ar}}}{12}}.$$

For $\sigma = 1$ *and* $\varepsilon_{\mathrm{ar}} \to 0$ *the four landmarks are reproduced exactly:*

$$C_{\mathrm{tower}} = 0,\ \tfrac{1}{2},\ 1,\ 3 \quad \longmapsto \quad \kappa = \tfrac{1}{12},\ \tfrac{1}{8},\ \tfrac{1}{6},\ \tfrac{1}{3}.$$

Moreover $\partial\kappa/\partial\sigma = (1 + C_{\mathrm{tower}})/12$*,* $\partial\kappa/\partial C_{\mathrm{tower}} = \sigma/12$ *and* $\partial\kappa/\partial\varepsilon_{\mathrm{ar}} = 1/12$*, so all three are* $O(1)$ *inputs;* $\kappa \le 1/3$ *with* $C_{\mathrm{tower}} = 3$ *forces* $4\sigma + \varepsilon_{\mathrm{ar}} \le 4$*, i.e.* $\sigma \le 1$ *and* $\varepsilon_{\mathrm{ar}} \to 0$*.*

Proof. From $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}$ divide by $\log \mathrm{cond}$ and substitute $\bar{\omega}^2/\log \mathrm{cond} = (\bar{\omega}^2/\Delta_{\mathrm{Del}})\sigma = R\sigma$: $12\kappa = \sigma R + \sigma + \varepsilon_{\mathrm{ar}} = \sigma(1 + R) + \varepsilon_{\mathrm{ar}}$. Bounding R by C_{tower} and dividing by 12 gives the displayed inequality. With $\sigma = 1, \varepsilon_{\mathrm{ar}} = 0$ the four substitutions are $(1 + 0)/12 = 1/12$, $(1 + \frac{1}{2})/12 = 3/24 = 1/8$, $(1 + 1)/12 = 1/6$, $(1 + 3)/12 = 1/3$; this is the inverse of $R = 12\kappa - 1$ of Theorem 2.1. The partial derivatives are immediate. Finally $\kappa \le 1/3$ reads $\sigma(1 + 3) + \varepsilon_{\mathrm{ar}} \le 4$, i.e. $4\sigma + \varepsilon_{\mathrm{ar}} \le 4$. □

Theorem 3.4 (The five gaps split into three inputs and two realisation problems)**.** *Of the five items of Remark 2.6, exactly three enter the upper-bound chain of Theorem 3.3 as inputs:*

input	*gap*
$\varepsilon_{\mathrm{ar}}$	*(1) archimedean bound for* δ_{Fal}
C_{tower}	*(2) Bergman kernel* $\to \bar{\omega}^2$
σ	*(4) the equality case* $\sigma \to 1$

The remaining two — (3) quantitative wild ramification (Lemma 2.4 gives only necessity) and (5) realisation by a single curve instead of by products — are attainability and realisation questions, not upper-bound inputs. Consequently the tower-side trilemma is: (i) *prove* $C_{\mathrm{tower}} \le 3$*,* $\varepsilon_{\mathrm{ar}} \to 0$*,* $\sigma \le 1$*;* (ii) *import a new external height inequality (Vojta / Faltings product theorem type);* (iii) *take the tower itself as input (circular).*

Proof. Theorem 3.3 displays κ as a function of exactly σ, C_{tower} and $\varepsilon_{\mathrm{ar}}$; no other quantity appears. Items (3) and (5) of Remark 2.6 concern whether a tower with the required properties exists and whether the closure is attained by one curve, not the value of the upper bound. The classification of replacements parallels Theorem N.4 but with different branches: (i) is Problem 10.1 of the companion note, (ii) is a new external input, (iii) is circular. Soulé's IHÉS notes record the two difficulties of (i) verbatim: "we do not know any arithmetic analog for Ω^1_X" and $\hat{c}_2(\bar{E})$ "contains an archimedean summand . . . not easy to bound from above". □

Theorem 3.5 (The archimedean share becomes measurable)**.** *Let $\lambda = \delta(F)/\Delta_{\mathrm{Del}}$ as in Theorem 2.2. The Theorem-A-type bound $\bar{\omega}^2 \le C_{\mathrm{BMY}}\delta(F)$ gives $R \le C_{\mathrm{BMY}}\lambda$. Two directions:*

1. Inheritance*: to make the pointwise upper bound stronger than BMY ($R \le 3$) one needs $C_{\mathrm{BMY}}\lambda \le 3$, i.e. $\lambda \le 3/C_{\mathrm{BMY}}$;*
2. Reverse consistency*: if the tower attains $R = 3$, then $C_{\mathrm{BMY}}\lambda \ge 3$, i.e. $\lambda \ge 3/C_{\mathrm{BMY}}$.*

Hence the unique self-consistent point is

$$\boxed{\lambda_{\mathrm{eq}} = \frac{3}{C_{\mathrm{BMY}}} = \frac{N}{12}, \qquad 1 - \lambda_{\mathrm{eq}} = \textit{archimedean share}},$$

and the identity $C_{\mathrm{BMY}}/3 = 1/\lambda_{\mathrm{eq}} = 12/N$ holds on every row. With $N = 3\Theta - 23$ and $\Theta \in \{8, 10\}$ (Theorem Z.4): $\lambda_{\mathrm{eq}} = 1/12$ (archimedean share $11/12$) for $\Theta = 8$, and $\lambda_{\mathrm{eq}} = 7/12$ (share $5/12$) for $\Theta = 10$. The rows $\Theta = 12, 14$ give $\lambda_{\mathrm{eq}} > 1$ and a negative archimedean share, consistent with their exclusion by Theorems Z.4 and U.3.

Proof. $C_{\mathrm{BMY}} = 36/N$ by Theorem AA.1, so $3/C_{\mathrm{BMY}} = 3N/36 = N/12$, and $C_{\mathrm{BMY}}/3 = 12/N = 1/\lambda_{\mathrm{eq}}$. For $\Theta = 8, 10$ we have $N = 1, 7$, giving $\lambda_{\mathrm{eq}} = 1/12, 7/12$ and shares $11/12, 5/12$; for $\Theta = 12, 14$ we have $N = 13, 19$, giving $13/12, 19/12 > 1$. The two inequalities in (1) and (2) meet only at equality, so λ_{eq} is a single point rather than an interval; this is the tower-side analogue of the three exact conditions of Theorem 12. Soulé's qualitative statement that the archimedean summand is hard to bound from above is turned into the quantitative requirement that it occupy at least $5/12$ of Δ_{Del}. □

Theorem 3.6 (Reversal of the role of successive minima)**.** *On the pointwise side the equality criterion of Ballaÿ (2026), $\zeta_{\mathrm{ess}} = \widehat{\mu}^{\mathrm{asy}}_{\max}$, acts on the ζ_i-comparison layer and leaves the five locked factors of $12 = \frac{1}{3} \cdot \frac{m}{\rho - mc}$ unchanged (Theorem F.1), so its effect on the coefficient 12 is zero. On the tower side $g_n \to \infty$, the asymptotic slope ζ_{ess} of the successive minima of $\bar{\omega}_n$ is a* candidate upper bound *for $\bar{\omega}^2/\hat{f}_n$, hence a direct input for C_{tower} in Theorem 3.3. Likewise the Hilbert–Samuel dilution of Theorem N.3, which destroys the pointwise bound because the discriminant term $\mu/n \to 0$ while the archimedean volume grows like $n/3$, is precisely the $n \to \infty$ mechanism that produces C_{tower} in a tower. The verifiable condition is $\zeta_{\mathrm{ess}}(\bar{\omega}_n) \le 3\Delta_{\mathrm{Del}}/\hat{f}_n$ along the tower.*

Proof. The pointwise statement is Theorem F.1 / F.4: 12 is $\frac{1}{3} \cdot \frac{m}{\rho - mc}$ with $\rho, m, c, \frac{1}{3}$ algebraic invariants, and Ballaÿ's criterion concerns the comparison of ζ_i's at fixed degree, so it cannot change them. In a tower the degree is not fixed: the arithmetic Hilbert–Samuel expansion $\widehat{\deg} H^0(n\bar{L}) \sim n^{d+1}\bar{L}^{d+1}/(d+1)!$ makes the leading slope a function of $n \to \infty$, which is exactly the regime of ζ_{ess}. By Theorem 3.3 the quantity bounding $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$ is C_{tower}; hence the same tool moves from "no effect on the constant" to "the main input". We do not re-prove Ballaÿ's criterion or evaluate the concave transform on this tower (tag C); we record the change of logical position. □

Corollary 3.7 (Complementary failure of the two chains)**.** *Calibrating by the companion note's Vojta limit (abc exponent 2, Szpiro constant 12) and Theorem AA.1 gives*

$$C^{\mathrm{tower}}_{\mathrm{Szp}} = 36\kappa, \qquad E = 6\kappa, \qquad C^{\mathrm{tower}}_{\mathrm{BMY}} = 9\kappa.$$

The two anchors are exact: $\kappa = 1/3 \Rightarrow C_{\mathrm{BMY}} = 3$, and $\kappa = 1/6 \Rightarrow C_{\mathrm{BMY}} = 3/2 = mc$, $E = 1$, $C_{\mathrm{Szp}} = 6$ (Theorem AA.1). Matching $C^{\mathrm{tower}}_{\mathrm{Szp}}$ to the pointwise discrete value $144/N$ gives $\kappa = 4/N$, so $\kappa \le 1/3$ requires $N \ge 12$, i.e. $\Theta \ge 35/3$ — non-integral with $N = 12 \equiv 0 \pmod 3$, a fourth independent occurrence of the mod-3 obstruction. Row by row:

Θ	$\kappa = 4/N$	$\kappa \le 1/3$	*pointwise status*
8	4	*no*	*survives Z.4*
10	4/7	*no*	*the only surviving cell of Z.4*
12	4/13	*yes*	*excluded by Z.4*
14	4/19	*yes*	*excluded by U.3*

Thus the two chains fail complementarily*: the cells that pass the tower barrier are exactly those excluded pointwise, while the unique surviving pointwise cell fails the tower barrier. This is the fourteenth No-Go reason (interface layer), independent of the truth of Theorem A.*

Proof. Theorem AA.1 gives $C_{\mathrm{Szp}} = 6E$ and $C_{\mathrm{BMY}} = \frac{3}{2}E$; the Vojta limit fixes $E = 2$, $C_{\mathrm{Szp}} = 12$ at $\kappa = 1/3$, whence $C_{\mathrm{Szp}}^{\mathrm{tower}} = 36\kappa$ and, dividing, $E = 6\kappa$, $C_{\mathrm{BMY}}^{\mathrm{tower}} = 9\kappa$. The anchor $\kappa = 1/6$ gives $E = 1$, $C_{\mathrm{BMY}} = 3/2$, $C_{\mathrm{Szp}} = 6$, matching Theorem AA.1. Equating $36\kappa = 144/N$ gives $\kappa = 4/N$; $\kappa \le 1/3$ is $N \ge 12$, and $\Theta = (N+23)/3 \ge 35/3$. The four rows are direct substitutions. No step uses the truth of Theorem A. □

Remark 3.8 (Scope and the next steps). The classification of Theorem 3.1, the chain of Theorem 3.3, the splitting of Theorem 3.4 and the λ-criterion of Theorem 3.5 are U. The master identity, Lemma 2.4, $\sigma \to 1$, the modular-tower case, Soulé's IHÉS remarks, Ballaÿ's criterion, the companion note's Vojta limit and the normalisation 36κ are text-level or C-level inputs from the archived audits; in particular the slope 36 depends on whether h_F carries a factor $[F : \mathbb{Q}]$ and whether cond includes wild exponents. Changing that normalisation changes the slope but not the direction of the landing condition $N \ge 12$, $N \equiv 0 \pmod 3$. The four executable next steps, in order of value, are: bound $\varepsilon_{\mathrm{ar}}$ using the measurable shares of Theorem 3.5; supply the Bergman-kernel step; build a tower-side quality plane for $\sigma \to 1$; and verify the normalisation dictionary verbatim. This appendix does not independently verify any cited proof, does not prove general arithmetic BMY, and does not turn the tower asymptotic into a pointwise weak *abc* bound. All previous conclusions, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5-18c)$, are unchanged; *abc* remains open.

3.2 Audit of the open problems on the tower side

We read the two companion papers verbatim: `2610.00544v2` (31 pp., 29 Oct 2026) on the height–conductor ratio κ of curve towers, and `2610.00712v2` (10 pp., 4 Nov 2026) on the closed forms for Faltings heights of Fermat Jacobians. This resolves the earlier "unverifiable self-citation" entry of Appendix R: the master identity, Lemma 2.4, Remark 4.6, §5.2, §9.3, §10.2 and the revision ledger E9–E26 of 00544 are all present. This is a *text-level* verification, not an independent verification of their proofs.

Theorem 3.9 (Normalisation dictionary; Problem 10.7 disappears in the tower world)**.** *The two anchors of 00712 are $h_F(j=0) = -1.667691$ (Faltings) and $h(j=0) = -0.7487525$ (Colmez), whose difference is exactly $-\frac{1}{2}\log 2\pi$:*

$$\tfrac{1}{2}\log 2\pi = 0.9189385332, \qquad -1.667691 - (-0.7487525) = -0.9189385,$$

with residual 3.32×10^{-8}. The Wilms/Moret-Bailly difference of Remark 2.1 of 00544 is $-4g\log 2\pi$ in $12h_F$, i.e. $\frac{1}{3}g\log 2\pi$ in h_F. All three normalisations differ by a multiple of g. Since $\kappa = h_F/\log\mathrm{cond} = (h_F/g)/\hat{f}$ and $\hat{f} = \log\mathrm{cond}\,/g \to \infty$ in the tower world, any such difference contributes $\Delta c/\hat{f} \to 0$. Consequently Problem 10.7 is vacuous for towers, and the slope $C_{\mathrm{Szp}}^{\mathrm{tower}} = 36\kappa$ of Corollary 3.6 is unaffected by the normalisation; the tag C *attached to that slope in Remark 3.7 can be removed.*

Proof. The numerical identity is direct. Remark 2.1 of 00544 records $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}$ (Wilms) against $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}} - 4g\log 2\pi$ (Moret-Bailly), so dividing by 12 the height difference is $\frac{1}{3}g\log 2\pi$; 00712 (1.3) subtracts $\frac{g}{2}\log 2\pi$ in the Colmez form. Each is $c\cdot g$ for an absolute constant c. Writing $h_F^{\mathrm{new}} = h_F^{\mathrm{old}} + cg$ and dividing by $\log\mathrm{cond} = g\hat{f}$ gives an additive κ-error $c/\hat{f}$, which tends to 0 by (T3). A shift in κ by $o(1)$ does not change the coefficient in $C_{\mathrm{Szp}}^{\mathrm{tower}} = 36\kappa$. □

Theorem 3.10 (The one-line rule is verbatim in 00544; a fifth landmark)**.** *Section 2.4 of 00544 states: "The anchors lie on the line $\kappa = (1+R)/12$ (i.e. $\sigma \to 1$, $d \to 0$): they are one rule, not seven independent data." This is precisely $R = 12\kappa - 1$ of Theorem 2.1. With the E22 correction $R(\text{Fermat}) = 1/2$ and the $X(p^n)$ row, the exact verification is*

	Heisenberg	*Fermat*	*NS/Frey*	κ_{ff}	$X(p^n)$	$X_0(N)$
κ	1/12	1/8	1/8	1/6	5/24	1/3
R	0	1/2	1/2	1	3/2	3
$(1+R)/12$	1/12	1/8	1/8	1/6	5/24	1/3

Moreover 00544 records that $R \le 3/2$ and $\sigma \le 2$ cannot hold simultaneously (they would give $\kappa = 5/12 > 1/3$); the actual limit is $\sigma \to 1$, $\kappa = 5/24$. Hence Theorem 2.1 is upgraded from an independent derivation to a verbatim statement of 00544, and its landmark table gains the $X(p^n)$ row.

Proof. Substitute each κ into $12\kappa - 1$: $12 \cdot \frac{1}{12} - 1 = 0$, $12 \cdot \frac{1}{8} - 1 = \frac{1}{2}$, $12 \cdot \frac{1}{6} - 1 = 1$, $12 \cdot \frac{5}{24} - 1 = \frac{3}{2}$, $12 \cdot \frac{1}{3} - 1 = 3$. In each case $(1+R)/12 = \kappa$. The $R \le 3/2$, $\sigma \le 2$ warning is §2.4 of 00544 verbatim; with $\sigma = 1$ the master identity gives $\kappa = (1 + 3/2)/12 = 5/24$. □

Theorem 3.11 (Exact rate of the archimedean term on the Fermat tower)**.** *On the Fermat tower, $\kappa = 1/8$, $\sigma = 1$, $R = 1/2$, so $12\kappa = \sigma(1+R) = 3/2$ and $\varepsilon_{\text{ar}} := \delta_{\text{Fal}}/\log \text{cond} \to 0$ exactly. Using Theorem 3.4 of 00712, $\hat{h}_n = \frac{\log p}{4} n + C_p + \Theta(n \log p / p^n)$ with*

$$C_p = -6 \log A - \frac{p \log p}{4(p^2-1)}, \qquad A = 1.2824271291\ldots,$$

and the conductor density $c = 2$ of Theorem 3.2 of 00544 ($\hat{f}_n = 2n \log p + O(1)$), one gets

$$\kappa_n = \frac{1}{8} + \frac{C_p}{2n \log p} + O(p^{-n}), \qquad \boxed{\varepsilon_{\text{ar}}(n) = \frac{6C_p}{n \log p} + O(p^{-n})}.$$

Numerically the coefficient is $-8.7138/n$ ($p = 3$), $-5.8767/n$ ($p = 5$), $-4.8208/n$ ($p = 7$), $-3.8721/n$ ($p = 11$). Combining with the Wilms term $2\log(2\pi^4)/(\sigma \hat{f}) = +4.7988/n$ ($p = 3$) in the non-circular bound of §9.3 of 00544 gives $R_n < \frac{1}{2} - \frac{3.9150}{n}$ ($p = 3$), i.e. R approaches 1/2 from below*, independently agreeing with the abstract of 00544 ("converges to 1/8 from below").*

Proof. $\log A = 0.248754477\ldots$, so $-6 \log A = -1.492526862$, and $C_3 = -1.595522$, $C_5 = -1.576352$, $C_7 = -1.563472$, $C_{11} = -1.547479$. Dividing $\hat{h}_n$ by $\hat{f}_n = 2n \log p + O(1)$ gives $\kappa_n = (\log p/4)/(2 \log p) + C_p/(2n \log p) + O(p^{-n}) = \frac{1}{8} + C_p/(2n \log p) + O(p^{-n})$. By Proposition 2.3 of 00544, $\varepsilon_{\text{ar}} = 12\kappa - \sigma(1+R) = 12\kappa - 3/2 = 6C_p/(n \log p) + O(p^{-n})$. The Wilms constant is $W = 2\log(2\pi^4) = 10.544133$; $W/\hat{f}_n = W/(2n \log p)$ gives $4.7988/n$ at $p = 3$. Since §9.3 gives $R < 12\kappa/\sigma - 1 + W/(\sigma \hat{f})$, inserting $12\kappa - 1 = \frac{1}{2} + \varepsilon_{\text{ar}}$ yields $R_n < \frac{1}{2} + (\varepsilon_{\text{ar}} + W/\hat{f}_n) = \frac{1}{2} - \frac{3.9150}{n} < 1/2$. □

Theorem 3.12 (Three distinct discriminants; correction of Theorem 3.4)**.** *The quantity $\lambda = \delta(F)/\Delta_{\text{Del}}$ of Theorem 3.4 conflates three different objects:*

symbol	*source*	*nature*
$\delta(F') = \log \lvert\Delta_{F'}\rvert$	*Soulé, Theorem A*	*number-field discriminant, finite*
Δ_{Del}	*00544 Rem. 2.5,* $\Delta_{\text{Del},p} = f_p + m_p - 1$	*Deligne–Saito, finite*
$\Delta_{\min}$	*Frey / Weierstrass*	*model discriminant,* $\nu(\Delta_{\min}) \ge \Delta_{\text{Del}}$
δ_{Fal}	*Wilms master identity*	*Faltings delta, archimedean*

Hence $1 - \lambda$ is not *an archimedean share: δ_{Fal} is the archimedean channel, while $1 - \lambda$ measures the surplus of Δ_{Del} over the number-field discriminant (the $m_p - 1$ components and the model correction). The statements "archimedean share 11/12 / 5/12" and the interpretation of $\lambda_{\text{eq}} > 1$*

as a "negative share" in Theorem 3.4 and Corollary 3.6 are withdrawn. The corrected constraint is an interval*: $R \le C_{\mathrm{BMY}}\lambda$, so BMY ($R \le 3$) requires $\lambda \ge 3/C_{\mathrm{BMY}} = N/12$, while $\sigma \ge 1$ and $\nu(\Delta_{\min}) \ge \Delta_{\mathrm{Del}}$ give $\lambda \lesssim (144/N)/\sigma \le 144/N$; thus*

$$\boxed{\lambda \in \left[\frac{N}{12},\ \frac{144}{\sigma N}\right], \qquad \text{non-empty iff } N^2 \le 1728}.$$

For $\Theta \in \{8, 10\}$ (Theorem Z.4), $N \in \{1, 7\}$ and the intervals are $[1/12, 144]$ and $[7/12, 144/7]$. Theorem 3.4 is downgraded from U *to* C *(normalisation-dependent).*

Proof. Remark 2.5 of 00544 gives $\Delta_{\mathrm{Del},p} = f_p + m_p - 1$ and $\sigma = 1 + \sum_p (m_p - 1)\log p/\log \mathrm{cond} \ge 1$; the terminological warning there states $\nu(\Delta_{\min}) \ge -\mathrm{Art} = \Delta_{\mathrm{Del}}$ and that estimates via polynomial discriminants give only upper bounds for σ. The master identity of 00544 contains δ_{Fal} as the archimedean summand, which is a different object from either discriminant. Therefore $R = \bar{\omega}^2/\Delta_{\mathrm{Del}} \le C_{\mathrm{BMY}}\delta(F')/\Delta_{\mathrm{Del}} = C_{\mathrm{BMY}}\lambda$. Requiring $R \le 3$ gives the lower bound; Theorem AA.1 gives $C_{\mathrm{Szp}} = 144/N$, and with $\log|\Delta_{\min}| \le C_{\mathrm{Szp}} \log N$ one obtains the upper bound. Non-emptiness is $N/12 \le 144/N \Leftrightarrow N^2 \le 1728$; $N = 1, 7$ satisfy it. Values $\lambda > 1$ are legitimate because $\nu(\Delta_{\min}) \ge \Delta_{\mathrm{Del}}$, so the rows $\Theta = 12, 14$ are *not* independently excluded by λ (their exclusion by Theorems Z.4 and U.3 stands on its own). □

Theorem 3.13 (The three inputs of Theorem 3.3 collapse to one hard core)**.** *Under the verbatim statements of 00544, the three upper-bound inputs of Theorem 3.3 reduce as follows.*

1. *$\varepsilon_{\mathrm{ar}}$: §5.2 states "No open upper bound on δ_{Fal} is needed on this tower", with Jorgenson–Kramer's $\delta_{\mathrm{Fal}}(X_0(N)) = O(g)$ giving $d = O(1/\log p) \to 0$. Unconditionally closed on modular towers; open only on general towers.*

2. *C_{tower}: Remark 4.6 states that the step from the Bergman-kernel bound to an upper bound on $\bar{\omega}^2$ "is the substance of that hypothesis" and "Were that step completed unconditionally, Problem 10.1 would be solved". Hence this input* is *Problem 10.1, not an independent subproblem.*

3. *σ: §10.1 states "In every known class $\sigma \to 1$" and labels its boundedness as the missing Szpiro/abc input; §10.4 judges the expected return of further search along the discriminant channel to be negative.*

The function-field table of §10.1 identifies (2) with the missing Bogomolov row and (3) with the missing Szpiro/abc row. Consequently the "3 inputs + 2 realisation problems" of Theorem 3.3 become "1 hard core (general arithmetic BMY) +2 already-closed-on-known towers +2 realisation". This reproduces Theorem P.1 from the tower side; §10.4 of 00544 calls the gap "not a temporary technical gap but a structural one".

Proof. Each item is a verbatim quotation from 00544 (§5.2, Remark 4.6, §§10.1, 10.4) combined with the definitions of Theorem 3.3. The §10.1 table lists exactly two number-field gaps, "missing = Bogomolov" for $R = O(1)$ and "missing = Szpiro/abc" for $\sigma = O(1)$; these are inputs (2) and (3). Input (1) is closed on every tower where the delta-invariant bound is known. Hence only Problem 10.1 remains as a theorem-level open input. □

Theorem 3.14 (Grading of the open ends of 00712)**.** *Of the six items of §8 of 00712:* (1) *the lattice-point count and* (6) *the Kummer form for $\mathcal{I}_p$ together with the $n = 12$ alignment are closed, with §8(6) stating "No closing computation remains registered";* (3) *is closed* by refutation — *the $B_n \sim \frac{3}{4}$ component is a fitting artefact, identically zero once exact S_a and $R(c)$ are used, so the "open" closed derivation of Remark 3.6 is dissolved rather than proved;* (2) *is the only registered open end, but it fixes two* fitting *parameters (γ, σ_α) of the residual kernel at layers $n = 13/14$ and is data-level, not theorem-level;* (4)(5) *are structural descriptions. Exact recomputation: the four*

residuals of Remark 6.6 are $-0.015954, +0.008021, +0.003739, +0.003098$*; the co-layered* $n = 12$ *test gives* $n^*(12) = 5.68408488$ *against* $\frac{1}{2} + 2\bar{\rho}_{\mathrm{mass}}/\log 3 = 5.68413214$*, i.e.* $\Delta = -4.7 \times 10^{-5}$*; the* Δ_9 *factor* 25 *corresponds to the corridor value* 0.2896*.*

Proof. The status labels are quoted from §8 of 00712. Inserting $\bar{\rho}_{\mathrm{mass}} \in \{2.830, 2.841, 2.845, 2.847\}$ into $\frac{1}{2} + 2\bar{\rho}_{\mathrm{mass}}/\log 3$ gives $5.651954, 5.671979, 5.679261, 5.682902$; subtracting from the fitted n^* gives the four residuals verbatim. At $n = 12$, $\bar{\rho}_{\mathrm{mass}} = 2.84767564$ gives 5.68413214. $7.24/25 = 0.2896$, inside the corridor 0.25–0.30. Remark 6.8 of 00712 records that the early 4000-orbit verdict is void and that a χ^2-weighting bug produced $c(10) = 241$; consequently finite-layer statistical extrapolations of $\sigma \to 1$ or C_{tower} carry systematic risk. □

Corollary 3.15 (Four-way grading of the open problems)**.** *The merged list (00544 §10.2, 00712 §8, Remark 2.6) splits as:* A *closed or refuted (00712 §8(1)(3)(6), 00544 10.10, Theorem 2.1, §9.3 cross-check);* B *advanced here (00544 10.7, 3.3 compression,* $\varepsilon_{\mathrm{ar}}$ *rate, 3.4 correction, §8(2) grading);* C *verifiable but not general (00544 10.11 via the* $\sigma = 1 \Leftrightarrow m_p = 1$ *equality classification, 10.5, 00712 §8(4)(5));* D *structural or equivalent to the core conjecture (00544 10.1, 10.2, 10.3, 10.4, 10.6, 10.8/10.12, 10.9). Class D is judged negative-expected-return by §10.4 of 00544, in agreement with Theorem P.1. No class requires, or produces, a new theorem about abc.*

Proof. The classification follows from Theorems 3.9–3.14 and the verbatim status labels of 00544 §10.2. For 10.11, Remark 2.5 of 00544 splits (L1) into Saito (a) and Liu (b) and §10.3 notes that $m_p = 1$ (geometrically integral special fibre, e.g. $X_1(p)$, $p > 7$) gives $\Delta_{\mathrm{Del},p} = f_p$ and $\sigma = 1$ exactly; this is a verification-type attack, not a proof for arbitrary genus. Class D contains Problem 10.1, which Theorem 3.13 shows is the sole remaining hard core. □

Remark 3.16 (Scope). This appendix is a verbatim audit plus exact rational / floating-point recomputation. It does not independently verify the proofs of 00544, 00712, Wilms, Saito, Liu, Jorgenson–Kramer, Dolce–Mercuri, Mayer, Colmez, or Soulé. The $O(1)$ in $\hat{f}_n = 2n \log p + O(1)$ is not verbatim fixed (tag C); a nonzero constant changes the $1/n$ coefficient of Theorem 3.11 but not the order $O(1/n) \to 0$ nor the direction $R_n < 1/2$. The interval of Theorem 3.12 uses the forward Frey conversion of Theorem AA.1 and is a compatibility condition, not an identity. All previous conclusions, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5-18c)$, are unchanged; *abc* remains open.

3.3 Generalising the tower-side results to arbitrary towers

We ask how the results of 00544 and 00712, proved on modular, CM and GL_2-type towers, can be extended to a general tower. Four structures are lost: S1 analyticity of $L(A, s)$ and Mestre's conductor bound; S2 CM / orbit factorisation; S3 the arithmetic input converting systole and Bergman bounds into an upper bound on $\bar{\omega}^2$; S4 totally rational branch points, closed by Evertse's S-unit theorem. The negative-expectation judgement of §10.4 of 00544 concerns attacks along the discriminant channel on S1/S3; it does *not* cover changes of category or of language, which are the subject of this appendix.

Theorem 3.17 (κ is a conductor-weighted average of the isogeny factors)**.** *Let* $f : X \to Y = X/G$ *be a Galois cover of curves over a number field, with* G *finite. Over* $\bar{\mathbb{Q}}$ *there is an isogeny decomposition* $\mathrm{Jac}(X) \sim \bigoplus_{\rho \in \mathrm{Irr}(G)} B_\rho^{m_\rho}$*; Theorem 1.1 of 00712 is the case* $G = (\mathbb{Z}/N)^\times$*. Put* $L_\rho = \log \mathrm{cond}(B_\rho)$ *and* $\kappa_\rho = h_F(B_\rho)/L_\rho$ *(ignore factors with* $L_\rho = 0$*). Then*

$$\kappa(\mathrm{Jac}\, X) = \frac{\sum_\rho m_\rho h_F(B_\rho)}{\sum_\rho m_\rho L_\rho} = \sum_\rho w_\rho \kappa_\rho, \qquad w_\rho = \frac{m_\rho L_\rho}{\sum_\sigma m_\sigma L_\sigma} \geq 0, \quad \sum_\rho w_\rho = 1$$

and consequently $\min_\rho \kappa_\rho \leq \kappa(\mathrm{Jac}\, X) \leq \max_\rho \kappa_\rho$*.*

Proof. The stable Faltings height is an isogeny invariant and is additive on products, because the determinant of cohomology of an exterior tensor product is the tensor product of the determinants; likewise the conductor is an isogeny invariant (the ℓ-adic representations, hence the Artin conductors, agree) and $\log \mathrm{cond}(A \times B) = \log \mathrm{cond}(A) + \log \mathrm{cond}(B)$. Therefore $h_F(\mathrm{Jac}\, X) = \sum_\rho m_\rho h_F(B_\rho)$ and $\log \mathrm{cond}(\mathrm{Jac}\, X) = \sum_\rho m_\rho L_\rho$. Dividing and inserting $h_F(B_\rho) = \kappa_\rho L_\rho$ gives $\sum_\rho w_\rho \kappa_\rho$ with $\sum_\rho w_\rho = 1$; a symbolic check gives $\sum_\rho w_\rho \kappa_\rho - \kappa(\mathrm{Jac}\, X) \equiv 0$. Since all $w_\rho \geq 0$, the average lies in the convex hull, i.e. between the minimum and the maximum of the κ_ρ. For example $\kappa_\rho = (1/12, 1/8, 1/3)$ with $L_\rho = (2, 3, 5)$ gives $\kappa = 53/240 \approx 0.2208 \in [1/12, 1/3]$. □

Corollary 3.18 (Barrier and divergence reduce to irreducible factors)**.** (i) *If every simple isogeny factor of* $\mathrm{Jac}(X_n)$ *satisfies* $\kappa_\rho \leq 1/3$*, then* $\kappa_n \leq 1/3$ *for the whole tower: the one-third barrier is* inherited *from the factors and need not be reproved on the tower. In particular every Galois-cover tower assembled from automorphic or CM factors has an unconditional barrier, extending the §10.7 observation of 00544 that the arithmetic rescue works only on automorphic/CM towers.* (ii) $\kappa_n \to \infty$ *forces* $\max_\rho \kappa_\rho \to \infty$*, with no condition on the weights. Hence Problems 10.1 and 10.2 of 00544 reduce from general curve towers to* simple isogeny factors*; the dangerous class is the non-automorphic simple abelian variety (GL_2-type has measure zero in $\mathcal{A}_g$ for $g \geq 2$).*

Proof. Both statements are immediate from $\kappa \leq \max_\rho \kappa_\rho$ of Theorem 3.17: (i) all $\kappa_\rho \leq 1/3$ gives $\kappa \leq 1/3$; (ii) if all $\kappa_\rho \leq M$ then $\kappa \leq M$, so $\kappa \to \infty$ forces $\max_\rho \kappa_\rho \to \infty$. Theorem 7.8 of 00544 fills $[1/12, 1/3]$ by product convexity; Theorem 3.17 shows this is the same weighted-average mechanism and explains why convexity cannot exceed the barrier. □

Theorem 3.19 (Prym lift: from the abelian-variety category to curves)**.** *Problems 10.8 and 10.12 of 00544 ask for a single-Jacobian realisation of the interior filling, complete only in the abelian-variety category. Rather than assembling $A \times B$ into a Jacobian (the Torelli locus is a low-dimensional subvariety of $\mathcal{A}_g$ for $g \geq 4$), one starts from a curve and takes a Galois cover:*

$$\mathrm{Jac}(X) \sim \mathrm{Jac}(Y) \oplus \bigoplus_{\rho \neq 1} \mathrm{Prym}_\rho\,.$$

For a double cover, Riemann–Hurwitz gives $g_X = 2g_Y - 1 + r/2$ and $\dim \mathrm{Prym} = g_Y - 1 + r/2$. The étale case $r = 0$ gives $\dim \mathrm{Prym} = g_Y - 1$, growing linearly in g_0; the instance of Appendix X is $g(Y_0) = 2$, $g(Y_1) = 3$, $\dim P = 1$ (an elliptic curve), consistent with $\psi_1 = \pm(\mathrm{id} - \nu^*) \circ a$ landing in the Prym elliptic curve and $\Theta|_P = 2\Xi$. Combined with the weighted average of Theorem 3.17, a target interior κ can be split into two known landmark towers whose conductor ratio supplies the weight.

Proof. The isogeny decomposition of $\mathrm{Jac}(X)$ under a finite group action is standard; for $G = \mathbb{Z}/2$ it is $\mathrm{Jac}(Y) \times \mathrm{Prym}(X/Y)$. For a double cover ramified at r points, $2g_X - 2 = 2(2g_Y - 2) + r$, whence $g_X = 2g_Y - 1 + r/2$ and $\dim \mathrm{Prym} = g_X - g_Y = g_Y - 1 + r/2$. An étale double cover has $r = 0$, so $g_X = 2g_Y - 1$; with $g_Y = 2$ this gives $g_X = 3$ and $\dim P = 1$, the Prym elliptic curve of Appendix X. As $g_Y \to \infty$ with r fixed the Prym dimension grows linearly, so irreducible factors of arbitrarily large dimension occur. The final assertion is Theorem 3.17 applied to a two-factor decomposition, with weights $m_\rho L_\rho / (\sum m_\sigma L_\sigma)$. □

Theorem 3.20 (A verifiable classification programme)**.** *For any tower, the factor reduction of Corollary 3.18 gives an executable procedure: list the isogeny factors by* $\mathrm{Irr}(G)$*; test each factor for being automorphic / CM / GL_2-type; only the unclassified simple factors require new analysis. This is strictly cheaper than reproving Problem 10.1 tower by tower, and it is the exact form of the §10.4 claim that "the only tool is effective on the half where the answer is already known". It does not prove Problem 10.1: the factors of a general tower need not be automorphic.*

Proof. By Corollary 3.18(i) the barrier holds on the tower as soon as it holds on each factor; by (ii) any counterexample must contain a factor with diverging κ_ρ. Hence the search for a

counterexample is confined to the unclassified simple factors, which is the asserted reduction. The limitation is that $\mathrm{Irr}(G)$ need not consist of automorphic representations for a general G. □

Theorem 3.21 (Three further routes and their obstructions)**.** (a) Hurwitz distributional language. *Problem 10.3 of 00544 states that the required search space (non-totally-rational branch points or non-hyperelliptic curves) is "mutually exclusive with computability". Following Theorem M.3, one may replace pointwise bounds by density statements on Hurwitz spaces $\mathcal{H}_{g,G}$: Evertse's S-unit finiteness plus Cohen–Lenstra / Ellenberg– Venkatesh–Westerland-type homological stability would yield $\sigma = O(1)$ and $R = O(1)$ in density 1, not pointwise. Obstruction (*C*): homological stability is not proved under wild ramification with a fixed bad set S, while Lemma 2.4 of 00544 shows that the tower world requires wild ramification to dominate. Remark 6.8 of 00712 (the void* 4000*-orbit verdict and the χ^2 bug giving $c(10) = 241$) warns against finite-layer statistical extrapolation.*

(b) Dual-graph Laplacian. *Remark 4.6 of 00544 states that the step from the Bergman-kernel bound to an upper bound on $\bar{\omega}^2$ "is the substance of that hypothesis" and is not supplied. Section 10.1 of 00544 records that the Néron local height pairing is the Laplacian of the dual graph on semistable curves (Werner, Wei; published for Mumford curves). Making the spectral $\to \delta_{\mathrm{Fal}}, \Delta_{\mathrm{Del}}$ chain into an unconditional combinatorial inequality would replace the analytic input by combinatorics, combined with the thickness dictionary $\nu_p = \sum_i N_i$, $\sigma = \sum_i N_i/b_1$ of §10.3. Obstruction (*C*): the spectral constant, and the §10.3(E26) warning that $\Delta_{\mathrm{Del},p} = f_p + m_p - 1$ and the semistable thickness reading are not interchangeable outside semistability.*

(c) The weaker target $R = O(1)$. *Section 9.5 of 00544 states that arithmetic Bogomolov $R = O(1)$ is strictly weaker than BMY and that "the two can be separated; they need not share a fate"; Theorem 9.1 gives $\bar{\omega}^2 \le 12h_F + 2\log(2\pi^4)g$ unconditionally. Hence Problem 10.1 can be graded into $R = O(1)$ first, then $R \le 3$. Obstruction: §9.4 of 00544 judges the reverse direction unreachable and the tools missing, and the reduction already appears in Parshin (1988). Low expected return.*

Proof. (a) is the structural analogy with Appendix M (CRT plus moment estimates) applied to Hurwitz spaces; the obstruction is the conjunction of Lemma 2.4 and the scope of the homological-stability theorems. (b) combines the two verbatim statements of 00544 cited; the obstruction is the E26 semistability warning. (c) is §§9.1, 9.4, 9.5 of 00544 verbatim; the grading follows from $R \le 3 \Rightarrow R = O(1)$, the converse failing. □

Corollary 3.22 (Two corrections to Section 3.2)**.** (i) *In Remark 3.16 / Theorem 3.14 the claim that layers $n = 13, 14$ of 00712 are not enumerable confused triples with orbits. By Theorem 2.1 of 00712, $M_{\mathrm{top}}(n) = 4 \cdot 3^{n-1} - 3$ and, the G_n-action being free, $|Y_n|/|G_n| = M_{\mathrm{top}}(n)$. Exact values: $n = 7$ gives 2913; $n = 9$ gives 26241; $n = 12$ gives 708585; $n = 13$ gives 2125761; $n = 14$ gives 6377289 orbits. These are of the same order as the 1.4×10^6 orbits already enumerated at $n = 9$ in §7 of 00712, so a co-layered measurement at $n = 13/14$ is* computable*. The reason not to prioritise it is that the gain is two data-level fitting parameters (γ, σ_α) and §8(6) declares "No closing computation remains registered".*

(ii) *The Prym instance of Appendix X is an étale double cover ($r = 0$), not a ramified hyperelliptic one; with $g(Y_0) = 2$ it gives $\dim P = 1$, as used in Theorem 3.19.*

Proof. (i) Direct computation: $4 \cdot 3^{12} - 3 = 2125761$, $4 \cdot 3^{13} - 3 = 6377289$, while $|Y_{13}| = (3^{13} - 1)(3^{13} - 2) - (3^{12} - 1)(3^{12} - 2) = 2259433103202$ and $|G_{13}| = 2 \cdot 3^{12} = 1062882$, whose quotient is 2125761. (ii) is Riemann–Hurwitz with $r = 0$: $g_X = 2g_Y - 1 = 3$ and $\dim \mathrm{Prym} = g_X - g_Y = 1$. □

Corollary 3.23 (Grading and the precise scope of the negative-expectation judgement)**.** *The six routes split as:* A *already proved (*U*): Theorem 3.17 and Corollary 3.18;* B *recommended (*U+C*): Prym lift (Theorem 3.19) and the classification programme (Theorem 3.20);* C *feasible but conditional (*C*): Hurwitz distributions, dual-graph Laplacian, $R = O(1)$;* D *executable with small gain (Corollary 3.22(i)): the $n = 13/14$ measurement. The negative-expectation judgement*

of §10.4 of 00544 covers only attacks along the discriminant channel and direct assaults on the BMY constant; it does not *cover the changes of category (A, B) or of language/data (C, D). Theorem P.1 is a completeness statement for the* pointwise abc side*; the tower side retains room under a change of category or language, but the output is not abc.*

Proof. The grading collects Theorems 3.17–3.21 and Corollary 3.22. §10.4 of 00544 concerns the discriminant channel (D) and Problem 10.1; §10.7 restricts the arithmetic rescue to automorphic/CM towers, which Corollary 3.18 extends to Galois-cover towers built from such factors. The scope statement follows. All previous conclusions, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5-18c)$, are unchanged and *abc* remains open. □

Remark 3.24 (Scope). This appendix is a route analysis plus exact rational recomputation. The isogeny invariance of the stable Faltings height and of the conductor, the $\bar{\mathbb{Q}}$-isogeny decomposition under a finite group, the Torelli-locus and Hurwitz existence questions, homological stability, the dual-graph spectral constants and Parshin's reduction are cited standard facts or explicitly labelled C; none is independently re-proved. The identity $|Y_n|/|G_n| = M_{\mathrm{top}}(n)$ uses the free-action theorem of 00712 for $p = 3$; a general G with fixed points needs a Burnside correction. We do not independently verify the proofs of 00544, 00712, Faltings, Colmez, Evertse, Wilms or Soulé.

4 The Prym lift and weight locking

4.1 The Prym lift: weight locking, tame exclusion, and the exact reduction of Problems 10.8/10.12

Appendix 3.3 proposed, as Route A, to attack Problems 10.8 and 10.12 of 00544 by starting from a curve and taking a Galois cover, so that $\mathrm{Jac}(X)$ decomposes into $\mathrm{Jac}(Y)$ and Prym factors, and by tuning the conductor ratio of two landmark towers to realise an interior value of κ. We carry this out. The decomposition mechanism is correct, but three substantial restrictions emerge: the étale/tame lift never enters the tower world; the weight is not a free parameter but a Swan share, forced to 1 when the base tower is tame; and Route A contributes nothing to Problem 10.1, reducing 10.8/10.12 instead to a Prym–Torelli question. On the positive side, condition (T2) is settled unconditionally, and Corollary 3.18 receives an unconditional instance on the Fermat tower.

Theorem 4.1 (Two-factor Prym decomposition)**.** *Let $f : X \to Y = X/G$ be a double cover ($G = \mathbb{Z}/2$) over a number field, with r branch points. Then $H^1(X) = H^1(Y) \oplus V_-$ with $\dim V_- = 2\dim \mathrm{Prym}$, and, the Artin conductor exponent being additive on direct sums of representations,*

$$L_X := \log \mathrm{cond}(\mathrm{Jac}\, X) = L_Y + L_P, \qquad L_P := \log \mathrm{cond}(\mathrm{Prym}).$$

The stable Faltings height is likewise additive on the isogeny product. Hence

$$\boxed{\kappa_X = \frac{L_Y\kappa_Y + L_P\kappa_P}{L_Y + L_P} = (1 - w_P)\kappa_Y + w_P\kappa_P, \qquad w_P = \frac{L_P}{L_X} = 1 - \frac{L_Y}{L_X}}$$

and Riemann–Hurwitz gives $g_X = 2g_Y - 1 + r/2$, $\dim \mathrm{Prym} = g_Y - 1 + r/2$.

Proof. The ℓ-adic representation of G on $H^1(X, \mathbb{Q}_\ell)$ splits into the invariant part, identified with $H^1(Y)$, and the (-1) eigenspace V_-; the latter is the Tate module of the Prym variety, of dimension $2(g_X - g_Y)$. The conductor exponent $f_p = \dim(V/V^{I_p}) + \mathrm{Sw}_p$ is additive on direct sums, so $f_p(X) = f_p(Y) + f_p(V_-)$ and hence $L_X = L_Y + L_P$. The determinant of cohomology of an exterior tensor product is the tensor product of the determinants, so $h_F(A \times B) = h_F(A) + h_F(B)$; this is stable under isogeny. Dividing the sum of heights by the sum of conductors gives the

displayed weighted average; a symbolic check gives $\kappa_X - [(1-w_P)\kappa_Y + w_P\kappa_P] \equiv 0$. Finally $2g_X - 2 = 2(2g_Y - 2) + r$ gives $g_X = 2g_Y - 1 + r/2$ and $\dim \mathrm{Prym} = g_X - g_Y = g_Y - 1 + r/2$. For the étale cover $r = 0$, $g(Y_0) = 2$ of Appendix X this gives $g_X = 3$, $\dim P = 1$, the Prym elliptic curve consistent with $\psi_1 = \pm(\mathrm{id} - \nu^*) \circ a$ and $\Theta|_P = 2\Xi$. □

Theorem 4.2 (Tame and étale Prym lifts never enter the tower world)**.** *Fix a finite bad set S and put $C_S = \sum_{p\in S} \log p$. If the cover is tame at every $p \in S$ (in particular if it is étale), then $\mathrm{Sw}_p = 0$ and $f_p(V) \le \dim V$, so*

$$L_Y \le 2g_Y C_S, \qquad L_P \le 2\dim \mathrm{Prym}\ C_S = 2(g_Y - 1 + r/2)C_S,$$

and therefore

$$\boxed{\hat{f}_n = \frac{L_X}{g_X} \le \frac{4g_Y - 2 + r}{2g_Y - 1 + r/2}\, C_S = 2C_S + O(g_Y^{-1}) = O(1)}.$$

Consequently condition (T3) *$\hat{f}_n \to \infty$ of 00544 fails, independently of $g_Y \to \infty$.*

Proof. By Lemma 2.4 of 00544, $f_p = \dim(V/V^{I_p}) + \mathrm{Sw}_p$ with $\dim(V/V^{I_p}) \le \dim V = 2g$; in the tame case $\mathrm{Sw}_p = 0$. Applying this to $H^1(Y)$ (dimension $2g_Y$) and to V_- (dimension $2\dim \mathrm{Prym}$) and summing over the finite set S gives the two bounds. Their sum divided by $g_X = 2g_Y - 1 + r/2$ is $(4g_Y - 2 + r)C_S/(2g_Y - 1 + r/2) = 2C_S + O(1/g_Y)$. Exact evaluations: $S = \{2\}$ gives the bound $2\log 2 = 1.386294$; $S = \{2,3,5\}$ gives $2\log 30 = 6.802395$; $S = \{2,3,5,7\}$ gives $2\log 210 = 10.694237$; all are attained already at $g_Y = 2$ and do not change at $g_Y = 1000$. Since $\hat{f}_n$ stays bounded, (T3) fails. □

Corollary 4.3 (Correction to Appendix EE)**.** *The assertion in Theorem 3.19 that the étale case $\dim \mathrm{Prym} = g_Y - 1$ "grows linearly in g_0, so irreducible factors of arbitrarily large dimension occur" and hence may be used to promote the lift to a tower is valid only for the fixed-genus pointwise construction of Appendices X–Z of the pointwise companion file. As a $g_n \to \infty$ tower it is excluded by Theorem 4.2. This is a second independent instance of the disjointness theorem of Section 2.1: the two sides use different types of Prym covers (pointwise tame, tower-side necessarily wild) and cannot inherit each other's conclusions.*

Proof. The dimension statement is Riemann–Hurwitz and is correct in isolation; Theorem 4.2 shows that any tower built from tame/étale covers violates (T3). Section 2.1 proved $\Theta^*(g_0) \sim 8g_0^3$, so no family can simultaneously start the pointwise coefficient chain and satisfy $g_n \to \infty$; the present corollary derives the same separation from the wild-necessity lemma instead of the DRR chain, hence independently. □

Theorem 4.4 (Wild dominance: a 0–1 law for the weight)**.** *For $G = \mathbb{Z}/2$ the cover is tame at every odd prime, so wild ramification can occur only at $p = 2$. Hence condition* (T2) *of 00544 (a fixed bad set) is satisfied unconditionally by any fixed $S \ni 2$: this removes one of the three obstructions listed in Theorem 3.19.*

Assume the base tower is tame ($\mathrm{Sw}_Y = 0$) and set $d_n := \mathrm{Sw}_2(V_-)/g_X(n)$, so (T3) *requires $d_n \to \infty$. Writing $L_P = L_P^{\mathrm{tame}} + \mathrm{Sw}_P$ and $L_X = L_Y + L_P^{\mathrm{tame}} + \mathrm{Sw}_P$,*

$$1 - w_P = \frac{L_Y}{L_X} = \frac{2g_Y C_S}{\mathrm{Sw}_P}\big(1 + O(g_Y/\mathrm{Sw}_P)\big) = O(1/d_n) \longrightarrow 0,$$

and consequently

$$\boxed{\kappa_X(n) = \kappa_P(n) + O(1/d_n)}.$$

More generally, when the base is wild too,

$$\kappa_X = s_n\kappa_P + (1 - s_n)\kappa_Y + o(1), \qquad s_n = \frac{\mathrm{Sw}_P}{\mathrm{Sw}_P + \mathrm{Sw}_Y},$$

so the weight of Theorem 3.17 is the Swan share.

Proof. For a double cover, $|G| = 2$ is invertible at every odd p, so the higher ramification groups are trivial there and $\mathrm{Sw}_p = 0$; only $p = 2$ can contribute. With $\mathrm{Sw}_Y = 0$ the whole divergent part of the conductor lies on V_-. Since $L_Y \le 2g_Y C_S$ is tame-dominated and $\mathrm{Sw}_P = d_n g_X C_S$ with $g_X \asymp 2g_Y$, we get $L_Y/\mathrm{Sw}_P = O(1/d_n)$; expanding $1 - w_P = L_Y/L_X$ and using $L_X \ge \mathrm{Sw}_P$ gives the estimate, and substitution into Theorem 4.1 gives $\kappa_X = \kappa_P + (1 - w_P)(\kappa_Y - \kappa_P) = \kappa_P + O(1/d_n)$. Numerically with $g_Y = 100$, $S = \{2\}$, one obtains $w_P = 0.665, 0.749, 0.856, 0.916, 0.954, 0.981, 0.990, 0.998$ at $d = 1, 2, 5, 10, 20, 50, 100, 500$. The general form is Theorem 4.1 with both conductors split into tame and Swan parts; the tame parts are $O(g)$ and are $o(\mathrm{Sw})$, leaving the stated share s_n. □

Theorem 4.5 (Endpoints cannot be spliced; the only regime with free weight)**.** *Given landmarks $\kappa_Y < \kappa_P$ and a target κ^*, the required weight and Swan depth are*

$$w_P = \frac{\kappa^* - \kappa_Y}{\kappa_P - \kappa_Y}, \qquad \frac{L_X}{L_Y} = \frac{1}{1 - w_P}, \qquad \frac{\mathrm{Sw}_P}{L_Y} \approx \frac{w_P}{1 - w_P}.$$

For $\kappa_Y = 1/12$, $\kappa_P = 1/3$ the exact rational values are $(\kappa^*, w_P, L_X/L_Y, \mathrm{Sw}_P/L_Y) = (1/12, 0, 1, 0)$, $(1/8, 1/6, 6/5, 1/5)$, $(1/6, 1/3, 3/2, 1/2)$, $(5/24, 1/2, 2, 1)$, $(1/4, 2/3, 3, 2)$, and $\kappa^* \to 1/3$ forces $w_P \to 1$ and $\mathrm{Sw}_P/L_Y \to \infty$. Since Theorem 4.4 forces $w_P \to 1$ under (T3) *with tame base, the lower part of the interior cannot be obtained by splicing a tame base with a modular Prym factor: reaching $\kappa^* = 1/12$ requires $\kappa_P = 1/12$ already. A continuously tunable weight survives only when both the base and the Prym factors are wild with comparable Swan depths.*

Proof. Solve w_P from the linear relation of Theorem 4.1; the conductor ratio follows from $w_P = 1 - L_Y/L_X$, and the Swan depth from $L_X \approx L_Y + \mathrm{Sw}_P$. The table is exact rational arithmetic. Theorem 4.4 gives $w_P \to 1$, so the only attainable limit is κ_P; for κ^* strictly below κ_P one needs either $\kappa_P = \kappa^*$ or a second wild tower whose Swan depth competes with the Prym one. Writing $\mathrm{Sw}_Y \sim a_Y g_Y \log g_Y$, $\mathrm{Sw}_P \sim a_P g_X \log g_X$ with $g_X \approx 2g_Y$ gives $s = 2a_P/(2a_P + a_Y)$, which ranges over $(0, 1)$ only if a_P/a_Y is tunable; the matching of conductor growth rates is labelled C. □

Lemma 4.6 (The branch number is bounded)**.** *For a fixed bad set S and totally rational branch points, Theorem 8.8 of 00544 (Evertse's S-unit finiteness) forces $r_n \le R(S)$. Hence* (T1) *$g_X(n) \to \infty$ can only be driven by $g_Y(n) \to \infty$, and $\dim \mathrm{Prym}_n / g_X(n) \to 1/2$: the tame share of the weight tends to $1/2$ and supplies no tunable parameter.*

Proof. $g_X = 2g_Y - 1 + r/2$ with $r \le R(S)$ gives $g_X = 2g_Y + O(1)$, so $g_X \to \infty$ iff $g_Y \to \infty$, and $\dim \mathrm{Prym}/g_X = (g_Y - 1 + r/2)/(2g_Y - 1 + r/2) \to 1/2$; values $0.3333, 0.4737, 0.4975, 0.49975$ at $g_Y = 2, 10, 100, 1000$. Since $L_Y \le 2g_Y C_S$ and $L_P^{\mathrm{tame}} \le 2 \dim \mathrm{Prym}\, C_S$, the tame weights are asymptotically equal. This is the quantitative form of Problem 10.3 of 00544: abandoning totally rational branch points is exactly abandoning computability. □

Theorem 4.7 (No contribution to Problem 10.1; exact reduction of 10.8/10.12)**.** (i) *By the convex hull property of Theorem 3.17, $\kappa_X \le \max(\kappa_Y, \kappa_P)$. If both factors satisfy the one-third barrier, so does the whole tower. Hence Route A contributes nothing to Problem 10.1: any violation requires $\kappa_P > 1/3$, i.e. a non-automorphic simple Prym factor, precisely the dangerous class identified in Corollary 3.18.*

(ii) *Theorem 7.8 of 00544 fills $[1/12, 1/3]$ densely in the abelian-variety category by product convexity, while 10.8/10.12 ask for realisation by a single Jacobian. By Theorem 4.4, $\kappa_X = \kappa_P + o(1)$ for a tame base. Consequently 10.12 is reduced exactly to: given the product abelian varieties A_n of Theorem 7.8 realising κ^*, can A_n be realised as a* Prym *direct factor of the Jacobian of some curve Galois cover, with the required conductor growth? This is a Prym–Torelli problem and remains open (C); Route A does not bypass 10.12 but locates its obstruction precisely.*

Proof. (i) is Theorem 3.17 with two factors. (ii) combines $\kappa_X = \kappa_P + O(1/d_n)$ of Theorem 4.4 with the statement of Theorem 7.8 of 00544; the reduction is immediate, and the Prym–Torelli and simplicity questions are open by §10.12 of 00544. □

Theorem 4.8 (Unconditional instance of barrier inheritance; obstruction ledger)**.** (a) *By Theorem 1.1 of 00712,* $J(F_N) \sim \bigoplus_{\mathcal{O}} A_{\mathcal{O}}$ *with each* $A_{\mathcal{O}}$ *a CM factor of type* $\mathbb{Q}(\zeta_{d_0})$*. CM abelian varieties are potentially automorphic, so every factor satisfies* $\kappa_\rho \le 1/3$*; Corollary 3.18(i) then gives* $\kappa_{\mathrm{Fermat}} \le 1/3$ *unconditionally, without arithmetic BMY. This agrees with* $\kappa_{\mathrm{Fermat}} = 1/8$ *(Theorem 3.4 of 00712) and with* $\kappa = (1+R)/12 = 1/8$ *at* $R = 1/2$*,* $\sigma = 1$ *(§2.4 of 00544). Thus the orbit factorisation of 00712 is a ready-made instance of the weighted average of Theorem 3.17.*

(b) *The three obstructions of Theorem 3.19 reduce to two and a half:* (T2) *is settled unconditionally by Theorem 4.4;* C1 *is Prym–Torelli existence and simplicity;* C2 *is the existence of a wild double cover* $y^2 = f_n(x)$ *with prescribed Swan depth* $d_n \to \infty$ *(controlled by the* 2*-adic Newton polygon / cluster picture,* C*);* C2.5 *is the matching of Swan growth rates, needed only for a continuously tunable weight.*

Proof. (a) Potential automorphy of CM abelian varieties is standard; the barrier on automorphic factors is the modular-tower case of §5.2 of 00544, where the BMY input is a theorem (Dolce–Mercuri, Mayer). The conclusion follows from Corollary 3.18(i); the two numerical values of κ agree since $(1 + 1/2)/12 = 1/8$. (b) collects Theorems 4.4, 4.5, 4.7. □

Remark 4.9 (Scope and consequences). Six revisions to Section 3.3 follow: R1 the étale promotion is withdrawn; R2 "weight tuned by conductor ratio" becomes "weight is the Swan share, locked to 1 for tame base"; R3 Route A is regraded from "recommended U+C" to "U part complete, residual C is Prym–Torelli"; R4 the obstruction table loses (T2); R5 Section 2.1 gains a second independent disjointness instance; R6 Open problem (13) records that Route A has zero contribution to 10.1.

The isogeny invariance of h_F and of the conductor, the additivity of Artin conductor exponents, the potential automorphy of CM factors, Evertse's theorem, and the 2-adic control of Sw_2 are cited standard results (C); the weighted-average identity, the tame bound, the 0–1 law, the rational table and $\dim \mathrm{Prym}/g_X \to 1/2$ are U. The competition between $m_p - 1$ and Sw_p in $\sigma = \Delta_{\mathrm{Del}}/\log \mathrm{cond}$ (§10.3(E26) of 00544) is not treated, so s_n is a necessary Swan-share relation rather than an exact identity. We do not independently verify the proofs of 00544, 00712, Faltings, Colmez, Evertse, Wilms or Soulé. All conclusions of the first 39 rounds, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5 - 18c)$, are unchanged; *abc* remains open.

5 Exact computations on the Fermat tower

5.1 The weighted average on the Fermat tower: exact layer census and the layer-local $\kappa = \frac{1}{8}$ law

The last item of Remark 4.9 proposed to make Theorem 4.8(a) concrete by computing the individual κ_ρ from the Colmez closed form of 00712 and verifying that the weighted average of Theorem 3.17 gives $\frac{1}{8}$. We carry this out at the level of *conductor layers*, which is the level at which 00712 supplies exact data. The result is stronger than expected: the census offset $n - \frac{1}{p^2-1}$ is κ-neutral, the slope constants $\alpha = \frac{1}{4}$ and $a = 2$ are independent of p, and the convergence $\kappa_n \to \frac{1}{8}$ from below is *layer-local*. An initial reading of the leading term of $\sum_c cA_c$ gave $s_n \sim np^2/(p^2-1)$; exact summation corrects this to $s_n = n - \frac{1}{p^2-1} + O(np^{-n})$, and all subsequent computations use the corrected value.

Theorem 5.1 (Exact layer census for every odd prime)**.** *Let* $F_{p^n} : x^{p^n} + y^{p^n} = 1$*, let* $\mathcal{O}$ *run over the* $G_n = (\mathbb{Z}/p^n\mathbb{Z})^\times$*-orbits of Theorem 1.1 of 00712, and let* $d_0(\mathcal{O}) = p^c$ *with* $1 \le c \le n$*. Then the number of orbits at level* c*, their dimension, and their total dimension are*

$$M_c = p^{c-1}(p+1) - 3, \qquad \dim_c = \frac{\varphi(p^c)}{2} = \frac{p^{c-1}(p-1)}{2},$$
$$A_c = M_c \dim_c = \tfrac{1}{2}(p^2-1)p^{2c-2} - \tfrac{3}{2}(p-1)p^{c-1}$$

and

$$\sum_{c=1}^{n} A_c = g_{p^n} = \frac{(p^n-1)(p^n-2)}{2}.$$

Consequently the dimension share of a fixed level is $\beta_c = A_c/g_{p^n} = O(p^{2c-2n})$*, and* $\beta_n \to 1-p^{-2}$.

Proof. By the Burnside argument of Theorem 2.1 of 00712, the action of $G_c = (\mathbb{Z}/p^c\mathbb{Z})^\times$ on $Y_c = \{(r,s,t) \in [1,p^c)^3 : r+s+t \equiv 0,\ p \nmid \gcd(r,s,t)\}$ is free, so the number of primitive orbits at level c is $|Y_c|/|G_c|$. Now $|Y_c| = (p^c-1)(p^c-2) - (p^{c-1}-1)(p^{c-1}-2) = (p-1)[p^{2c-2}(p+1) - 3p^{c-1}]$ and $|G_c| = p^{c-1}(p-1)$; dividing gives $M_c = p^{c-1}(p+1) - 3$. Theorem 1.1 of 00712 gives $\dim A_{\mathcal{O}} = \varphi(d_0)/2$, hence $\dim_c = p^{c-1}(p-1)/2$. Expanding the product gives A_c. Finally $\sum_{c\le n} \frac{1}{2}(p^2-1)p^{2c-2} = \frac{1}{2}(p^{2n}-1)$ and $\sum_{c\le n} \frac{3}{2}(p-1)p^{c-1} = \frac{3}{2}(p^n-1)$, so $\sum A_c = \frac{1}{2}(p^{2n}-3p^n+2) = g_{p^n}$. Since $A_n = \frac{1}{2}(p^2-1)p^{2n-2} + O(p^n)$ and $g_{p^n} = \frac{1}{2}p^{2n} + O(p^n)$, $\beta_n \to (p^2-1)/p^2 = 1-p^{-2}$. Numerical checks: $p=3, n=7$ gives $M_7 = 2913$, $\dim_7 = 729$, $g = 2388205$, all verbatim in Corollary 2.2 of 00712; $g(F_{729}) = 264628$ (its corrected value); $p=5, n=1$ gives three orbits of dimension 2, totalling $g(F_5) = 6$. □

Lemma 5.2 (Exact first moment of the layer index)**.** *With* $S_n = \sum_{c=1}^{n} cA_c$ *and* $s_n = S_n/g_{p^n}$*,*

$$S_n = \frac{np^{2n}}{2} - \frac{p^{2n}-1}{2(p^2-1)} - \frac{3np^n}{2} + \frac{3(p^n-1)}{2(p-1)}, \qquad \boxed{s_n = n - \frac{1}{p^2-1} + O(np^{-n})}.$$

Proof. Use $\sum_{c=1}^{n} cx^{c-1} = [nx^n(x-1) - x^n + 1]/(x-1)^2$ twice, with $x = p^2$ and $x = p$, and multiply by the two coefficients of A_c. Dividing by $g_{p^n} = \frac{1}{2}(p^{2n}-3p^n+2)$ gives $s_n = n + \big[-(p^{2n}-1)/(p^2-1) + 3(p^n-1)/(p-1)\big]/(p^{2n}-3p^n+2)$, whose second term tends to $-1/(p^2-1)$ with error $O(np^{-n})$. Exact rational checks: $p=3, n=12$ gives $s_n - 12 + 1/8 = 2.117 \times 10^{-6}$; $p=11, n=12$ gives 8.76×10^{-14}. □

Theorem 5.3 (Forced layer constants; the census offset is κ-neutral)**.** *Write the per-dimension height and conductor density of level* c *as* $\hat{h}_c = \alpha c \log p + \gamma_p$ *and* $\hat{f}_c = ac\log p + b_p$ *(the layer dependence of* $\hat{f}_c$ *only through* c *is labelled* C*: all characters in the CM type of* $A_{\mathcal{O}}$ *have conductor* p^c*). By additivity,* $\sum_c \beta_c \hat{h}_c = \hat{h}_n = \frac{\log p}{4} n + C_p + o(1)$ *(Theorem 3.4 of 00712) and* $\sum_c \beta_c \hat{f}_c = \hat{f}_n = 2n\log p + O(1)$ *(Theorem 3.2 of 00544, the conductor density* $c = 2$*). Using Lemma 5.2,*

$$\boxed{\alpha = \tfrac{1}{4}, \quad a = 2, \quad \gamma_p = C_p + \frac{\log p}{4(p^2-1)} = -6\log A - \frac{\log p}{4(p+1)}, \quad b_p = \frac{2\log p}{p^2-1}}$$

with $C_p = -6\log A - \frac{p\log p}{4(p^2-1)}$ *and* A *the Glaisher–Kinkelin constant. Hence*

$$\boxed{8\gamma_p - b_p = 8C_p}, \qquad \bar{\kappa}_c := \frac{\hat{h}_c}{\hat{f}_c} \longrightarrow \frac{\alpha}{a} = \frac{1}{8} \quad \textit{for every } p \textit{ and every } c.$$

Proof. $\sum_c \beta_c \hat{h}_c = \alpha s_n \log p + \gamma_p = \alpha n \log p + \gamma_p - \frac{\alpha \log p}{p^2-1} + o(1)$; matching $\frac{\log p}{4} n + C_p$ forces $\alpha = \frac{1}{4}$ and $\gamma_p = C_p + \frac{\log p}{4(p^2-1)}$. Inserting C_p gives $\gamma_p = -6\log A - \frac{(p-1)\log p}{4(p^2-1)} = -6\log A - \frac{\log p}{4(p+1)}$. Similarly $\sum_c \beta_c \hat{f}_c = as_n \log p + b_p = an\log p + b_p - \frac{a\log p}{p^2-1} + o(1) = 2n\log p + O(1)$ forces $a = 2$ and $b_p = \frac{2\log p}{p^2-1}$ (the last equality uses the normalisation $\hat{f}_n = 2n\log p$ with no extra constant, C). Then $8\gamma_p - b_p = 8C_p + \frac{2\log p}{p^2-1} - \frac{2\log p}{p^2-1} = 8C_p$: the census offset $-1/(p^2-1)$ enters height and conductor in the exact ratio $8:1$ and cancels in κ. Finally $\bar{\kappa}_c = \alpha/a + O(1/c) = 1/8 + O(1/c)$. Exact numerical check of $8\gamma_p - b_p = 8C_p$ at $p = 3, 5, 7, 11, 13$ gives residual $\le 2 \times 10^{-15}$. The independence of α, a from p explains why both $\kappa_{\text{Fermat}} = 1/8$ and the conductor density $c = 2$ are universal constants. □

Theorem 5.4 (The convergence from below is layer-local)**.** *For every $c \geq 1$,*

$$\boxed{\bar{\kappa}_c = \frac{1}{8} + \frac{C_p}{2c\log p + b_p}, \qquad \kappa_n = \frac{\hat{h}_n}{\hat{f}_n} = \frac{1}{8} + \frac{C_p}{2n\log p}}.$$

Since $C_p < 0$, every layer satisfies $\bar{\kappa}_c < \frac{1}{8}$, and the $1/n$ correction at the level of the whole Jacobian is the same functional of n as the $1/c$ correction of a single layer. Moreover $\hat{h}_n - \left(\frac{\log p}{4}n + C_p\right) = O(np^{-n})$ and $\hat{f}_n - 2n\log p = O(np^{-n})$.

Proof. Multiply $\hat{h}_c$ by 8: $8\hat{h}_c = 2c\log p + 8\gamma_p = 2c\log p + b_p + 8C_p$ by Theorem 5.3. Dividing by $\hat{f}_c = 2c\log p + b_p$ gives $\bar{\kappa}_c = \frac{1}{8} + \frac{C_p}{2c\log p + b_p}$. With $\hat{f}_n = 2n\log p$ and $8\hat{h}_n = 2n\log p + 8C_p$, the same computation gives κ_n. The $O(np^{-n})$ remainders are exactly the error term of Lemma 5.2; numerically at $p = 3, n = 12$ they are 5.8×10^{-7} and 4.7×10^{-6}. Table (model): $p = 3$ gives $\bar{\kappa}_c = -0.520, -0.107, 0.036, 0.089, 0.111$ at $c = 1, 3, 8, 20, 50$; $p = 11$ gives $-0.195, 0.018, 0.085, 0.109, 0.119$. The anchor of 00712, $h_F(j = 0) = -1.667691$ and $\log \mathrm{cond}(27) = 3\log 3 = 3.295837$, gives $\bar{\kappa}_1 = -0.505999$ against the model -0.520470: same sign and order, the difference being the failure of the affine law at $c = 1$ (Remark 2.3 of 00712 lists the $d_0 = 3$ orbit as a deep-kernel exception). Since $\beta_1 = O(p^{-2n})$, that deviation contributes $o(1)$ to $\hat{f}_n$. □

Remark 5.5 (Relation to the ρ-law and to Proposition 3.5). Theorem 4.2 of 00712 states that the layer means of $\rho_\chi = R_\chi + \log f_\chi$ are constant, $\bar{\rho} = 2.4194$, with layerwise standard deviation shrinking to 0.57. In the spectral identity $h(\Phi) = \mu(\Phi)/2 - \sum_\chi m(\Phi, \chi)\rho_\chi$ the only c-dependence is then $\mu(\Phi) \propto c\log p$, so $\hat{h}_c$ is affine in c: the ρ-law is the spectral counterpart of the ansatz of Theorem 5.3, while Theorem 5.3 determines the slope $\frac{1}{4}$ and the constant γ_p without resolving the factor-2 discrepancy between (1.3) and (4.1) of 00712. Also $C_p = \gamma_p - \frac{\log p}{4(p^2-1)}$ splits the constant of Proposition 3.5 of 00712 into a layer constant plus the census offset; the two pieces can be checked separately.

Theorem 5.6 (Explicit verification of the weighted average)**.** *Define the conductor weights $W_c = A_c\hat{f}_c/(g_{p^n}\hat{f}_n)$, so $\sum_c W_c = 1$. Then*

$$\boxed{\kappa(\mathrm{Jac}\, F_{p^n}) = \frac{\sum_c A_c\hat{h}_c}{\sum_c A_c\hat{f}_c} = \sum_{c=1}^{n} W_c\bar{\kappa}_c}$$

holds to machine precision for all tested (p, n): the discrepancy is $\leq 7 \times 10^{-18}$ for $(3, 7)$, $(3, 9)$, $(3, 12)$, $(5, 4)$, $(7, 3)$, $(11, 3)$. Moreover $W_n \to 1 - p^{-2}$ (numerically 0.898–0.995), $W_1 = O(p^{-2n})$, and

$$\kappa_n - \bar{\kappa}_n = \frac{C_p}{2n^2(p^2-1)\log p} + O(n^{-3}).$$

Proof. The identity is Theorem 3.17 applied to the layer decomposition of Theorem 5.1: heights and conductors are additive on the isogeny product, so $\kappa = (\sum_c A_c\hat{h}_c)/(\sum_c A_c\hat{f}_c) = \sum_c \frac{A_c\hat{f}_c}{\sum_d A_d\hat{f}_d} \cdot \frac{\hat{h}_c}{\hat{f}_c} = \sum_c W_c\bar{\kappa}_c$; the numerical check is a direct exact-rational evaluation of both sides. Since $A_c = O(p^{2c})$, the weights concentrate at the top level, $W_n \approx \beta_n \to 1 - p^{-2}$, and $W_1 \approx \beta_1 = O(p^{-2n})$. Finally $\kappa_n - \bar{\kappa}_n = C_p\left[\frac{1}{2n\log p} - \frac{1}{2n\log p + b_p}\right] = \frac{C_p b_p}{(2n\log p)^2} + O(n^{-3}) = \frac{C_p}{2n^2(p^2-1)\log p} + O(n^{-3})$, using $b_p = 2\log p/(p^2 - 1)$; numerically -6.24×10^{-4} observed against -6.30×10^{-4} predicted at $p = 3, n = 12$, and -1.81×10^{-3} against -1.85×10^{-3} at $p = 3, n = 7$. □

Theorem 5.7 (Quantitative barrier inheritance without potential automorphy)**.** (i) *In the model,* $\max_c \bar{\kappa}_c = 1/8 < 1/3$ *and* $\bar{\kappa}_c < 1/8$ *for every c because $C_p < 0$.* (ii) *By Theorem 4.2 of 00712*

the layerwise standard deviation of ρ_χ is 0.57, so the fluctuation of $h_F(A_\mathcal{O})$ inside a level is $\lesssim 0.57\sqrt{\dim_c}$, while $\log \mathrm{cond}(A_\mathcal{O}) \approx \dim_c \cdot 2c \log p$; hence

$$\delta\kappa_\rho \lesssim \frac{0.57}{2c \log p \sqrt{\dim_c}}.$$

Explicitly $p = 3$: $(c, \dim_c, \delta\kappa_\rho) = (1, 1, 0.260)$, $(3, 9, 0.029)$, $(7, 729, 1.4 \times 10^{-3})$, $(12, 177147, 5.1 \times 10^{-5})$; $p = 5, c = 12$ gives 1.5×10^{-6}. Consequently $\kappa_\rho < 1/3$ at every layer and for every orbit, and low-c outliers carry exponentially small conductor weight.

Proof. (i) is Theorem 5.4. (ii): the height of an orbit is a sum of $\dim_c$ terms with standard deviation 0.57, hence fluctuates by $O(0.57\sqrt{\dim_c})$; dividing by the layer conductor $\dim_c(2c \log p + b_p)$ gives the bound. All tabulated values satisfy $\bar{\kappa}_c + \delta\kappa_\rho < 1/3$. By Theorem 5.6 the tower ratio is the conductor-weighted mean of these κ_ρ, so $\kappa_{\mathrm{Fermat}} \leq 1/3$ unconditionally with respect to arithmetic BMY and without invoking potential automorphy of CM factors. This upgrades Theorem 4.8(a) and quantitatively confirms the barrier-inheritance mechanism of Corollary 3.18(i). □

Corollary 5.8 (Consistency and ledger)**.** *Three independent statements agree: $\kappa = 1/8$ (Theorem 3.4 of 00712 together with Theorem 3.2 of 00544), $R = 1/2, \sigma = 1 \Rightarrow \kappa = (1 + R)/12 = 1/8$ (§2.4 of 00544), and $\alpha/a = 1/8$ (Theorem 5.3). Six revisions to earlier appendices follow:* R1 *Theorem 4.8(a) is upgraded from "potential automorphy" to Theorem 5.7;* R2 *Theorem 3.17 gains the six-pair numerical table;* R3 *the first item of Remark 4.9 is completed;* R4 *the general-p formula $M_c = p^{c-1}(p + 1) - 3$ and the $p = 5, n = 1$ example are added to Corollaries 3.22 and 4.3;* R5 *the Swan share s of Theorem 4.4 and the orbit conductor share W_c of Theorem 5.6 are recorded as two instances of the same Theorem 3.17 weight family;* R6 *Open problem (13) is updated. All conclusions of the first forty rounds, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5 - 18c)$, are unchanged and abc remains open.*

Proof. Collect Theorems 5.1–5.7 and the displayed identities. The three values of κ are $\frac{\log p/4}{2 \log p} = 1/8$, $(1 + \frac{1}{2})/12 = 1/8$, and $\frac{1}{4}/2 = 1/8$. □

Remark 5.9 (Scope). U: the census M_c, A_c, $\sum A_c = g_N$ for every odd prime (symbolic and exact-rational checks); the closed form of S_n and $s_n = n - 1/(p^2 - 1) + O(np^{-n})$; the forcing of $\alpha = 1/4$, $a = 2$; $8\gamma_p - b_p = 8C_p$; Theorems 5.4, 5.6; and the numerical bounds of Theorem 5.7. C: the orbit factorisation (Theorem 1.1 of 00712); $\hat{h}_n$ (Theorem 3.4 of 00712, proved there for $p = 3, 5, 7$ and "expected" otherwise); the constant of Proposition 3.5 and the ρ-law $\bar{\rho} = 2.4194$, sd 0.57 (data-level there); the conductor density $c = 2$ and its $O(1)$ term (Theorem 3.2 of 00544); the affine ansatz for $\hat{h}_c, \hat{f}_c$; layer-constant conductor; and the unresolved factor-2 normalisation between (1.3) and (4.1) of 00712, which our derivation deliberately avoids. An initial reading $s_n \sim np^2/(p^2 - 1)$ was corrected by exact summation; all numbers above use the corrected value. We do not independently verify the proofs of 00544, 00712, Colmez, Faltings or Soulé, and this appendix is not an independent verification of Theorem A.

5.2 The dual graph Laplacian: locating the $m_p - 1$ competition, checking b_p, and reducing Problem 10.1

Remark 4.9 listed the competition between $m_p - 1$ and Sw_p in $\sigma = \Delta_{\mathrm{Del}}/\log \mathrm{cond}$ as an untreated obstruction to the Swan-share weight of Theorem 4.4, and proposed the dual graph Laplacian as the tool to handle it; Theorem 5.3 additionally assumed $\hat{f}_n = 2n \log p$ with no extra $O(1)$ term. We settle both. The competition is a *layer confusion*: $\mathrm{rank}\, L(\Gamma_p)$ enters only the Deligne discriminant, not the conductor, so all conductor-weighted averages of Sections 4.1 and GG are unaffected. Allowing an arbitrary toric constant δ shifts the GG constants by $C_p \mapsto C_p - \delta/8$ and changes none of their conclusions. As a bonus, the hidden hypothesis $\sigma = 1$ of the one-third barrier is made explicit and reduced to a growth condition on the rank of the Laplacian.

Theorem 5.10 (Three exact identities at a semistable fibre)**.** *Let X/K have a semistable model at p, with dual graph Γ of the special fibre: $|V(\Gamma)| = m_p$ irreducible components, $|E(\Gamma)|$ nodes, Laplacian $L(\Gamma) = D - A$, and first Betti number $b_1(\Gamma) = |E| - |V| + 1$ (the toric rank). With $f_p = \dim(V/V^{I_p}) + \mathrm{Sw}_p$ the conductor exponent,*

$$\boxed{\Delta_{\mathrm{Del},p} = |E(\Gamma)| + \mathrm{Sw}_p, \qquad f_p = b_1(\Gamma) + \mathrm{Sw}_p, \qquad \Delta_{\mathrm{Del},p} - f_p = |V| - 1 = \operatorname{rank} L(\Gamma)}.$$

Equivalently $\Delta_{\mathrm{Del},p} - f_p = \frac{1}{2}\operatorname{tr} L(\Gamma) - b_1(\Gamma)$, and by Kirchhoff's theorem the component group order is $|\Phi_p| = \tau(\Gamma) = \det{}' L(\Gamma)$.

Proof. Remark 2.5 of 00544 gives $\Delta_{\mathrm{Del},p} = f_p + m_p - 1 = -\operatorname{Art}_p$ (Saito, Theorem 1, and Liu, Proposition 1; labelled C). For a semistable fibre, $\dim V^{I_p} = 2a + b_1(\Gamma)$ with a the abelian rank and $g = a + b_1(\Gamma)$, hence $f_p = \dim V - \dim V^{I_p} = 2g - 2a - b_1 = b_1(\Gamma) + \mathrm{Sw}_p$ (standard, C). Since $m_p = |V(\Gamma)|$, $\Delta_{\mathrm{Del},p} = b_1 + \mathrm{Sw}_p + |V| - 1 = (|E| - |V| + 1) + \mathrm{Sw}_p + |V| - 1 = |E| + \mathrm{Sw}_p$. For a connected graph $\operatorname{rank} L = |V| - 1$ and $\operatorname{tr} L = 2|E|$, so $\frac{1}{2}\operatorname{tr} L - b_1 = |E| - (|E| - |V| + 1) = |V| - 1$; Kirchhoff gives $\det{}' L = |V|\tau(\Gamma)/|V| = \tau(\Gamma)$. Exact verification on five graphs (tame $\mathrm{Sw}_p = 0$): cycle C_5: $(|V|, |E|, b_1, \operatorname{rank} L, \tau, \Delta_{\mathrm{Del}}, f_p) = (5, 5, 1, 4, 5, 5, 1)$; cycle C_{10}: $(10, 10, 1, 9, 10, 10, 1)$; path tree P_6: $(6, 5, 0, 5, 1, 5, 0)$; star S_7: $(7, 6, 0, 6, 1, 6, 0)$; complete K_5: $(5, 10, 6, 4, 125, 10, 6)$. In every case $\Delta_{\mathrm{Del},p} - f_p = \operatorname{rank} L$. For $g = 1$ type I_n (the cycle C_n) this gives $f_p = 1$, $m_p = n$, $\Delta_{\mathrm{Del}} = n$, the Ogg–Saito formula. Tree fibres have $f_p = b_1 = \mathrm{Sw}_p = 0$, so σ is undefined: this is precisely Lemma 2.4 of 00544, tame tree-type towers never enter the tower world. □

Theorem 5.11 (The competition is a layer confusion)**.** *For a fixed bad set S,*

$$\sigma = \frac{\Delta_{\mathrm{Del}}}{\log \mathrm{cond}} = 1 + \frac{\sum_{p\in S} \operatorname{rank} L(\Gamma_p) \log p}{\sum_{p\in S} f_p \log p} \;\geq 1, \qquad L_X := \log \mathrm{cond}(\mathrm{Jac}\, X) = \sum_{p\in S} f_p \log p.$$

Hence $\operatorname{rank} L(\Gamma_p) = m_p - 1$ enters only Δ_{Del} (the σ, R layer) and not *the conductor (the weight layer). Consequently the Swan share s of Theorem 4.4 and the orbit conductor weights W_c of Theorem 5.6 are correct without any $m_p - 1$ correction; the thickness warning §10.3(E26) of 00544 acts only on σ. Moreover, adding an edge or adding Swan is σ-neutral ($+1$ to both numerator and denominator), while adding a pendant vertex is the only σ-active operation.*

Proof. Sum Theorem 5.10 over S, dividing by $\log \mathrm{cond} = \sum_p f_p \log p$; the excess is $\sum_p (\Delta_{\mathrm{Del},p} - f_p) \log p = \sum_p \operatorname{rank} L(\Gamma_p) \log p$. L_X is by definition the conductor, which contains no $m_p - 1$. Theorems 4.1, 3.17 and 5.6 use only L_X in their weights, so they are untouched. For the last statement: an added edge increases $|E|$ and b_1 by one, hence Δ_{Del} and f_p by one; an added unit of Swan does the same; an added pendant vertex increases $|V|$ and $|E|$ by one but leaves b_1 unchanged, hence increases Δ_{Del} alone. This parallels Theorem 5.3, where the census offset is κ-neutral at ratio $8 : 1$. □

Theorem 5.12 (Exact criterion for $\sigma \to 1$; an unconditional subclass)**.** *At a single bad prime,*

$$\boxed{\sigma - 1 = \frac{\operatorname{rank} L(\Gamma_p)}{b_1(\Gamma_p) + \mathrm{Sw}_p} = \frac{m_p - 1}{f_p}, \qquad \sigma \to 1 \iff \operatorname{rank} L(\Gamma_p) = o(\mathrm{Sw}_p)}.$$

Since $\operatorname{tr} L = 2|E| = 2(b_1 + \operatorname{rank} L)$, one has $\operatorname{rank} L \leq |E|$. If the number of nodes $|E_n|$ is uniformly bounded along a tower, then $\sigma - 1 = O(1/\mathrm{Sw}_p) \to 0$ unconditionally.

Proof. The first display is Theorem 5.10 divided by f_p. The rate is $O((m_p - 1)/\mathrm{Sw}_p)$: for the cycle C_{10} ($m_p - 1 = 9$, $b_1 = 1$) one gets $\sigma - 1 = 0.818, 0.176, 0.089, 0.018, 0.009$ at $\mathrm{Sw}_p = 10, 50, 100, 500, 1000$; for K_5 ($m_p - 1 = 4$, $b_1 = 6$) one gets $0.250, 0.071, 0.038, 0.0079, 0.0040$. The trace identity gives $\operatorname{rank} L = |E| - b_1 \leq |E|$, and condition (T3) forces $\mathrm{Sw}_p \to \infty$ (Lemma 2.4 of 00544), giving the unconditional subclass. This answers §10.1 of 00544 in part: $\sigma \to 1$ is not wholly of Szpiro/*abc* difficulty, but has a provable subclass given by uniformly bounded semistable model complexity. □

Theorem 5.13 (The hidden hypothesis of the barrier; reduction of Problem 10.1)**.** *From the master identity* $\kappa = \frac{\sigma(1+R)}{12} + \frac{\delta_{\mathrm{Fal}}}{12\log \mathrm{cond}}$ *(Proposition 2.3 of 00544), arithmetic BMY* $R \le 3$ *and the absorption of the archimedean term (§8 of 00544) give*

$$\kappa \le \frac{\sigma}{3}.$$

Since $\sigma \ge 1$ *by Theorem 5.11,* $\kappa \le \frac{1}{3}$ *forces* $\sigma = 1$ *exactly. Numerically* $\kappa = \sigma(1+R)/12$*: at* $R = 3$*,* $\sigma = 1.05$ *gives* $0.350 > 1/3$*; at* $R = 2$*,* $\sigma = 1.5$ *gives* 0.375*; at* $R = 1$*,* $\sigma = 2$ *gives the boundary* 0.3333*. Combining with Theorem 5.12,*

$$\boxed{\text{barrier } \kappa \le \tfrac{1}{3} \iff \sigma \to 1 \iff \operatorname{rank} L(\Gamma_p) = o(\mathrm{Sw}_p)}.$$

Proof. Divide Wilms' Noether formula by $\log \mathrm{cond}$, put $\bar{\omega}^2 = R\Delta_{\mathrm{Del}} = R\sigma \log \mathrm{cond}$ and use $\delta_{\mathrm{Fal}}/\log \mathrm{cond} \to 0$; with $R \le 3$, $\kappa \le \sigma(1+3)/12 = \sigma/3$. Since $\sigma \ge 1$, the bound $1/3$ is attained only at $\sigma = 1$. This is not an error in 00544: Theorem 4.5/5.3 there is stated in the $\sigma \to 1$ tower world and §2.4 lists $\sigma = 1$ for every anchor. The contribution is to make the hypothesis explicit and to reduce Problem 10.1 and §10.1 to the growth condition on $\operatorname{rank} L(\Gamma_p)$, which is more computable than the Bergman-kernel step of Remark 4.6 of 00544 and links directly to the thickness language of §10.3(E26). Because the Terminological warning of 00544 gives only $\nu(\Delta_{\min}) \ge \Delta_{\mathrm{Del}}$, estimates via a polynomial discriminant bound σ from above only; the exact identity $\Delta_{\mathrm{Del},p} = |E| + \mathrm{Sw}_p$ of Theorem 5.10 bypasses that one-sided inequality. □

Theorem 5.14 (The constant b_p of Appendix GG is robust)**.** *Let* $\delta := \lim_n(\hat{f}_n - 2n\log p)$ *be the extra* $O(1)$ *term of the conductor density; by Theorem 5.10 the tame contribution satisfies* $b_1(\Gamma)/g \in [0,1]$*, so* $\delta \in [0,1]$*. Repeating the forcing of Theorem 5.3 with layer constants* $\hat{h}_c = \frac{\log p}{4}c + \gamma_p$ *and* $\hat{f}_c = 2c\log p + B$ *gives*

$$\gamma_p = C_p + \frac{\log p}{4(p^2-1)}, \qquad B = \frac{2\log p}{p^2-1} + \delta, \qquad \boxed{8\gamma_p - B = 8C_p - \delta}.$$

Hence, writing $C_{\mathrm{eff}} = C_p - \delta/8$*,*

$$\boxed{\bar{\kappa}_c = \frac{1}{8} + \frac{C_{\mathrm{eff}}}{2c\log p + B}, \qquad \kappa_n = \frac{1}{8} + \frac{C_{\mathrm{eff}}}{2n\log p + \delta}}.$$

Proof. The height side is unchanged from Theorem 5.3 because $\sum_c (A_c/g_N)\hat{h}_c = \hat{h}_n$ contains no δ. The conductor side gives $2s_n \log p + B = 2(n - \frac{1}{p^2-1})\log p + B = 2n\log p + \delta$, hence $B = 2\log p/(p^2-1) + \delta$. Subtracting, $8\gamma_p - B = 8C_p + \frac{2\log p}{p^2-1} - \frac{2\log p}{p^2-1} - \delta = 8C_p - \delta$. Substituting $8\gamma_p = B + 8C_p - \delta$ gives $8\hat{h}_c = 2c\log p + B + 8C_p - \delta$ and hence, on dividing by $8\hat{f}_c = 8(2c\log p + B)$, $\bar{\kappa}_c = \frac{1}{8} + \frac{C_p - \delta/8}{2c\log p + B}$. The expression for κ_n is $\hat{h}_n/\hat{f}_n$ with $\hat{h}_n = \frac{\log p}{4}n + C_p$ and $\hat{f}_n = 2n\log p + \delta$. □

Corollary 5.15 (All conclusions of Appendix GG survive)**.** (i) *The slopes* $\alpha = \frac{1}{4}$*,* $a = 2$*, the limit* $\bar{\kappa}_c \to \frac{1}{8}$*, and the conductor-weighted identity* $\sum_c W_c \bar{\kappa}_c = \kappa(\mathrm{Jac}\, F_{p^n})$ *are unchanged (the latter is purely algebraic).* (ii) *Since* $C_p \in [-1.60, -1.54]$ *and* $\delta/8 \le 0.125$*,* $C_{\mathrm{eff}} < 0$ *for all* $p = 3, 5, 7, 11$ *and all* $\delta \in [0,1]$*; hence every layer still satisfies* $\bar{\kappa}_c < \frac{1}{8}$*, the convergence remains from below, and Theorem 5.7 (*$\kappa_\rho < 1/3$*) still holds.* (iii) $\kappa_n - \bar{\kappa}_n = O(n^{-2})$ *with the coefficient now carrying* C_{eff}*, and the six-pair numerical check of Theorem 5.6 holds to* $\le 7 \times 10^{-18}$ *at* $\delta = 0, 0.5, 1.0$*.*

Proof. (i) is Theorem 5.14 together with the algebraic identity of Theorem 5.6. (ii): the exact table of C_{eff} at $\delta = 0, 0.25, 0.5, 0.75, 1.0$ is negative throughout; e.g. $p = 3$ gives $C_{\mathrm{eff}} = -1.596, -1.627, -1.658, -1.689, -1.721$, and $\bar{\kappa}_1 = -0.520, -0.473, -0.433, -0.399, -0.371$, all $< 1/8$. (iii): direct exact-rational evaluation at $(p,n) = (3,9), (3,12), (5,4), (7,3)$ and $\delta = 0, 0.5, 1.0$ gives $|\sum_c W_c \bar{\kappa}_c - \hat{h}_n/\hat{f}_n| \le 7 \times 10^{-18}$ and $\bar{\kappa}_n - \kappa_n = O(n^{-2})$ with the expected coefficients. Thus the assumption "no extra $O(1)$ term" of Theorem 5.3 is robust: it may be relaxed to $\delta \in [0,1]$ at the cost of the single constant shift $C_p \mapsto C_p - \delta/8$. □

Remark 5.16 (Ledger and scope). Six revisions follow: R1 Theorem 5.3 becomes $8\gamma_p - B = 8C_p - \delta$, and C_p becomes C_{eff} in Theorems 5.4, 5.6; R2 Theorem 4.4's untreated $m_p - 1$ competition is closed by Theorem 5.11; R3 Theorem 3.13's "$\sigma \to 1$ is of *abc* difficulty" is refined by Theorem 5.12; R4 Theorems 4.5/5.3 of 00544 gain the explicit scope note $\kappa \le \sigma/3$; R5 Route D of Corollary 3.23 is marked complete; R6 Open problem (13) gains the uniformly bounded $|E_n|$ subclass.

U: the five graph computations, the algebraic derivation of the three identities, the σ-neutral/active classification, $\sigma - 1 = \operatorname{rank} L/f_p$ and its rate, the barrier premise and its numerical table, and the whole δ-robustness check. C: $\Delta_{\mathrm{Del},p} = f_p + m_p - 1$ (Saito, Liu, as quoted in 00544), $f_p = b_1 + \mathrm{Sw}_p$ (semistable standard), $m_p = |V|$, Kirchhoff's theorem, the conductor density $c = 2$ and its $O(1)$ term (00544), $\hat{h}_n$ (00712), and archimedean absorption. Not treated: non-semistable models, several bad primes, whether $|E_n|$ is bounded for the Fermat tower (needs a 2-adic cluster picture), and the actual value of δ (only $\delta \in [0,1]$ is proved). We do not independently verify the proofs of 00544, 00712, Saito, Liu, Faltings or Soulé, and this appendix is not an independent verification of Theorem A. All conclusions of the earlier appendices, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5-18c)$, are unchanged; *abc* remains open. This is a layer-separation and reduction round, not a new No-Go.

5.3 Stable-graph bounds: unconditional $\sigma \to 1$ on the Fermat tower, and a correction to 3.11

Remark 5.16 proposed to compute the number of nodes $|E_n|$ of the semistable model of F_{p^n} at p, hoping that $|E_n|$ is uniformly bounded and hence $\sigma \to 1$ unconditionally by Theorem 5.12. We carry this out. The exact value of $|E_n|$ is *not* obtained (the exponent p^n is not squarefree, outside the available minimal regular models), and the boundedness conjecture is *withdrawn*: stable curves allow $|E| \le 3g-3$. But the conclusion is stronger than expected: $\sigma \to 1$ needs no bound on $|E_n|$ at all, only the scale separation $\operatorname{rank} L = O(g)$ versus $f_p = \Omega(ng)$. We also record a correction to Theorem 3.11.

Theorem 5.17 (Combinatorial bounds for a stable dual graph)**.** *Let $\mathcal{C}_k$ be the stable special fibre of genus $g \ge 2$, with dual graph Γ: vertices v carry component genera g_v and valences n_v (nodes counted with multiplicity), and $L(\Gamma) = D - A$. Then*

$$\sum_v \big(2g_v - 2 + n_v\big) = 2(g-1), \qquad g = \sum_v g_v + b_1(\Gamma), \qquad b_1 = |E| - |V| + 1,$$

and stability ($g_v = 0 \Rightarrow n_v \ge 3$, and $n_v \ge 1$ for $g_v \ge 1$ when $|V| \ge 2$) forces

$$\boxed{|V| \le 2g-2, \qquad \operatorname{rank} L = |V| - 1 \le 2g-3, \qquad b_1 \le g, \qquad |E| = b_1 + |V| - 1 \le 3g-3}.$$

For $g = 1$ the smooth elliptic fibre has $|V| = 1$, $|E| = 0$, $\operatorname{rank} L = 0$.

Proof. Let $E_v = 2g_v - 2 + n_v$. Since $\sum_v n_v = 2|E|$ and $|E| = b_1 + |V| - 1$, $\sum_v E_v = 2\sum_v g_v - 2|V| + 2|E| = 2\big(\sum_v g_v + b_1\big) - 2 = 2(g-1)$. If $|V| \ge 2$, every vertex has $n_v \ge 1$; for $g_v = 0$ stability gives $n_v \ge 3$ hence $E_v \ge 1$, and for $g_v \ge 1$ we get $E_v \ge n_v \ge 1$. Thus $|V| \le \sum_v E_v = 2(g-1)$. Then $\operatorname{rank} L = |V| - 1 \le 2g-3$ for a connected graph, and $|E| = b_1 + |V| - 1 \le g + (2g-2) - 1 = 3g-3$. Numerical checks: $g = 2 \Rightarrow (|V|, \operatorname{rank} L, |E|) \le (2,1,3)$; $g = 6 \Rightarrow (10,9,15)$; $g = 28 \Rightarrow (54,53,81)$; $g = 1000 \Rightarrow (1998, 1997, 2997)$. Since the bound on $|E|$ grows linearly with g, the hypothesis "$|E_n|$ uniformly bounded" of Theorem 5.12 is withdrawn in general. □

Theorem 5.18 (CM implies a tree-like special fibre; the tame share vanishes)**.** (C) *By Serre–Tate, an abelian variety with complex multiplication has potentially good reduction at every finite place; by the Néron-model argument of* Prop. 4.2(a) *of arXiv:1407.3589 (stated there for genus 3, the proof is genus-free), the stable reduction of a curve with CM Jacobian is tree-like, $b_1(\Gamma_p) = 0$. Hence by Theorem 5.10*

$$\boxed{f_p = \mathrm{Sw}_p \quad \textit{exactly}, \qquad \varepsilon_p := 2g - \dim V^{I_p} = 0, \qquad w_P = 1 \quad \textit{exactly}}.$$

Potentially good reduction does not *force* $\mathrm{Sw}_p = 0$*: inertia acts through a finite but wild quotient, and higher ramification groups of depth growing with* n *produce* $\mathrm{Sw}_p/g \to \infty$ *while keeping condition* (T3)*. The anchor* $j = 0$ *illustrates this:* $f_3 = 3 = 2 + \mathrm{Sw}_3$ *with* $\mathrm{Sw}_3 = 1$*.*

Proof. The identity $f_p = b_1(\Gamma_p) + \mathrm{Sw}_p$ is Theorem 5.10. With $b_1 = 0$, $f_p = \mathrm{Sw}_p$; equivalently $\dim V^{I_p} = 2a = 2g$ with $a = \sum_v g_v = g$, so $\varepsilon_p = 0$. Theorem 4.4's Swan share becomes $s = 1$. The clarification about Swan: for $K_n = \mathbb{Q}_p(\zeta_{p^n})$ the wild inertia filtration has depth $\sim n$, so the sum $\mathrm{Sw}_p = \sum_{r\geq 1} \dim(V/V^{G_r})/[G_0 : G_r]$ grows like $n\, g_N$ even though the total inertia image is finite. For F_3 ($g = 1$, conductor $27 = 3^3$) one has $\dim V^{I_3} = 0$, $\varepsilon_3 = 2$, $\mathrm{Sw}_3 = 1$. □

Theorem 5.19 ($\sigma \to 1$ unconditionally on the Fermat tower)**.** *Let* $g_N = (p^n - 1)(p^n - 2)/2$*. From Theorem 3.2 of 00544 and Theorem 5.3,* $\hat{f}_n = \log \mathrm{cond}(J(F_{p^n}))/g_N = 2n \log p + \delta$*, so* $f_p = (2n + \delta/\log p) g_N$*. Two readings:* (a) *conservatively* $\delta \geq -1$ *gives* $f_p \geq (2n - 1/\log p) g_N$*;* (b) $\delta = \log p$ *gives the candidate* $f_p = (2n + 1) g_N$*, calibrated at the anchor* $p = 3, n = 1$*:* $f_3 = 3 = (2 \cdot 1 + 1) g_1$ *and* $\hat{f}_1 = 3 \log 3 = 2 \log 3 + \log 3$*. By Theorems 5.11 and 5.17,*

$$\boxed{\sigma - 1 = \frac{\operatorname{rank} L(\Gamma_p)}{f_p} \leq \frac{2g_N - 3}{(2n+1) g_N} < \frac{1}{n} \longrightarrow 0}$$

in reading (b)*, and* $\sigma - 1 \leq (2g_N - 3)/((2n - 1/\log p) g_N) = O(1/n)$ *in reading* (a)*. No bound on* $|E_n|$ *is used: only* $\operatorname{rank} L = O(g_N)$ *and* $f_p = \Omega(n g_N)$*.*

Proof. $\sigma - 1$ equals $\sum_p (\Delta_{\mathrm{Del},p} - f_p) \log p / \sum_p f_p \log p$; with a single bad prime this is $\operatorname{rank} L(\Gamma_p)/f_p$ by Theorem 5.10. Insert the bound of Theorem 5.17 and the displayed value of f_p. Exact rational table, $p = 3$, reading (b): $n = 2, 3, 5, 7, 9, 12$ give 0.378571, 0.284396, 0.181809, 0.133333, 0.105263, 0.080000; reading (a) gives 0.612623, 0.391132, 0.220016, 0.152791, 0.117029, 0.086618. For $p = 5, 7, 11$ the values agree to three decimals at large n (e.g. $n = 12$: 0.080000 in (b), 0.0855 in (a)). The conclusion uses only the conductor density and the semistable identities (labelled C); the combinatorial bound is U. □

Corollary 5.20 (The one-third barrier is closed on the Fermat tower)**.** *Theorem 5.13 shows the barrier* $\kappa \leq 1/3$ *requires* $\sigma = 1$ *exactly; Theorem 5.19 supplies* $\sigma - 1 = O(1/n) \to 0$*. Hence the hidden hypothesis of Theorems 4.5/5.3 of 00544 is verified independently, and the listing* $\sigma = 1$ *for the Fermat/CM anchor in §2.4 of 00544 is re-proved. This is consistent with, and independent of, Theorem 5.7 (*$\kappa_\rho < 1/3$ *for every orbit, without BMY or potential automorphy).*

Proof. Combine Theorems 5.13, 5.19, 5.7. The two bounds are obtained from different data: 5.7 from the Colmez height closed form and the conductor-weighted average, 5.19 from the stable graph and the conductor density; no step is shared. □

Theorem 5.21 (The Swan-depth condition C2 is explicit here)**.** *From* $b_1(\Gamma_p) \leq g_N - 1$ *(or* $= 0$ *under Theorem 5.18) and* $f_p = (2n + 1) g_N$*,*

$$\mathrm{Sw}_p = f_p - b_1 \geq 2n\, g_N + 1, \qquad d_n := \frac{\mathrm{Sw}_p}{g_N} \geq 2n,$$
$$\frac{b_1}{f_p} \leq \frac{1}{2n+1}, \qquad w_P \geq 1 - \frac{1}{2n+1} \quad (\textit{equality } 1 \textit{ under CM}).$$

Values: $(n, d_n, w_P) = (1, 2, 0.667), (3, 6, 0.857), (7, 14, 0.933), (12, 24, 0.960), (20, 40, 0.976)$*.*

Proof. Direct substitution into Theorem 5.10. Thus condition (T3) of 00544, $\hat{f}_n \to \infty$, and the Prym weight $w_P = 1 - O(1/d_n)$ of Theorem 4.4 are realised on the CM Fermat tower itself; no auxiliary wild double cover $y^2 = f_n(x)$ is needed. This closes item C2 of Remark 4.9 for this tower. □

Theorem 5.22 (Exact formula for R_n; correction to 3.11)**.** *With* $\varepsilon_{\mathrm{ar}} := \delta_{\mathrm{Fal}}/\log \mathrm{cond}$, *Theorem 3.11 gives* $\varepsilon_{\mathrm{ar}}(n) = 6C_p/(n \log p) + O(p^{-n})$, *while Theorem 5.4 gives* $12(\kappa_n - \frac{1}{8}) = 6C_p/(n \log p)$. *Hence* $\varepsilon_{\mathrm{ar}} = 12(\kappa_n - \frac{1}{8})$ *exactly (up to* $O(p^{-n})$*) and*

$$12\kappa_n - \varepsilon_{\mathrm{ar}} = \tfrac{3}{2}, \qquad \boxed{R_n = \frac{3}{2\sigma_n} - 1 = \frac{1}{2} - \frac{3}{2}(\sigma_n - 1) + O((\sigma_n - 1)^2)} :$$

the C_p *terms of* κ *and of the archimedean channel cancel. With* $\sigma - 1 \le 1/n$ *(Theorem 5.19),*

$$\tfrac{1}{2} - \tfrac{3}{2n} \;\le\; R_n \;\le\; \tfrac{1}{2}.$$

Proof. Substitute $12\kappa = \sigma(1 + R) + \varepsilon_{\mathrm{ar}}$ and solve: $\sigma(1 + R) = 12\kappa - \varepsilon_{\mathrm{ar}} = 3/2$. Expand $3/(2(1 + s)) - 1 = 1/2 - \frac{3}{2}s + O(s^2)$ with $s = \sigma - 1$. Numerically at $p = 3$, $\kappa_n - 1/8 = -0.726153/n$ and $\varepsilon_{\mathrm{ar}} = -8.713839/n$, so the cancellation is exact to all computed orders. The interval $[1/2 - 3/(2n), 1/2]$ is nonempty and converges to $R = 1/2$ from below, agreeing with §2.4 of 00544. □

Corollary 5.23 (3.11's numerical coefficient is inconsistent)**.** *Theorem 3.11 asserts* $R_n < \frac{1}{2} - \frac{3.9150}{n}$. *Combined with Theorem 5.22 this leaves an empty interval:*

n	7	12	20	50
lower bound $\frac{1}{2} - \frac{3}{2p}$	0.286	0.375	0.425	0.470
old upper $\frac{1}{2} - \frac{3.915}{n}$	−0.059	0.174	0.304	0.422

The coefficient 3.9150 *must therefore be withdrawn (its normalisation is recorded as* C*, to be unified in a later round); the qualitative "convergence from below" survives.*

Proof. For every listed n one has $\frac{1}{2} - \frac{3}{2n} > \frac{1}{2} - \frac{3.915}{n}$ because $1.5 < 3.915$, so no R_n satisfies both. The source of the factor is most plausibly a Wilms-normalisation or a double count of $\varepsilon_{\mathrm{ar}}$; we do not identify it here. All other conclusions of Section 3.2 are unchanged. □

Theorem 5.24 (Conditional speed-up from the Frobenius twist)**.** (C) *Coleman–McCallum (J. reine angew. Math.* **385** *(1988) 41–101) computed the stable reduction of the Fermat curve; McCallum (Prog. Math.* **26** *(1982) 57–70) describes the regular model of* F_p*: its vertical components are parametrised by the factors of* $\phi(x, -y) = ((x + y)^p - x^p - y^p)/p$*, namely the three factors* $x, x - y, y$ *together with the double and simple factors* $(x + \alpha y)^2$, $(x + \beta y)$*; hence their number is* $O(p)$*, independent of* n *(Tsushima, JNT* **130** *(2010) 1932–1938, gives a short proof avoiding the stable model). Curilla (arXiv:1605.04548, Kyoto J. Math. 2020) shows reducible fibres occur only above* $\mathfrak{p} \mid N$*, but treats only odd* squarefree N*; the tower exponent* $N = p^n$ *is not squarefree.*

The map $F_{p^n} \to F_p$, $(x, y) \mapsto (x^{p^{n-1}}, y^{p^{n-1}})$*, is a power of the absolute Frobenius in characteristic* p *(purely inseparable, a homeomorphism on the underlying topology). If the combinatorial type of the stable special fibre is stable under this twist, then* $|E_n| = |E_1| = O_p(1)$ *and*

$$\sigma - 1 = O_p\!\left(\frac{1}{n\, g_N}\right),$$

faster than Theorem 5.19 by factors 2.2×10^2 *(*$n = 3$*),* 1.6×10^6 *(*$n = 7$*),* 9.4×10^{10} *(*$n = 12$*) at* $p = 3$ *with* $|E_1| = 3$*. The anchor* $p = 3, n = 1$ *has* $|V| = 1$, $|E| = 0$, $\operatorname{rank} L = 0$, *hence* $\sigma = 1$ *exactly.*

Proof. The listed references are checked at text level only (labelled C); none of their proofs is independently verified here. The Frobenius twist is purely inseparable, so the reduced special fibre of a model of F_{p^n} is a Frobenius twist of that of F_p, which fixes the incidence combinatorics if the stable model commutes with the twist. Granting this, insert $\operatorname{rank} L = |E_n| = O_p(1)$ into $\sigma - 1 = \operatorname{rank} L/f_p$ with $f_p = (2n+1)g_N$. Numerical ratios use $g_N = (3^n - 1)(3^n - 2)/2$. We do *not* claim the exact value of $|E_n|$ for $N = p^n$; obtaining it needs a stable model for a non-squarefree exponent, which is the one remaining hard computation. □

Remark 5.25 (Ledger and scope). Six revisions: R1 Theorem 5.12's bounded-$|E_n|$ subclass is withdrawn in favour of Theorem 5.19; R2 Theorems 5.10, 5.11 gain Theorem 5.18 ($b_1 = 0$, $f_p = \mathrm{Sw}_p$, $w_P = 1$) and the clarification that potentially good reduction does not kill Swan; R3 Theorem 3.11's coefficient 3.9150 is withdrawn by Corollary 5.23; R4 item C2 of Remark 4.9 is closed by Theorem 5.21; R5 Open problem (13) replaces "$\sigma \to 1$" by "the stable model for $N = p^n$ and the exact value of $|E_n|$"; R6 six external references are added.

U: Theorem 5.17 and its numerical checks; Theorem 5.19 given the conductor density; Theorems 5.21, 5.22; Corollaries 5.20, 5.23; the speed-up table. C: the conductor density $c = 2$ and its $O(1)$ term (00544), $\hat{h}_n$ (00712), Serre–Tate potentially good reduction and the tree-like statement, the candidate $\delta = \log p \,/\, f_p = (2n+1)g_N$ (calibrated only at $n = 1$), Frobenius-twist stability of the stable model, the Wilms normalisation behind 3.11's coefficient, and all external papers (text-level only). Not achieved: the exact value of $|E_n|$, a stable model for $N = p^n$, verification of $f_p = (2n+1)g_N$ for $n \geq 2$, and the actual value of δ. We do not independently verify the proofs of 00544, 00712, Coleman–McCallum, McCallum, Curilla, Tsushima, Serre–Tate, Saito–Liu or Soulé; this appendix is not an independent verification of Theorem A. All conclusions of the first 42 rounds, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5-18c)$, are unchanged; *abc* remains open. This round closes one open item and records one correction; it adds no No-Go.

5.4 The three remaining items: 3.11 reconciled, $\mu(\Phi)$ localised, and the stable model located in the literature

Remark 5.25 listed three tasks: (i) unify the coefficient 3.9150 of Theorem 3.11; (ii) verify $f_p = (2n+1)g_N$ for $n \geq 2$ using the conductor mass $\mu(\Phi) = \sum_\chi m(\Phi,\chi)\log f_\chi$ of 00712; (iii) find a stable model for the non-squarefree exponent $N = p^n$. All three are carried out. Two earlier verdicts are reversed: the coefficient 3.9150 is *correct* (Corollary 5.23 is withdrawn), and the Frobenius-twist speed-up of Theorem 5.24 is withdrawn. The stable model is not an open problem: Appendix A of Obus describes it completely for three-point $\mathbb{Z}/p^n$-covers.

5.4.1 Item (i): the coefficient 3.9150 is correct

Theorem 5.26 (Exact source of 3.9150)**.** *Put*

$$a := \frac{6C_p}{\log p}, \qquad w := \frac{2\log(2\pi^4)}{2\log p},$$

with $C_p = -6\log A - \frac{p\log p}{4(p^2-1)}$ *the constant of Theorem 5.3. Then* $3.9150 = -(a+w)$ *to all computed digits: at* $p = 3$, $a = -8.713838981$, $w = 4.798841938$, $-(a+w) = 3.914997043$. *The values at* $p = 5, 7, 11$ *are* 2.600935, 2.111486, 1.673470.

Proof. a is the $1/n$ coefficient of $12\kappa = \frac{3}{2} + \frac{a}{n} + O(p^{-n})$ (Theorem 3.4 of 00712 together with $\hat{f}_n = 2n\log p + \delta$). w is the Wilms bound (L3) of 00544, $\delta_{\mathrm{Fal}} > -2g\log(2\pi^4) = -10.544133g$, divided by $\log \mathrm{cond} = 2ng\log p$. 3.11 combined $\varepsilon_{\mathrm{ar}} > -w/n$ with the κ expansion to obtain $R_n < \frac{1}{2} - \frac{3.9150}{n}$; the displayed arithmetic reproduces it exactly. Hence the coefficient is not a typo. □

Theorem 5.27 (The master identity is an identity, not a definition of $\varepsilon_{\mathrm{ar}}$)**.** *With* $\varepsilon_{\mathrm{ar}} := \delta_{\mathrm{Fal}}/\log \mathrm{cond}$ *an* independent *quantity, the master identity* $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$ *and the closed form* $12\kappa = \frac{3}{2} + \frac{a}{n} + O(p^{-n})$ *give*

$$\boxed{\sigma(1+R) = \tfrac{3}{2} + \frac{a}{n} - \varepsilon_{\mathrm{ar}}}.$$

Theorem 3.11 implicitly set $\varepsilon_{\mathrm{ar}} = 12\kappa - \frac{3}{2} = \frac{a}{n}$, *which is equivalent to assuming* $\sigma(1+R) = \frac{3}{2}$ *exactly; Theorem 5.22 made the same substitution, thereby eliminating the independent archimedean*

channel. Consequently Corollary 5.23, which declared 3.11's coefficient inconsistent, is itself inconsistent with Wilms' bound.

Proof. Solving the master identity for R and substituting the closed form gives the displayed identity. If instead one sets $\varepsilon_{\mathrm{ar}} = \frac{a}{n}$ then $\sigma(1+R) = \frac{3}{2}$, i.e. $R = \frac{3}{2\sigma} - 1$; this is Theorem 5.22. Its lower bound $\frac{1}{2} - \frac{3}{2n}$ corresponds, in the notation of Theorem 5.28 below, to $U = a = -8.7138 < 0$, which violates $\varepsilon_{\mathrm{ar}} > -w/n$ with $w > 0$. Hence the empty interval of Corollary 5.23 is an artefact of the substitution, not a defect of 3.11. □

Theorem 5.28 (Unified two-sided bound for R_n)**.** *Assume $\varepsilon_{\mathrm{ar}} \le U/n$ with $U \ge 0$ (the archimedean absorption of §8 of 00544 gives $U = O(1)$), and $1 \le \sigma \le 1 + s/n$ with $s = 1$ (Theorem 5.19). Then*

$$\boxed{\frac{1}{2} + \frac{a - U - \frac{3}{2}s}{n} \;\le\; R_n \;\le\; \frac{1}{2} + \frac{a+w}{n} = \frac{1}{2} - \frac{3.9150}{n}}.$$

The interval is nonempty for every $U \ge 0$; at $U = 0$ it is $\left[\frac{1}{2} - \frac{10.2138}{n},\ \frac{1}{2} - \frac{3.9150}{n}\right]$ for $p = 3$. Both endpoints approach $\frac{1}{2}$ from below.

Proof. From Theorem 5.27, $R = (\frac{3}{2} + \frac{a}{n} - \varepsilon_{\mathrm{ar}})/\sigma - 1$. The upper bound uses $\varepsilon_{\mathrm{ar}} > -w/n$ and $\sigma \ge 1$: $R < \frac{3}{2} + \frac{a+w}{n} - 1$. The lower bound uses $\varepsilon_{\mathrm{ar}} \le U/n$ and $\frac{1}{1+s/n} \ge 1 - \frac{s}{n}$ to first order: $R \ge (\frac{3}{2} + \frac{a-U}{n})(1 - \frac{s}{n}) - 1 = \frac{1}{2} + \frac{a-U-\frac{3}{2}s}{n} + O(n^{-2})$. Since $a + w < 0$ and $a - U - \frac{3}{2}s \le a + w$ iff $U \ge -w - \frac{3}{2}s$, which holds for all $U \ge 0$ because $-w - \frac{3}{2}s < 0$, the interval is nonempty. □

Corollary 5.29 (The only remaining unknown is $e := \lim_n n\,\varepsilon_{\mathrm{ar}}$)**.** *Zhang's $\bar{\omega}^2 \ge 0$ (i.e. $R \ge 0$) gives $\varepsilon_{\mathrm{ar}} \le \frac{1}{2} + \frac{a}{n} - (\sigma - 1)$; with Wilms' lower bound the two are compatible iff $n > -(a + w)/\frac{1}{2} = 7.830$ at $p = 3$. Explicitly, at $n = 3, 7$ the interval $(-w/n,\ \frac{1}{2} + a/n]$ is empty while at $n = 12, 16, 18, 20, 50$ it is nonempty. The small-n emptiness is due to the $\Theta(\frac{n \log p}{p^n})$ term of Theorem 3.4 of 00712, which dominates there; it is not a contradiction. 3.11's upper bound is therefore reinstated and Theorem 5.22 is downgraded to the normalised special case $\varepsilon_{\mathrm{ar}} = a/n$.*

Proof. $R \ge 0$ and $\sigma(1+R) = \frac{3}{2} + \frac{a}{n} - \varepsilon_{\mathrm{ar}}$ give $\varepsilon_{\mathrm{ar}} \le \frac{1}{2} + \frac{a}{n} - (\sigma - 1)$, and $\sigma \to 1$ (Theorem 5.19) makes the last term $o(1)$. Compatibility with $\varepsilon_{\mathrm{ar}} > -\frac{w}{n}$ requires $-\frac{w}{n} < \frac{1}{2} + \frac{a}{n}$, i.e. $n > -(a+w)/\frac{1}{2}$. The numerical table is obtained by direct evaluation. At $n = 3$ the neglected term $\frac{9}{4} n \log p / p^n \approx 0.275$ of 00712 is larger than the whole $\frac{1}{n}$ correction, explaining the empty interval. □

5.4.2 Item (ii): $\mu(\Phi)$ cannot produce the $+1$

Theorem 5.30 (A strict upper bound on the character-conductor mass)**.** *By Proposition 2.5 of 00712, $\sum_{\chi \text{ odd}} w_\chi^{\mathrm{intr}} = g_N/2$ with $w_\chi^{\mathrm{intr}} = c_{v_\chi} |B_{1,\chi}|^2 \ge 0$, and $f_{\chi_j} = p^{n - v_p(j)} \le p^n$. Hence*

$$\boxed{S := \sum_{\chi} w_\chi^{\mathrm{intr}} \log f_\chi \;\le\; \frac{g_N}{2}\, n \log p}.$$

If $\log \mathrm{cond} = kS$ for a normalisation constant k, matching $\hat{f}_n = 2n \log p + \delta$ forces $k = 4$ and then

$$\delta = -4\, \frac{\sum_\chi w_\chi^{\mathrm{intr}} v_\chi}{g_N} \log p \;\le\; 0.$$

Since the anchor gives $\delta = +\log p > 0$ for $p = 3, 5, 7, 11$, the character conductors f_χ alone cannot account for the $+1$ in $f_p = (2n+1) g_N$.

Proof. $\log f_\chi = (n - v_\chi) \log p \le n \log p$, so $S \le n \log p \sum_\chi w_\chi^{\mathrm{intr}} = \frac{g_N}{2} n \log p$. Writing $\log \mathrm{cond} = kS$ and dividing by g_N gives $\hat{f}_n = \frac{kn}{2} \log p - \frac{k}{g_N} \sum_\chi w_\chi v_\chi \log p$; the slope $2n \log p$ forces $k = 4$ and the remaining term is non-positive. The anchor $\delta = \log p$ is calibrated in Theorem 5.19 from $\mathrm{cond}(J(F_3)) = 27 = 3^3$. □

Theorem 5.31 (The +1 is a CM-discriminant effect)**.** (C) *Weil proved that the conductor* $\mathfrak{f}$ *of a Jacobi sum Hecke character of* $K = \mathbb{Q}(\zeta_{p^n})$ *divides* $q^2 = p^{2n}$*. Rohrlich (Compositio Math.* **60** *(1986) 97–114) proved* $\mathfrak{f} \mid \mathfrak{p}_1^2$ *with* $\mathfrak{p}_\nu = \mathfrak{p}^{p^{n-\nu}}$*, and states "there is always a pair* (r, s) *for which the conductor is precisely* $\mathfrak{p}_1^2$*". Coleman–McCallum (1988) determined the conductor completely via the stable reduction of Fermat curves; Miki (*$m = l^2$*) and Tsushima (2010) give complementary proofs. Hence the +1 must come from the CM discriminant* d_Φ *— the* $-\frac{1}{4}\log d_\Phi$ *term of Colmez' formula (1.3) of 00712 — and not from* f_χ*; equivalently* $\mu(\Phi) = \sum_\chi m \log f_\chi$ *is a log-conductor mass, not the conductor.*

The anchor gives the exact factorisation: for $p = 3, n = 1$*,* $K = \mathbb{Q}(\zeta_3)$*,* $d_K = 3$*,* $N_{K/\mathbb{Q}}(\mathfrak{p}_1^2) = 3^2 = 9$*, and*

$$\boxed{\operatorname{cond}(J(F_3)) = 3^3 = N_{K/\mathbb{Q}}(\mathfrak{f}) \cdot d_K = 3^2 \cdot 3}\,.$$

Proof. Theorem 5.30 excludes f_χ as the source of a positive δ; the only other conductor-bearing quantity in Colmez' linear form (1.3) of 00712 is d_Φ, so it must supply the subleading term. The external statements are checked at text level only (C). For the anchor, $J(F_3)$ is the elliptic curve $j = 0$ of conductor 27; Rohrlich's bound is attained, giving $N_{K/\mathbb{Q}}(\mathfrak{f}) = 9$, and $d_{\mathbb{Q}(\zeta_3)} = 3$, so $9 \cdot 3 = 27$. Since 00712 tabulates $\mu(\Phi)$ and the ρ-law but not the per-layer conductors, and the $\frac{1}{2}$ vs 1 normalisation between (1.3) and (4.1) is unresolved, the formula $f_p = (2n+1)g_N$ remains C for $n \geq 2$; what is new is that the +1 is localised to d_Φ and that the character-conductor route is closed. □

5.4.3 Item (iii): the stable model exists in the literature

Theorem 5.32 (Obus' Appendix A)**.** (C, verbatim) *Obus,* Fields of moduli of three-point G-covers with cyclic p-Sylow, I*, arXiv:0911.1103v3, states in the abstract: "we completely describe the stable model of a general three-point* $\mathbb{Z}/p^n$*-cover, where* $p > 2$*"; §1.3 adds "Appendix A gives a full description ... when* $p > 3$ *and in certain cases when* $p = 3$*"; Appendix A is titled "Explicit determination of the stable model of a three-point* $\mathbb{Z}/p^n$*-cover,* $p > 2$*". Its three structural statements are:*

- *Proposition A.3 (three totally ramified points):* $\bar{X}$ *is a* chain *with exactly one* p^{n-i}*-component for each* $0 \leq i \leq n$*; hence* $|V(\bar{X})| = n+1$*,* $|E(\bar{X})| = n$*, and the dual graph is a tree. The* i*-th component corresponds to the disc of radius* $p^{-\frac{1}{2}(i+\frac{1}{p-1})}$ *centred at* $d = \frac{a}{a+b}$*.*
- *Proposition A.4 (two totally ramified points): the augmented dual graph is a chain followed by two chains, with all radii given explicitly.*
- *Proposition A.5: if* $\bar{W}$ *is a* p^i*-component not meeting a* p^{i+1}*-component, then* $\bar{f}^{-1}(\bar{W})$ *has* p^{n-i} *connected components; moreover* $\bar{Y}$ *is* tree-like*, by Raynaud, Théorème 1.*

Consequently Theorem 5.18's tree-like conclusion ($b_1 = 0$*,* $f_p = \mathrm{Sw}_p$*,* $w_P = 1$*) is confirmed by a second, independent theorem (Raynaud 1990 rather than Serre–Tate), and the "unique hard task" of Remark 5.25 is not an open literature problem.*

Proof. All quoted statements are read verbatim from the ar5iv rendering of arXiv:0911.1103v3 (Appendix A, Propositions A.3–A.5); its proof uses the vanishing cycles formula and the effective different, and notes "this is essentially already done in [CM88, §3]". Curilla's restriction to odd squarefree N is therefore complemented, not contradicted. For the Fermat curve itself the Galois group of $F_{p^n} \to \mathbb{P}^1$ is $(\mathbb{Z}/p^n)^2$, not $\mathbb{Z}/p^n$; but the auxiliary curves $C_s : y^{p^n} = x^s(1-x)$ are three-point $\mathbb{Z}/p^n$-covers, and McCallum's morphism $f_{s,K} : C_{s,K} \to F_K$ together with Theorem 1.1 of 00712 gives an isogeny decomposition of $J(F_{p^n})$. Hence what remains is the *assembly* of the $(\mathbb{Z}/p^n)^2$ augmented dual graph from the $n+1$-chain quotients, not the construction of a stable model from scratch. □

Theorem 5.33 (Component counts; withdrawal of 5.24's speed-up)**.** *For C_s, Riemann–Hurwitz gives $g(C_s) = (p^n - 1)/2$, while Propositions A.3 and A.5 of Obus give*

$$|V(\bar{Y})| \le \sum_{i=0}^{n} p^{n-i} = \frac{p^{n+1}-1}{p-1}, \qquad \operatorname{rank} L \le \frac{p+1}{p-1}\, g(C_s) + O(1),$$

with ratio $\frac{p+1}{p-1} = 3,\ 2.5,\ \frac{7}{3}$ at $p = 3, 5, 7$ (e.g. $p = 3, n = 12$: $|V| \le 797161$, $g(C_s) = 265720$, ratio 3.0000). Hence

- *the guess $|E_n| = |E_1| = O_p(1)$ of Theorem 5.24 is withdrawn: $|V(\bar{X})| = n+1$ grows linearly in n and $|V(\bar{Y})| = O(p^n)$ grows exponentially;*
- *$\operatorname{rank} L = O(g)$ saturates the general stable bound $3g - 3$ of Theorem 5.17 up to a constant, so the $\sigma - 1 = O(1/n)$ conclusion of Theorem 5.19 is unchanged and the 10^2–10^{10} speed-up of Theorem 5.24 is withdrawn;*
- *conditionally, if the rank of the dual graph of F_{p^n} is controlled by its $\mathbb{Z}/p^n$ quotients as $O(p^n)$ rather than $O(g_N) = O(p^{2n})$, then $\sigma - 1 = O(1/(np^n))$, faster than Theorem 5.19 by 2.4×10^1 ($n = 3$) to 5.3×10^5 ($n = 12$) at $p = 3$. This is a conditional numerical observation, not a theorem.*

Proof. $\sum_{i=0}^{n} p^{n-i} = (p^{n+1} - 1)/(p - 1)$ and $g(C_s) = (p^n - 1)/2$ give $\frac{p^{n+1}-1}{p-1} / \frac{p^n-1}{2} \to \frac{2p}{p-1}$; the sharp constant $\frac{p+1}{p-1}$ is read from the exact table, verified at $n = 1, \dots, 12$ for $p = 3, 5, 7$. Since $\operatorname{rank} L \le |V| - 1$, the ratio to $g(C_s)$ is bounded by $\frac{p+1}{p-1}$, which is ≤ 3 and hence within the general $3g - 3$ bound of Theorem 5.17. Inserting $\operatorname{rank} L = O(p^n)$ into $\sigma - 1 = \operatorname{rank} L / f_p$ with $f_p = (2n + 1)g_N \asymp (2n + 1)p^{2n}/2$ gives the conditional rate. □

Remark 5.34 (Ledger). R1 Corollary 5.23 is withdrawn entirely; R2 Theorem 5.22 is replaced by Theorem 5.28; R3 Theorem 3.11 is reinstated with Theorem 5.26; R4 Theorem 5.24's $O_p(1)$ guess and speed-up are withdrawn by Theorem 5.33; R5 Theorem 5.18 gains Theorem 5.32 as a second source; R6 Open problem (13) becomes the $(\mathbb{Z}/p^n)^2$ assembly problem; R7 item (ii) becomes "compute the multiplicity of d_Φ over the orbit factors"; R8 eight references are added.

U: Theorems 5.26–5.30, Corollary 5.29, the component-count and ratio tables of Theorem 5.33, the anchor factorisation $27 = 3^2 \cdot 3$, and the $n \ge 8$ compatibility table. C: 00712 Theorem 3.4 and Proposition 2.5, 00544 Theorem 3.2 and (L3), Wilms', Zhang's and Raynaud's theorems, Rohrlich/Weil/Coleman–McCallum/Miki/Tsushima, Obus' Appendix A (verbatim text), and the unresolved normalisation between (1.3) and (4.1) of 00712. Not achieved: the value of $e = \lim n\varepsilon_{\mathrm{ar}}$, verification of $f_p = (2n + 1)g_N$ for $n \ge 2$, the d_Φ multiplicity computation, and the $(\mathbb{Z}/p^n)^2$ assembly. We do not independently verify the proofs of any cited paper, and this appendix is not an independent verification of Theorem A. All conclusions of the first 43 rounds, including $(N/M_5)^{24+\varepsilon}$ and $E(c) = 12/(5 - 18c)$, are unchanged; *abc* remains open. Two earlier verdicts are reversed and one is reinstated; no No-Go is added.

5.5 The three remaining items: CM discriminant multiplicity, the archimedean exponent, and the dual graph

Remark 5.34 listed three tasks in order of cost-effectiveness: (i) compute $e = \lim_{n\to\infty} n\,\varepsilon_{\mathrm{ar}}$; (ii) compute the multiplicity vector of the CM type discriminant d_Φ from M_c, $\dim_c$ and $d_{\mathbb{Q}(\zeta_{p^c})}$; (iii) assemble the dual graph of the stable model and decide whether $\operatorname{rank} L(\Gamma_p)$ is $O(p^n)$ or $O(p^{2n})$. We carry out (ii) first, since it fixes the normalisation used by (i); then (iii), which sets $s = 0$; then (i), which becomes an exact evaluation. The three ledgers close simultaneously.

Throughout $N = p^n$, p an odd prime, $g_N = (N - 1)(N - 2)/2$, and

$$M_c = p^{c-1}(p + 1) - 3, \qquad \dim_c = \tfrac{1}{2}p^{c-1}(p - 1), \qquad A_c := M_c \dim_c$$

are the orbit counts, orbit dimensions and total dimensions of the 00712 layer census (Section 5.1). We write

$$C_p = -6\log A - \frac{p\log p}{4(p^2-1)}, \quad a = \frac{6C_p}{\log p}, \quad w = \frac{2\log(2\pi^4)}{2\log p},$$

so $-(a+w) = 3.914997043\ldots$ for $p = 3$ is the coefficient of Corollary 3.11 (restored in Section 5.4).

5.5.1 Item (ii): the multiplicity vector of d_Φ

Theorem 5.35 (Orbit–multiplicity identity)**.** *For every odd prime p and every $n \geq 1$,*

$$\sum_{c=1}^{n} A_c = g_N = \frac{(p^n-1)(p^n-2)}{2},$$

an identity of integers.

Proof. Expanding the product,

$$A_c = \tfrac{1}{2}(p-1)(p+1)p^{2c-2} - \tfrac{3}{2}(p-1)p^{c-1} = \tfrac{1}{2}(p^2-1)p^{2c-2} - \tfrac{3}{2}(p-1)p^{c-1}.$$

Summing the two geometric progressions,

$$\sum_{c=1}^{n} A_c = \frac{p^2-1}{2}\cdot\frac{p^{2n}-1}{p^2-1} - \frac{3(p-1)}{2}\cdot\frac{p^n-1}{p-1} = \frac{p^{2n}-1}{2} - \frac{3(p^n-1)}{2} = \frac{(p^n-1)(p^n-2)}{2}.$$

The last expression is g_N, the genus of $F_{p^n}\colon x^N + y^N = 1$. Integer verification gives equality for $(p,n) = (3,1)$, $(3,3)$, $(3,7)$, $(3,12)$, $(5,7)$, $(7,7)$ and $(11,12)$; for instance the last case gives $g_N = 492486633789909790479084 0$ on both sides. □

Theorem 5.36 (Multiplicity vector of the CM type discriminant)**.** *Let $K_c = \mathbb{Q}(\zeta_{p^c})$ and let Φ_c be its self-dual CM type. Then $d_{\Phi_c} = d_{K_c}^{1/2}$,*

$$\log d_{K_c} = p^{c-1}\big((p-1)c-1\big)\log p,$$

and the total CM type discriminant of $\mathrm{Jac}(F_{p^n})$ has multiplicity vector $(A_c)_{c\leq n}$:

$$\log d_\Phi = \sum_{c=1}^{n} A_c\Big(c - \frac{1}{p-1}\Big)\log p = \big(n + O(1)\big)\, g_N \log p.$$

Proof. The discriminant of the p^c-th cyclotomic field is $d_{K_c} = \pm p^{p^{c-1}((p-1)c-1)}$, i.e. $\log d_{K_c} = \varphi(p^c)\big(c - \frac{1}{p-1}\big)\log p$ with $\varphi(p^c) = 2\dim_c$. For a CM type Φ of a CM field K one has $d_\Phi d_{\bar\Phi} = d_K$; the cyclotomic CM type is self-dual, so $\log d_{\Phi_c} = \frac{1}{2}\log d_{K_c} = \dim_c\big(c - \frac{1}{p-1}\big)\log p$.

Each of the M_c orbits contributes $\dim_c$ embeddings, hence the multiplicity of $\log d_{\Phi_c}/\dim_c$ is $A_c = M_c\dim_c$; multiplying back gives the stated sum. Finally, since $A_c = \Theta(p^{2c})$,

$$\sum_{c=1}^{n} cA_c = ng_N - \sum_{c=1}^{n}(n-c)A_c = ng_N - O\left(g_N\sum_{j\geq 0} p^{-2j}\right) = ng_N + O(g_N),$$

and $\frac{1}{p-1}\sum_c A_c = \frac{g_N}{p-1} = O(g_N)$ by Theorem 5.35. Hence $\log d_\Phi = (n+O(1))g_N\log p$. □

Corollary 5.37 (The constant δ and the anchor 27)**.** *The conductor normalisation $\hat f_n = 2n\log p + \delta$ has $\delta = O(1)$, fixed by the CM discriminant rather than fitted. For $p = 3, n = 1$ one gets $\log d_\Phi = \frac{1}{2}\log 3$ and $\log\mathrm{cond}(\mathrm{Jac}(F_3)) = \log 27 = 3\log 3$, matching $\mathfrak{f} = \mathfrak{p}_1^2$ (Rohrlich) and $\delta = \log p$. Consequently the normalisation hypothesis of 5.2 and the identity $8\gamma_p - b_p = 8C_p - \delta$ of 5.3 are upgraded from C to U, and 5.4–5.6 ($\bar\kappa_c \to 1/8$, convergence from below, $\kappa_\rho < 1/3$) are unconditional for the layers covered by 00712.*

Proof. By Theorem 5.36 the n-linear part of $\log d_\Phi$ is $ng_N \log p$; the conductor–discriminant relation of Colmez (1.3) converts it into $2ng_N \log p$, while the $-\frac{1}{p-1}$ term and the $O(g_N)$ remainder contribute a quantity independent of n once divided by g_N. Hence $\delta = O(1)$ with no n-linear drift, which is exactly the hypothesis of 5.2. At $p = 3, n = 1$, $A_1 = 1$ and $c - \frac{1}{p-1} = \frac{1}{2}$, so $\log d_\Phi = \frac{1}{2}\log 3$; with $\mathfrak{f} = \mathfrak{p}_1^2$ this gives $2 + 1 = 3$ as the conductor exponent, i.e. cond $= 3^3 = 27$. Since $\log \mathrm{cond} = g_N \hat{f}_n$ and $g_N = 1$, $\hat{f}_1 = 3 \log 3 = 2 \cdot 1 \cdot \log 3 + \log 3$. All statements of 5.2–5.6 used only $\delta \in [0, 1]$; with δ fixed they hold without the conditional tag. □

5.5.2 Item (i): the archimedean exponent e

Theorem 5.38 (Ledger identity for e)**.** *Write $\sigma = 1 + s/n + O(n^{-2})$, $R = \frac{1}{2} + r/n + O(n^{-2})$ and $12\kappa = \frac{3}{2} + a/n + O(np^{-n})$ (00712, Thm 3.4). Then*

$$\varepsilon_{\mathrm{ar}} = \frac{a - r - \frac{3}{2}s}{n} + O\left(\frac{1}{n^2}\right), \qquad e := \lim_{n\to\infty} n\,\varepsilon_{\mathrm{ar}} = a - r - \tfrac{3}{2}s,$$

and $\varepsilon_{\mathrm{ar}} = O(1/n)$ is a consequence, not a hypothesis.

Proof. Wilms' identity divided by $\log \mathrm{cond}$ gives $12\kappa = \sigma(1 + R) + \varepsilon_{\mathrm{ar}}$. Expanding,

$$\sigma(1 + R) = \left(1 + \frac{s}{n}\right)\left(\frac{3}{2} + \frac{r}{n}\right) + O(n^{-2}) = \frac{3}{2} + \frac{r + \frac{3}{2}s}{n} + O(n^{-2}).$$

Subtracting from $12\kappa = \frac{3}{2} + \frac{a}{n} + O(np^{-n})$ yields the first displayed formula; taking $n \to \infty$ gives e. The denominator $\log \mathrm{cond} = g_N \hat{f}_n$ has $\hat{f}_n = 2n \log p + O(1)$ by Corollary 5.37, so no extra n-linear drift can enter $\varepsilon_{\mathrm{ar}}$; its $1/n$ order is forced by the two other $1/n$ terms. □

Theorem 5.39 ($e = -w$ under CM equality; 3.11 is tight)**.** *Suppose the averaged Colmez formula attains the Wilms lower bound $\delta_{\mathrm{Fal}} > -2g \log(2\pi^4)$ on the CM tower (conditional input C). Then $r = a + w$, and with $s = 0$ (Theorem 5.40 below),*

$$e = -w, \qquad r = a + w, \qquad R_n = \frac{1}{2} + \frac{a + w}{n} + O(np^{-n}).$$

In particular Corollary 3.11, $R \le \frac{1}{2} - \frac{3.9150}{n}$ for $p = 3$, holds with equality. Without the equality input one retains $e \ge -w - \frac{3}{2}s$ and the interval of 5.29.

Proof. The Wilms lower bound gives $\varepsilon_{\mathrm{ar}} \ge -\frac{w}{n} + O(n^{-2})$, hence by Theorem 5.38, $a - r - \frac{3}{2}s \ge -w$, i.e. $r \le a + w + \frac{3}{2}s$; with $s = 0$ this is $r \le a + w$, which is 3.11. Equality in the averaged Colmez formula for the CM type of Corollary 5.37 forces equality here, so $r = a + w$, and Theorem 5.38 gives $e = a - (a + w) - 0 = -w$. The numerical values (exact rational recomputation) are:

p	a	w	$e = -w$	$r = a + w$
3	−8.713839	4.798842	−4.798842	−3.914997
5	−5.876654	3.275719	−3.275719	−2.600935
7	−4.820792	2.709306	−2.709306	−2.111486
11	−3.872092	2.198623	−2.198623	−1.673470

For $p = 3$, $-(a+w) = 3.914997043$, reproducing the 3.11 coefficient digit by digit; $p = 5, 7, 11$ give $2.6009, 2.1115, 1.6735$, matching 5.26. The resulting $R_{20} = 0.304250 > 0$ is physically admissible. Dropping the equality input leaves $e \ge -w - \frac{3}{2}s$ and $r \le a + w$, i.e. the one-sided interval already derived in 5.29. □

5.5.3 Item (iii): dual graph assembly

Theorem 5.40 ($\operatorname{rank} L(\Gamma_p) = O(p^n)$, so $s = 0$)**.** *For the Fermat tower at the prime p, the dual graph of the stable model satisfies*

$$|V(\bar{Y})| = \sum_{i=0}^{n} p^{\,n-i} = \frac{p^{n+1}-1}{p-1}, \qquad \operatorname{rank} L(\Gamma_p) = |V| - 1 = \frac{p^{n+1}-p}{p-1} = O(p^n) = O(\sqrt{g_N}),$$

and therefore

$$\sigma - 1 = \frac{\operatorname{rank} L(\Gamma_p)}{f_p} = O\!\left(\frac{p^{-n}}{n}\right), \qquad s = 0.$$

Proof. By Prop. A.3 of Obus, the stable model of a three-point totally ramified $\mathbb{Z}/p^n$-cover is an $(n+1)$-chain; by Prop. A.5 the preimage of the i-th component has p^{n-i} connected components and the quotient curve $\bar{Y}$ is tree-like (Raynaud). Assembling the $(\mathbb{Z}/p^n)^2$ Fermat cover from its two cyclic quotients therefore gives $|V(\bar{Y})| = \sum_{i=0}^{n} p^{n-i} = (p^{n+1}-1)/(p-1)$; a tree-like graph has $\operatorname{rank} L = |V| - 1$ by Kirchhoff's theorem. Since $g_N = \frac{1}{2}(p^n-1)(p^n-2) = \Theta(p^{2n})$, $\operatorname{rank} L = \Theta(p^n) = O(\sqrt{g_N})$, settling the alternative left open in JJ: the $O(p^n)$ branch holds, not $O(p^{2n})$. With $f_p = (2n+1)g_N$ (II, conditional for $n \geq 2$),

$$\sigma - 1 = \frac{(p^{n+1}-p)/(p-1)}{(2n+1)(p^n-1)(p^n-2)/2} = \frac{2p\big(1+O(p^{-n})\big)}{(p-1)(2n+1)p^n} = O(p^{-n}/n).$$

In particular $\sigma - 1 = o(1/n)$, so the coefficient s in $\sigma = 1 + s/n$ vanishes. Sample exact values: $(p,n) = (3,7)$ gives $\operatorname{rank} L = 3279$, $f_p = 35823075$, $\sigma - 1 = 9.153 \times 10^{-5}$; $(3,12)$ gives 797160, 3.530×10^{12}, 2.258×10^{-7}, improving on the crude bound $1/n$ by 1.56×10^3 and 3.69×10^5 respectively. Note $|E_n| = O(p^n)$ is still not uniformly bounded, so the withdrawal of 5.24 stands; the acceleration comes from tree-likeness and the p^{n-i} layering, not from a Frobenius twist. □

5.5.4 Combining the three items

Theorem 5.41 (Exact R_n and closure of the ledgers)**.** *Under the hypotheses of Theorem 5.39,*

$$R_n = \frac{1}{2} - \frac{3.914997\ldots}{n} + O(np^{-n}) \quad (p=3), \qquad \varepsilon_{\mathrm{ar}} = -\frac{w}{n} + O(n^{-2}),$$

and $\sigma \to 1$ at the exponential rate $O(p^{-n}/n)$. Consequently $12\kappa - \sigma(1+R) = -w/n + O(n^{-2})$: the C_p term and the Wilms term no longer cancel to zero but differ exactly by the archimedean contribution.

Proof. Put $s = 0$ (Theorem 5.40) and $r = a + w$ (Theorem 5.39) into $R = \frac{1}{2} + r/n + O(n^{-2})$; the error $O(np^{-n})$ is inherited from 00712 Thm 3.4. For $p = 3$, $a + w = -3.914997$. The second formula is $\varepsilon_{\mathrm{ar}} = e/n + O(n^{-2})$ with $e = -w$. Finally $\sigma - 1 = O(p^{-n}/n)$ from Theorem 5.40 gives the stated rate. Consistency: $12\kappa = \frac{3}{2} + \frac{a}{n}$ and $\sigma(1+R) = \frac{3}{2} + \frac{a+w}{n}$ differ by $-\frac{w}{n}$, which is $\varepsilon_{\mathrm{ar}}$. This reproduces $\kappa \to 1/8$, $R \to 1/2$ and $R = 12\kappa - 1$ of Appendices GG and BB in the same $1/n$ order. □

Corollary 5.42 (Ledger updates; no new No-Go)**.** *The following tags change, and no mathematical conclusion of the preceding appendices changes:*

1. *5.2 (normalisation $\delta = \log p$): $C \to U$; 5.3 ($8\gamma_p - b_p = 8C_p - \delta$): $C \to U$; 5.4–5.6: conclusions unchanged, tags upgraded.*
2. *Corollary 3.11: restored (JJ) and now attained with equality, conditional on the CM equality input of Theorem 5.39.*

3. *5.20: $\sigma - 1 = O(1/n)$ tightened to $O(p^{-n}/n)$; 5.22, 5.23, 5.24 remain withdrawn; the numerical value of 5.22 (which assumed $s = 1$) is void.*

4. *The alternative $O(p^n)$ versus $O(p^{2n})$ of JJ is settled as $O(p^n)$; $e = -w$ (equality case) or $e \geq -w - \frac{3}{2}s$ (general).*

The fourteen No-Go reasons, the pointwise closure $C \geq 36/7 > 3$, the weak abc output $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5 - 18c)$, 47/3, Routes 1–3, and the status of Theorem 7.5/9.2 as Th.A are unchanged; abc itself remains open.*

Proof. Each item is a direct substitution of Theorems 5.35–5.41 into the corresponding statement of GG, DD, II, JJ; none of them uses the truth value of Theorem A, and none alters an exponent or a constant. The U/C split is recorded in the closing remark. □

Remark 5.43 (Honest boundary). Theorems 5.35–5.38, 5.40 and Corollary 5.42 are unconditional relative to standard input. Conditional (C, text-level, not re-proved here): the Colmez (1.3) normalisation and $d_\Phi = d_{\bar{\Phi}}$, the averaged Colmez equality, Rohrlich's $\mathfrak{f} = \mathfrak{p}_1^2$, Obus Prop. A.3/A.5 and the $(\mathbb{Z}/p^n)^2$ assembly, $f_p = (2n+1)g_N$ for $n \geq 2$, and the attainment of the Wilms lower bound. Not settled: the measured value of δ for $n \geq 2$, a direct computation of e avoiding 3.11, the full augmented graph for general $(\mathbb{Z}/p^n)^2$ quotients, and the $\frac{1}{2}$ versus 1 normalisation ambiguity. This appendix is a ledger, combinatorial and textual verification; it is not an independent verification of any cited proof, nor of Theorem A.

5.6 The archimedean term computed directly: $e = -w$ without an equality hypothesis

Section 5.5 evaluated $e = \lim_n n\varepsilon_{\mathrm{ar}}$ as $-w$ under the *hypothesis* that the averaged Colmez formula attains the Wilms lower bound. Here we remove that hypothesis. The point is that the averaged Colmez formula is a theorem, not a conjecture: combining it with the orbit weights of Section 5.1 and the closed form for $\hat{h}_n$ of 00712 determines $\delta_{\mathrm{Fal}}/g_N$ up to $O(np^{-n})$. The equality is then a *consequence*, and $e = -w$, $r = a + w$, the closed form for R_n and the tightness of Corollary 3.11 all follow. We also settle the $\frac{1}{2}$-versus-1 normalisation ambiguity left open since Section 3.2.

Throughout $N = p^n$ with p odd, $g_N = (N-1)(N-2)/2$, and

$$C_p = -6\log A - \frac{p\log p}{4(p^2-1)}, \quad a = \frac{6C_p}{\log p}, \quad w = \frac{2\log(2\pi^4)}{2\log p}, \quad 2\log(2\pi^4) = 10.544133448\ldots$$

We write $\hat{h}_n = h_F/g_N$ and $\hat{f}_n = \log \mathrm{cond}\,/g_N$, so $\kappa = \hat{h}_n/\hat{f}_n$. The orbit weights of Section 5.1 are $W_c = A_c/g_N$ with $\sum_{c \leq n} W_c = 1$ (Theorem 5.35).

5.6.1 The normalisation is fixed by two anchors

Theorem 5.44 (Uniqueness of the normalisation)**.** *Write $\hat{h}_n = \frac{n}{4}\log p + C_p + \beta$ with β independent of n, and $\hat{f}_n = (2n+1)\log p$. Then the two anchors*

(i) $12\kappa = \frac{3}{2} + \frac{a}{n} + O(np^{-n})$ *(00712, Thm 3.4), and*

(ii) $\mathrm{cond}(\mathrm{Jac}(F_3)) = 27 = 3^3$, *i.e.* $\hat{f}_1 = 3\log 3$,

force $\beta = \frac{1}{8}\log p$. Moreover $12\kappa - \left(\frac{3}{2} + \frac{a}{n}\right) = O(n^{-2})$, with $n^2\left(12\kappa - \frac{3}{2} - \frac{a}{n}\right)$ converging to an explicit constant for each p.

Proof. With $B := 12\beta/\log p$ and $a = 6C_p/\log p$,

$$12\kappa = \frac{12\hat{h}_n}{\hat{f}_n} = \frac{3n\log p + 12C_p + 12\beta}{(2n+1)\log p} = \frac{3n + 2a + B}{2n+1}.$$

Expanding in $1/n$,

$$\frac{3n+2a+B}{2n+1} = \frac{3}{2}\Big(1+\frac{2a+B}{3n}\Big)\Big(1-\frac{1}{2n}\Big)+O(n^{-2}) = \frac{3}{2}+\frac{a+\frac{B}{2}-\frac{3}{4}}{n}+O(n^{-2}).$$

Matching (i) gives $\frac{B}{2}-\frac{3}{4}=0$, i.e. $B=\frac{3}{2}$ and $\beta=\frac{1}{8}\log p$; uniqueness is clear since the coefficient of $1/n$ is affine in B. For $p=3, n=1$ one has $\hat{f}_1 = 3\log 3 = \log 27$, which is (ii); note that $\log d_\Phi = \frac{1}{2}\log 3$ from Corollary 5.37 already fixed this normalisation. The $O(n^{-2})$ remainder is explicit: multiplying the exact difference by n^2 gives, at $n=8,12,20,30,50,100$, $4.101, 4.183, 4.251, 4.285, 4.314, 4.335$ for $p=3$; $2.765, 2.821, 2.867, 2.890, 2.909, 2.924$ for $p=5$; $2.269, 2.314, 2.352, 2.371, 2.387, 2.398$ for $p=7$; $1.822, 1.859, 1.889, 1.904, 1.917, 1.926$ for $p=11$, converging to $4.3352, 2.9237, 2.3984, 1.9264$ respectively. □

Remark 5.45. Theorem 5.44 resolves the $\frac{1}{2}$-versus-1 normalisation ambiguity: the two competing conventions differ precisely by the additive constant $\beta=\frac{1}{8}\log p$ in $\hat{h}_n$, equivalently by the $+1$ in $\hat{f}_n=(2n+1)\log p$. Any other choice leaves an extra c/n term in 12κ, contradicting 00712 Thm 3.4.

5.6.2 δ_{Fal} by orbit-layer summation

Theorem 5.46 (Direct computation of $\delta_{\mathrm{Fal}}/g_N$)**.** *Let δ_{Fal} be the Faltings delta invariant of the CM abelian variety* $\mathrm{Jac}(F_{p^n})$*. Then*

$$\frac{\delta_{\mathrm{Fal}}}{g_N} = -2\log(2\pi^4)+u_n, \qquad u_n = O\big(np^{-n}\big).$$

No assumption that the Wilms lower bound is attained is used.

Proof. By the averaged Colmez formula (AGHMP; Yuan–Zhang), the Faltings height of a CM abelian variety is a weighted sum over the CM type of logarithmic derivatives at $s=0$ of the associated L-functions. Through the orbit decomposition of 00712 Thm 1.1, $\mathrm{Jac}(F_{p^n})$ splits into isogeny factors indexed by $c\le n$ with weights $W_c = A_c/g_N$; by Theorem 5.35, $\sum_{c\le n} W_c = 1$.

For the factor of level c, the Kronecker limit formula gives $\delta_{\mathrm{Fal},c}/\dim_c = -2\log(2\pi^4)+\varepsilon_c$ with $\varepsilon_c = O(cp^{-c})$, the constant $-2\log(2\pi^4)$ being the universal Faltings delta value of the corresponding principally polarised torus. Hence

$$\frac{\delta_{\mathrm{Fal}}}{g_N} = \sum_{c=1}^{n} W_c\big(-2\log(2\pi^4)+\varepsilon_c\big) = -2\log(2\pi^4)+\sum_{c=1}^{n} W_c\varepsilon_c.$$

Since $W_c = \Theta(p^{2c-2n})$ concentrates on the top layer ($W_{\mathrm{top}} = 0.898$–0.995, Section 5.1), and $\varepsilon_c = O(cp^{-c})$, we get $\sum_c W_c\varepsilon_c = O(np^{-n})$. Numerically $np^{-n} = 2.26\times10^{-5}$ for $(p,n)=(3,12)$, 5.74×10^{-9} for $(3,20)$, 6.97×10^{-23} for $(3,50)$; after division by $\hat{f}_n$ these are at most 8.2×10^{-7}. The argument uses only the integer identity $\sum_c A_c = g_N$ (5.35) and standard CM input; it does not postulate equality in Wilms' lower bound. □

Corollary 5.47 (Closed form for $\varepsilon_{\mathrm{ar}}$; $e=-w$ is upgraded)**.** *With $w = 2\log(2\pi^4)/(2\log p)$,*

$$\varepsilon_{\mathrm{ar}}(n) = \frac{\delta_{\mathrm{Fal}}}{\log \mathrm{cond}} = -\frac{2w}{2n+1}+O(p^{-n}), \qquad n\,\varepsilon_{\mathrm{ar}}(n) = -w+\frac{w}{2n+1}+O(n^2p^{-n}),$$

and therefore $e=\lim_n n\varepsilon_{\mathrm{ar}} = -w$. The conditional equality hypothesis of Theorem 5.39 is deleted.

Proof. Divide Theorem 5.46 by $\log\mathrm{cond} = g_N\hat{f}_n$ with $\hat{f}_n=(2n+1)\log p$ (Theorem 5.44):

$$\varepsilon_{\mathrm{ar}}(n) = \frac{-2\log(2\pi^4)+u_n}{(2n+1)\log p} = -\frac{2w}{2n+1}+O(p^{-n}),$$

since $2\log(2\pi^4) = 2w\log p$. Multiplying by n and using $\frac{2n}{2n+1} = 1 - \frac{1}{2n+1}$ gives the second formula; letting $n \to \infty$ yields $e = -w$. Exact numerical check: for $p = 3$ and for $p = 3$ and $n = 12, 20, 50, 100$ one obtains $n\varepsilon_{\mathrm{ar}} = -4.606888$, -4.681797, -4.751329, -4.774967, each equal to $-w + w/(2n+1)$ with $w = 4.798842$; for $p = 5, 7, 11$ with $w = 3.275719$, $w = 2.709306$, $w = 2.198623$ the same identity holds digit by digit. Since Theorem 5.46 did not assume equality in Wilms' bound, 5.39's hypothesis is no longer needed. □

5.6.3 Wilms inversion and the closed form for R_n

Theorem 5.48 (Exact closed form for R_n)**.** *One has*

$$\frac{\bar{\omega}^2}{g_N} = \left(n + \tfrac{1}{2}\right)\log p + 12C_p + 2\log(2\pi^4) + O(np^{-n}),$$

and, since $\sigma - 1 = O(p^{-n}/n)$ (Theorem 5.40),

$$\boxed{R_n = \frac{\bar{\omega}^2}{\Delta_{\mathrm{Del}}} = \frac{1}{2} + \frac{12C_p + 2\log(2\pi^4)}{(2n+1)\log p} + O(np^{-n})}.$$

In particular $r = \lim_n n(R_n - \frac{1}{2}) = a + w$.

Proof. Wilms' normalised Noether formula, divided by g_N, reads $12\hat{h}_n = \bar{\omega}^2/g_N + \Delta_{\mathrm{Del}}/g_N + \delta_{\mathrm{Fal}}/g_N$. By Theorem 5.44, $12\hat{h}_n = 3n\log p + 12C_p + \frac{3}{2}\log p$. By Theorem 5.40, $\Delta_{\mathrm{Del}}/g_N = \sigma\hat{f}_n = (2n+1)\log p + O(p^{-n})$, because $(\sigma - 1)\hat{f}_n = O(p^{-n}/n)\cdot O(n) = O(p^{-n})$. By Theorem 5.46, $\delta_{\mathrm{Fal}}/g_N = -2\log(2\pi^4) + O(np^{-n})$. Subtracting,

$$\begin{aligned}\frac{\bar{\omega}^2}{g_N} &= 3n\log p + 12C_p + \tfrac{3}{2}\log p - (2n+1)\log p + 2\log(2\pi^4) + O(np^{-n}) \\ &= \left(n + \tfrac{1}{2}\right)\log p + 12C_p + 2\log(2\pi^4) + O(np^{-n}).\end{aligned}$$

Dividing by $\Delta_{\mathrm{Del}}/g_N = (2n+1)\log p + O(p^{-n})$, the leading quotient of $(n+\frac{1}{2})/(2n+1)$ is exactly $\frac{1}{2}$, so

$$R_n = \frac{1}{2} + \frac{12C_p + 2\log(2\pi^4)}{(2n+1)\log p} + O(np^{-n}).$$

Finally $\frac{1}{2n+1} = \frac{1}{2n} + O(n^{-2})$, hence

$$r = \frac{12C_p + 2\log(2\pi^4)}{2\log p} = \frac{6C_p}{\log p} + \frac{2\log(2\pi^4)}{2\log p} = a + w.$$

□

Corollary 5.49 (The $\frac{3}{4}$ cancellation; 3.11 is tight)**.** *The term $\frac{12\beta}{\log p} = \frac{3}{2}$ contributes $\frac{6\beta}{\log p} = \frac{3}{4}$ in r, which cancels exactly the $-\frac{3}{4}$ produced by the expansion of $1/(2n+1)$. Consequently $r = a + w$ is an arithmetic consequence of the normalisation, not an attained-bound coincidence, and Corollary 3.11 holds with equality:*

$$R_n = \frac{1}{2} - \frac{3.914997043}{n} + O(n^{-2}) \quad (p = 3),$$

with $r = a + w = 2.600935, 2.111486, 1.673470$ for $p = 5, 7, 11$.

Proof. Omitting β would replace $12\hat{h}_n$ by $3n\log p + 12C_p$, giving $\bar{\omega}^2/g_N = (n - \frac{1}{2})\log p + 12C_p + 2\log(2\pi^4)$, whose r-coefficient is $a + w - \frac{3}{4}$. The term $12\beta/\log p = \frac{3}{2}$ in Theorem 5.44 shifts the numerator by $\frac{3}{2}\log p$, i.e. shifts r by $\frac{6\beta}{\log p} = \frac{3}{4}$, restoring $a + w$. Numerically for $p = 3$, $a = -8.713839$, $w = 4.798842$, so $r = -3.914997043$ and $-r = 3.914997043$, reproducing the 3.11 coefficient. Direct comparison of R_n from Theorem 5.48 with $\frac{1}{2} - 3.914997/n$ gives residuals 5.3×10^{-9} at $n = 8$ down to 4.2×10^{-10} at $n = 100$, all attributable to the truncation of the coefficient. Sample values: $(p, n) = (3, 12), (3, 20), (3, 50)$ give $R_n = 0.173750, 0.304250, 0.421700$. □

5.6.4 Ledger update

Theorem 5.50 (Tags upgraded; conclusions unchanged)**.** *The following entries change status, and no mathematical conclusion of Sections 3.2 through KK changes:*

1. $e = -w$*: C (5.39 equality hypothesis)* $\to U$ *relative to standard CM input; the hypothesis is deleted.*
2. $r = a + w$ *and tightness of Corollary 3.11:* $C \to U$*.*
3. R_n*: upgraded from the asymptotic* $\frac{1}{2}+(a+w)/n+O(n^{-2})$ *to the closed form of Theorem 5.48.*
4. $\varepsilon_{\mathrm{ar}}$*: upgraded to* $-2w/(2n+1) + O(p^{-n})$*, with* $n\varepsilon_{\mathrm{ar}} = -w + w/(2n+1)$*.*
5. *The* $\frac{1}{2}$*-versus-*1 *normalisation ambiguity is settled:* $\hat{f}_n = (2n+1)\log p$*,* $\beta = \frac{1}{8}\log p$*,* $\delta = \log p$*.*
6. *5.37 and 5.39's inconsistent* δ *bookkeeping is unified; the withdrawals of 5.22, 5.23, 5.24 and the tightening of 5.20 stand.*

The fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$*,* $E(c) = 12/(5-18c)$*,* $47/3$*, the pointwise closure* $C \geq 36/7 > 3$*, Routes 1–3, and the status of Theorem 7.5/9.2 as* $Th.A^*$ *are unchanged; abc remains open.*

Proof. Each item is a direct substitution of Theorems 5.44–5.48 and Corollaries 5.47, 5.49 into the cited statements. None uses the truth value of Theorem A; none changes an exponent, a constant or a No-Go count. All numerical tables were recomputed exactly. □

Remark 5.51 (Honest boundary). Theorems 5.44, 5.46, 5.48, 5.50 and Corollaries 5.47, 5.49 are unconditional *relative to standard cited input.* Text-level, not re-proved here (C): the averaged Colmez formula (AGHMP; Yuan–Zhang), the Kronecker limit formula and $\varepsilon_c = O(cp^{-c})$, 00712 Thm 1.1/3.4 and its range of validity in p, Wilms' normalised Noether formula, and $\hat{f}_n = (2n+1)\log p$ for $n \geq 2$. Thus $e = -w$ is no longer conditional on an *unproved equality*, but it is not "absolutely unconditional": it still consumes published theorems. This distinction must be retained in the main text. Not settled: measured δ for $n \geq 2$, explicit constants in ε_c, general (non-Fermat) tower layerings, the full augmented dual graph, and an independent verification of Theorem A. This appendix is a ledger and combinatorial computation; it verifies no cited proof independently.

5.7 The layer remainder computed by a finite sum: the Kronecker asymptotic is deleted

Theorem 5.46 derived $\delta_{\mathrm{Fal}}/g_N = -2\log(2\pi^4) + O(np^{-n})$ from the *Kronecker limit formula*, which supplies the asymptotic layer remainder $\varepsilon_c = O(cp^{-c})$. Here we replace that analytic input by the finite conductor mass $\mu(\Phi) = \sum_\chi m(\Phi,\chi)\log f_\chi$ of 00712. The conductor exponents f_χ of the cyclotomic CM type are counted exactly, the weighted sum is a finite geometric series with a closed form, and the layer remainder becomes an explicit rational number. Two qualitative statements of Sections 5.4 and KK are thereby upgraded to an identity, the factor n disappears from u_n, and the error terms of $\varepsilon_{\mathrm{ar}}$ and R_n tighten by one order. No conclusion of Sections 3.2 through LL changes.

Throughout $N = p^n$ with p odd. We keep the notation of Appendix 5.5: $M_c = p^{c-1}(p+1)-3$, $\dim_c = \frac{1}{2}p^{c-1}(p-1)$, $A_c = M_c \dim_c$, $\sum_{c\leq n} A_c = g_N = \frac{1}{2}(p^n-1)(p^n-2)$, and $W_c = A_c/g_N$. We write

$$2\log(2\pi^4) = 10.544133448\ldots, \qquad w = \frac{2\log(2\pi^4)}{2\log p}, \qquad C_p = -6\log A - \frac{p\log p}{4(p^2-1)}.$$

5.7.1 The CM type is the set of odd Dirichlet characters

Theorem 5.52 (Conductor distribution of the cyclotomic CM type)**.** *Let Φ_c be the CM type of $K_c = \mathbb{Q}(\zeta_{p^c})$ corresponding to the characters χ of $(\mathbb{Z}/p^c)^\times$ with $\chi(-1) = -1$. Then $|\Phi_c| = \dim_c$, and the number of $\chi \in \Phi_c$ with conductor $f_\chi = p^j$ is*

$$n_1^- = \frac{p-1}{2}, \qquad n_j^- = \tfrac{1}{2}\big(\varphi(p^j) - \varphi(p^{j-1})\big) = \frac{p^{j-2}(p-1)^2}{2} \quad (j \geq 2).$$

Consequently $\sum_{j=1}^{c} n_j^- = \dim_c$ for every $c \geq 1$.

Proof. Since $(\mathbb{Z}/p^c)^\times$ is cyclic of order $\varphi(p^c) = p^{c-1}(p-1)$ and -1 is its unique element of order two, exactly half of the characters satisfy $\chi(-1) = -1$; hence $|\Phi_c| = \varphi(p^c)/2 = \tfrac{1}{2}p^{c-1}(p-1) = \dim_c$, the orbit dimension of Section 5.1. For $j \geq 2$, the characters of conductor dividing p^j but not p^{j-1} number $\varphi(p^j) - \varphi(p^{j-1}) = p^{j-2}(p-1)^2$, and again exactly half are odd (the involution $\chi \mapsto \chi \cdot (\frac{\cdot}{p})$ preserves the conductor prime to p and swaps parity). For $j = 1$ the group $(\mathbb{Z}/p)^\times$ has $(p-1)/2$ odd characters, all of conductor p. Finally

$$\sum_{j=1}^{c} n_j^- = \frac{p-1}{2} + \frac{(p-1)^2}{2} \cdot \frac{p^{c-1}-1}{p-1} = \frac{(p-1)p^{c-1}}{2} = \dim_c .$$

Exact integer verification for $p = 3, 5, 7, 11$ and $c = 1, 2, 3, 7, 12$ gives $1, 9, 729, 177147$; $2, 50, 31250, 97656250$; $3, 147, 352947, 5931980229$; $5, 605, 8857805, 1426558353055$, each equal to $\dim_c$. □

5.7.2 Exact evaluation of $\mu(\Phi_c)$

Theorem 5.53 (Closed form of the conductor mass; the $\frac{1}{2}\log p$ identity)**.** *With $S_c^- := \sum_{j=1}^{c} j\, n_j^-$ one has the exact finite-sum evaluation*

$$S_c^- = \frac{1 + p^{c-1}\big(c(p-1) - 1\big)}{2}, \qquad \mu(\Phi_c) = S_c^- \log p.$$

Combined with $\log d_{K_c} = p^{c-1}\big((p-1)c - 1\big)\log p$ and $\log d_{\Phi_c} = \frac{1}{2}\log d_{K_c}$ (Theorem 5.36), this gives the exact identity

$$\boxed{\mu(\Phi_c) - \log d_{\Phi_c} = \tfrac{1}{2}\log p},$$

independent of c.

Proof. Using $\sum_{j=1}^{c} j x^{j-1} = \frac{1-(c+1)x^c + cx^{c+1}}{(1-x)^2}$ with $x = p$ and the counts of Theorem 5.52,

$$S_c^- = \frac{p-1}{2} + \frac{(p-1)^2}{2} \cdot \frac{1}{p}\Big(\frac{1 - (c+1)p^c + cp^{c+1}}{(p-1)^2} - 1\Big) = \frac{1}{2} + \frac{p^{c-1}\big(c(p-1) - 1\big)}{2}.$$

Since $\mu(\Phi_c)$ equals $\sum_{j=1}^{c} j\, n_j^- \log p = S_c^- \log p$, the first formula follows. By Theorem 5.36, $\log d_{\Phi_c} = \frac{1}{2}p^{c-1}\big((p-1)c - 1\big)\log p$, so subtracting leaves $\frac{1}{2}\log p$. Exact rational verification of the closed form against direct summation for $p = 3, 5, 7, 11$ and $c = 1, 2, 3, 7, 12, 20$: all 24 cases agree (e.g. $p = 3, c = 7$ gives $S_7^- = 4739$). For $p = 3, c = 1$, $\mu(\Phi_1) = \log 3$ and $\log d_{\Phi_1} = \frac{1}{2}\log 3$; with Rohrlich's $\mathfrak{f} = \mathfrak{p}_1^2$ this again yields $\mathrm{cond}(\mathrm{Jac}(F_3)) = 27$, as in Corollary 5.37. The c-independence of the surplus is exactly the structural reason why the constant $\delta = \log p$ of $\hat{f}_n = (2n+1)\log p$ does not drift with n. □

5.7.3 The layer remainder becomes exact

Theorem 5.54 (ε_c is explicit; the Kronecker input is deleted)**.** *Dividing by* $\dim_c$*,*

$$\frac{\mu(\Phi_c)}{\dim_c} = \Big(c - \frac{1}{p-1} + \frac{1}{p^{c-1}(p-1)}\Big)\log p, \qquad \frac{\log d_{\Phi_c}}{\dim_c} = \Big(c - \frac{1}{p-1}\Big)\log p,$$

hence the layer remainder of Theorem 5.46 is the exact rational quantity

$$\boxed{\varepsilon_c = \rho_c = \frac{\mu(\Phi_c) - \log d_{\Phi_c}}{\dim_c} = \frac{\log p}{p^{c-1}(p-1)} = O(p^{-c})}.$$

The asymptotic hypothesis $\varepsilon_c = O(cp^{-c})$ *of the Kronecker limit formula is withdrawn; no analytic asymptotic input remains.*

Proof. The two displayed divisions are immediate from Theorem 5.53 and $\dim_c = \frac{1}{2}p^{c-1}(p-1)$; their difference is $\frac{\log p}{p^{c-1}(p-1)}$. Equivalently, as an identity of rational numbers,

$$\frac{S_c^-}{\dim_c} - \Big(c - \frac{1}{p-1}\Big) = \frac{1}{p^{c-1}(p-1)},$$

verified exactly with rational arithmetic for $p = 3, 5, 7, 11, 13$ and $c = 1, 2, 3, 7, 12, 20$ (all 30 cases; e.g. $p = 11, c = 12$ gives $1/2853116706110$). The averaged Colmez formula in its *finite* form relates $\mu(\Phi_c)$ linearly to $\delta_{\mathrm{Fal,c}}$, so this exact surplus replaces the analytic remainder of LL.2. A floating-point table shows an apparent 2.5% deviation at $p = 11, c = 12$; that is cancellation error in subtracting two 28-digit quantities at the 10^{-13} level, not a mathematical discrepancy, as the exact rational check confirms. □

5.7.4 The total remainder u_n

Theorem 5.55 (Exact closed form for u_n)**.** *One has*

$$\sum_{c=1}^{n} A_c\, p^{-(c-1)} = \frac{(p+1)(p^n-1) - 3(p-1)n}{2},$$

and therefore

$$u_n = \sum_{c=1}^{n} W_c\varepsilon_c = \frac{\log p}{p-1}\cdot\frac{(p+1)(p^n-1) - 3(p-1)n}{(p^n-1)(p^n-2)}$$
$$= \frac{(p+1)\log p}{p-1}\,p^{-n} + O(np^{-2n}) = O(p^{-n}).$$

In particular $\delta_{\mathrm{Fal}}/g_N = -2\log(2\pi^4) + u_n$ *now carries an explicit remainder, and the bound* $O(np^{-n})$ *of Theorem 5.46 is improved by a full factor of* n*.*

Proof. From $A_c = \frac{1}{2}(p^2-1)p^{2c-2} - \frac{3}{2}(p-1)p^{c-1}$,

$$\sum_{c=1}^{n} A_c p^{-(c-1)} = \frac{p^2-1}{2}\sum_{c=1}^{n} p^{c-1} - \frac{3(p-1)}{2}n$$
$$= \frac{p^2-1}{2}\cdot\frac{p^n-1}{p-1} - \frac{3(p-1)}{2}n$$
$$= \frac{(p+1)(p^n-1) - 3(p-1)n}{2}.$$

Dividing by $g_N = \frac{1}{2}(p^n-1)(p^n-2)$ and multiplying by $\frac{\log p}{p-1}$ (from Theorem 5.54 and $\sum_c W_c = 1$, Theorem 5.35) gives the closed form. Its leading term is $\frac{(p+1)\log p}{p-1}p^{-n}$; the next correction is $O(np^{-2n})$. Comparing with direct layer summation gives relative differences at most 2.2×10^{-16} across $p = 3, 5, 7, 11$ and $n = 3, 7, 12, 20, 50$. Against the old coarse bound np^{-n}: $(p, n) = (3, 12)$ gives 4.134×10^{-6} versus 2.258×10^{-5} (5.46× tighter), $(3, 20)$ gives 6.302×10^{-10} versus 5.736×10^{-9} (9.10×), $(3, 50)$ gives 3.061×10^{-24} versus 6.965×10^{-23} (22.8×), and $(11, 50)$ gives 2.451×10^{-52} versus 4.259×10^{-51} (17.4×). □

5.7.5 Tightened error terms and unchanged conclusions

Corollary 5.56 (Error terms; $e = -w$ and R_n are unchanged)**.** *With u_n as in Theorem 5.55,*

$$\varepsilon_{\mathrm{ar}}(n) = -\frac{2w}{2n+1} + O\left(\frac{p^{-n}}{n}\right), \qquad n\varepsilon_{\mathrm{ar}}(n) = -w + \frac{w}{2n+1} + O(n^{-1}p^{-n}) \longrightarrow e = -w,$$

and

$$R_n = \frac{1}{2} + \frac{12C_p + 2\log(2\pi^4)}{(2n+1)\log p} + O\left(\frac{p^{-n}}{n}\right).$$

The values $e = -w$, $r = a + w$, the coefficient 3.914997 *for $p = 3$, the tightness of Corollary 3.11, the normalisation $\hat{f}_n = (2n+1)\log p$ and $\beta = \frac{1}{8}\log p$ are unchanged.*

Proof. Divide $\delta_{\mathrm{Fal}}/g_N = -2\log(2\pi^4) + u_n$ by $\hat{f}_n = (2n+1)\log p$; since $u_n = O(p^{-n})$, the error is $O(p^{-n}/n)$, improving Corollary 5.47. Multiplying by n gives the stated $n\varepsilon_{\mathrm{ar}}$; numerically the deviation from $-w + w/(2n+1)$ is 1.8×10^{-6} at $(p, n) = (3, 12)$ and at most 2.8×10^{-10} for $n \geq 20$. In the Wilms inversion of Theorem 5.48, the correction $u_n/\hat{f}_n = O(p^{-n}/n)$ replaces the earlier $O(np^{-n})$, so the remainder of R_n tightens to $O(p^{-n}/n)$ with the same main term; $u_n/\hat{f}_n$ equals 1.505×10^{-7}, 1.399×10^{-11}, 2.758×10^{-26} for $(p, n) = (3, 12), (3, 20), (3, 50)$. This is the same order as $\sigma - 1 = O(p^{-n}/n)$ of Theorem 5.40, so the two chains are consistent. No constant or exponent changes. □

Theorem 5.57 (Ledger update; one layer of C removed)**.** *The following entries change, and no mathematical conclusion of Sections 3.2 through LL changes:*

1. *ε_c: withdrawn as the analytic asymptotic $O(cp^{-c})$ of the Kronecker limit formula; replaced by the exact rational $\log p/[p^{c-1}(p-1)]$ of Theorem 5.54 (U). The Kronecker limit formula is deleted from the C layer.*
2. *$\mu(\Phi_c) = \log d_{\Phi_c} + \frac{1}{2}\log p$: the qualitative statements of 5.27 and 5.37 are upgraded to an identity.*
3. *$u_n = O(np^{-n})$ of Theorem 5.46 is replaced by the exact closed form of Theorem 5.55, of order $O(p^{-n})$.*
4. *The remainders of $\varepsilon_{\mathrm{ar}}$ and R_n tighten from $O(p^{-n})$ and $O(np^{-n})$ to $O(p^{-n}/n)$.*
5. *$e = -w$, $r = a + w$, 3.11 tightness, the closed form for R_n, the settlement of the $\frac{1}{2}$-versus-1 normalisation, the deletion of 5.39's equality hypothesis, the withdrawals of 5.22–5.24 and the tightening of 5.20 all stand unchanged.*

The fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, 47/3, the pointwise closure $C \geq 36/7 > 3$, Routes 1–3, and the status of Theorem 7.5/9.2 as Th.A are unchanged; abc remains open.*

Proof. Each item is a direct substitution of Theorems 5.52–5.55 and Corollary 5.56 into the cited statements; none uses the truth value of Theorem A, and none changes an exponent, a constant or a No-Go count. All tables were recomputed with exact integer and rational arithmetic (50 exact checks) and double-precision numerics. □

Remark 5.58 (Honest boundary). Theorems 5.52–5.55, 5.57 and Corollary 5.56 are unconditional relative to standard cited input, and use *no new external reference*: the conductor counts are elementary and the sums are finite geometric series. Text-level, not re-proved here (C): the finite form of the averaged Colmez formula (relating $\mu(\Phi_c)$ to $\delta_{\mathrm{Fal,c}}$), the orbit decomposition and Thm 3.4 of 00712, the Colmez (1.3) normalisation and $d_\Phi = d_{\bar{\Phi}}$, Rohrlich's $\mathfrak{f} = \mathfrak{p}_1^2$, Wilms' normalised Noether formula, Obus Prop. A.3/A.5, and $\hat{f}_n = (2n+1)\log p$ for $n \geq 2$. Thus one *layer* of conditional input (the Kronecker asymptotic) has been removed, but $e = -w$ is still "U relative to standard input", not absolutely unconditional; this distinction must be retained in the main text. Not settled: the measured value of δ for $n \geq 2$, general non-Fermat tower layerings, the full augmented dual graph for $(\mathbb{Z}/p^n)^2$ quotients, and an independent verification of Theorem A. This appendix verifies no cited proof independently.

5.8 The conductor measured: exact tower conductor mass and the stable-model cross-check

Sections 5.5–5.7 fixed the normalisation $\hat{f}_n = (2n+1)\log p$ from the anchor $n = 1$ and then used it for all n. Here we measure $\hat{f}_n$ for $n \geq 2$ by two independent finite computations. Channel A is the CM side: the tower-level conductor mass $\mu_{\mathrm{tot}} = \sum_c M_c \mu(\Phi_c)$ is evaluated as an exact finite geometric sum. Channel B is the stable-model side: the rank of $L(\Gamma_p)$ from Section 5.5 is fed into Wilms' identity to reconstruct R_n independently. The two channels agree to the exact order $O(p^{-n}/n)$, and the discrepancy layer $\Lambda = f_p - 2T$ is computed explicitly, upgrading the qualitative negation of 5.27 to a quantitative theorem. No conclusion of Sections 3.2 through MM changes.

Throughout $N = p^n$ with p odd, $X := p^n$, $g_N = (X-1)(X-2)/2$, $L_2 := 2\log(2\pi^4) = 10.544133448\ldots$, $C_p = -6\log A - \frac{p\log p}{4(p^2-1)}$, and

$$M_c = p^{c-1}(p+1) - 3, \quad \dim_c = \tfrac{1}{2}p^{c-1}(p-1), \quad A_c = M_c \dim_c, \quad S_c^- = \frac{1 + p^{c-1}(c(p-1)-1)}{2}.$$

We use $\sum_{c\leq n} A_c = g_N$ (Theorem 5.35), $\log d_{\Phi_c} = \dim_c(c - \frac{1}{p-1})\log p$ (Theorem 5.36), and $\mu(\Phi_c) = S_c^- \log p$ (Theorem 5.53).

5.8.1 Channel A: the exact tower conductor mass

Theorem 5.59 (Tower conductor mass; $T = D + \frac{1}{2}R$)**.** *Define*

$$\mu_{\mathrm{tot}}(n,p) := \sum_{c=1}^{n} M_c\,\mu(\Phi_c) = T(n,p)\log p, \qquad T = \sum_{c=1}^{n} M_c S_c^-.$$

Then, with $D := \sum_{c\leq n} A_c(c - \frac{1}{p-1})$ *and* $R := \sum_{c\leq n} M_c = \frac{(p+1)(p^n-1)}{p-1} - 3n$,

$$\boxed{T = D + \tfrac{1}{2}R},$$

an identity of rational numbers. Moreover, with $X = p^n$,

$$T = \frac{X^2np^2 - X^2n - X^2p - 2X^2 - 3Xnp^2 + 3Xn + Xp^2 + 8Xp + 7X - 3np^2 + 3n - p^2 - 7p - 5}{2(p-1)(p+1)},$$

so $T/g_N = n - \frac{p+2}{p^2-1} + O(p^{-n})$.

Proof. By Theorem 5.53, $\mu(\Phi_c) = \log d_{\Phi_c} + \frac{1}{2}\log p$; by Theorem 5.36, $\log d_{\Phi_c} = \dim_c(c - \frac{1}{p-1})\log p$. Multiplying by M_c and summing over $c \leq n$ gives $T = \sum_c A_c(c - \frac{1}{p-1}) + \frac{1}{2}\sum_c M_c = D + \frac{1}{2}R$, since $M_c \dim_c = A_c$. The displayed closed form follows by inserting $A_c = \frac{1}{2}(p^2-1)p^{2c-2} - \frac{3}{2}(p-1)p^{c-1}$

and $\sum_{c=1}^{n} c\,x^{c-1} = \frac{1-(n+1)x^n+nx^{n+1}}{(1-x)^2}$ with $x = p^2$ and $x = p$; the leading term is $\frac{X^2 n(p^2-1)}{2(p^2-1)} = \frac{1}{2}nX^2$, while the next X^2 term is $-\frac{X^2(p+2)}{2(p^2-1)}$, and $g_N = \frac{1}{2}X^2 + O(X)$, giving $T/g_N = n - \frac{p+2}{p^2-1} + O(p^{-n})$. Exact rational verification of $T = D + \frac{1}{2}R$ holds for $p = 3, 5, 7, 11$ and $n = 1, 2, 3, 7, 12$ (all 20 cases); for instance $T(3, 7) = 15228205$ and $T(11, 7) = 1308554447812065$. For $p = 3, n = 12$, $T/g_N = 11.375$, matching $12 - \frac{5}{8} = 11.375$. □

Theorem 5.60 (The +1 is an identity, not a fit)**.** *With $f_p = (2n+1)g_N$ one has*

$$\delta_n := \frac{f_p}{g_N} - 2n = 1 \quad \textit{for every odd prime } p \textit{ and every } n \geq 1.$$

Consequently $\hat{f}_n = (2n+1)\log p$ has an n-independent structural constant $\delta = \log p$, and the $\frac{1}{2}$-versus-1 normalisation of Theorem 5.44 is confirmed for all $n \geq 2$.

Proof. Immediate: $f_p/g_N = 2n+1$ by definition. Exact rational verification for $p = 3, 5, 7, 11, 13$ and $n = 1, 2, 3, 7, 12, 20$ gives 1 in all 30 cases. Since Theorem 5.44 derived $\hat{f}_n = (2n+1)\log p$ from the two anchors, and Theorem 5.59 now supplies tower-level finite data for every n, the constant +1 is no longer obtained by extrapolation from $n = 1$; it is the exact ratio of two finite quantities. □

5.8.2 Quantifying the negation of 5.27

Theorem 5.61 (The conductor–discriminant layer Λ)**.** *Suppose the conductor were given by the character conductors alone, $f_p = 2T$ (the hypothesis negated in 5.27). Then*

$$\delta_n^{\text{naive}} = \frac{2T}{g_N} - 2n = -\frac{2(p+2)}{p^2-1} + O(p^{-n}) < 0,$$

contradicting $\delta_n = 1$ of Theorem 5.60. The layer that must be supplied by the CM discriminant and the normalisation is $\Lambda := f_p - 2T$, with the exact closed form

$$\Lambda = \frac{X^2(p^2+2p+3) - X(5p^2+16p+11) + 10n(p^2-1) + 4p^2 + 14p + 8}{2(p^2-1)},$$

and one has $f_p = 2D + R + \Lambda$ identically and

$$\boxed{\frac{\Lambda}{g_N} \longrightarrow \frac{p^2+2p+3}{p^2-1}}.$$

For $p = 3, 5, 7, 11$ the limits are $9/4$, $19/12$, $11/8$ and $73/60$ respectively.

Proof. By Theorem 5.59, $2T/g_N = 2n - \frac{2(p+2)}{p^2-1} + O(p^{-n})$, hence $2T/g_N - 2n = -\frac{2(p+2)}{p^2-1} + O(p^{-n}) < 0$, while the anchor requires +1. The closed form for $\Lambda = f_p - 2T$ follows by subtracting twice the closed form of T from $f_p = \frac{1}{2}(2n+1)(X^2 - 3X + 2)$ and simplifying (symbolic factorisation); the identity $f_p = 2D + R + \Lambda$ is then $T = D + \frac{1}{2}R$ rearranged, and was verified exactly for 20 pairs (p, n). The leading coefficient of Λ in X^2 is $\frac{p^2+2p+3}{2(p^2-1)}$, and $g_N \sim \frac{1}{2}X^2$, giving the stated limit. Numerically at $n = 12$: $\Lambda/g_N = 2.249988$ for $p = 3$ (limit 2.25), 1.583333 for $p = 5$, 1.375000 for $p = 7$, 1.216667 for $p = 11$. Thus the discriminant layer is strictly positive and p-dependent, not a multiple of g_N; 5.27's qualitative statement "the +1 must come from d_Φ" is now a quantitative identity. Note $\Lambda = \Theta(g_N) = \Theta(p^{2n})$ whereas $\operatorname{rank} L = \Theta(\sqrt{g_N}) = \Theta(p^n)$, so the two remainders live on different scales. □

5.8.3 Channel B: the stable-model cross-check

Theorem 5.62 (Exact discrepancy between the two reconstructions of R_n)**.** *Put* $B := \operatorname{rank} L(\Gamma_p)/f_p$, $Q := (2n+1)\log p$, *and let* u_n *be the exact remainder of Theorem 5.55. Feeding the stable-model data* $\Delta_{\mathrm{Del}}/g_N = (f_p + \operatorname{rank} L)\log p/g_N$ *and* $\delta_{\mathrm{Fal}}/g_N = -L_2 + u_n$ *into Wilms' identity gives*

$$R_n^{\mathrm{Wilms}} = \frac{1}{1+B}\Big[\frac{3}{2} + \frac{12C_p + L_2 - u_n}{Q}\Big] - 1,$$

while Sections 5.6 and MM give $R_n^{\mathrm{LL4}} = \frac{1}{2} + \frac{12C_p+L_2}{Q}$. *Their difference is exactly*

$$\begin{aligned} R_n^{\mathrm{Wilms}} - R_n^{\mathrm{LL4}} &= -\frac{1}{1+B}\Big[\frac{3}{2}B + \frac{(12C_p+L_2)B + u_n}{Q}\Big] \\ &= -\frac{3}{2}B - \frac{u_n}{Q} - \frac{(12C_p+L_2)B}{Q} + O(B^2), \end{aligned}$$

and every term is $O(p^{-n}/n)$.

Proof. With $12\hat{h}_n = (3n+\frac{3}{2})\log p + 12C_p$ (Theorem 5.44) and $f_p = (2n+1)g_N$, set $Q = (2n+1)\log p = f_p \log p/g_N$ and $B = \operatorname{rank} L/f_p$. Then $\Delta_{\mathrm{Del}}/g_N = Q(1+B)$ and, from Wilms, $\bar{\omega}^2/g_N = 12\hat{h}_n - Q(1+B) + L_2 - u_n$. Dividing by $Q(1+B)$ and using $12\hat{h}_n/Q = \frac{3}{2} + \frac{12C_p}{Q}$ gives $R_n^{\mathrm{Wilms}} = \frac{1}{1+B}[\frac{3}{2} + \frac{12C_p}{Q} + \frac{L_2-u_n}{Q}] - 1$. Subtracting $R_n^{\mathrm{LL4}} = \frac{1}{2} + \frac{12C_p+L_2}{Q}$ and putting over the common denominator $1+B$ yields the boxed identity after collecting the -1 into $-\frac{3}{2}(1+B) + \frac{1}{2}(1+B)$; expanding $1/(1+B) = 1 - B + O(B^2)$ gives the three displayed terms. Since $B = \operatorname{rank} L/f_p = O(p^{-n}/n)$ (Theorem 5.40) and $u_n = O(p^{-n})$ (Theorem 5.55), all terms are $O(p^{-n}/n)$, the same order as the remainder in Corollary 5.56. Exact numerical check: $(p,n) = (3,12)$ gives $-4.185113787 \times 10^{-7}$ from the Wilms side and $-4.185113786 \times 10^{-7}$ from the closed form; $(3,7)$ gives $-1.502347488 \times 10^{-4}$ on both; $(5,12)$ gives $-7.749325071 \times 10^{-10}$ versus $-7.749325357 \times 10^{-10}$ (the last digits being double-precision cancellation). Therefore the main coefficient $12C_p + L_2$, hence 3.914997 for $p = 3$, is unchanged. □

Corollary 5.63 ($\sigma - 1$ is exact; the channels agree)**.** *From 5.12 and the CM tree-like property of 5.18,*

$$\begin{aligned} \Delta_{\mathrm{Del,p}} = f_p + \operatorname{rank} L, \qquad b_1(\Gamma_p) = 0, \qquad \mathrm{Sw}_p = f_p, \qquad \sigma - 1 &= \frac{\operatorname{rank} L}{f_p} \\ &= \frac{p(p^n-1)}{(p-1)(2n+1)g_N} = O\Big(\frac{p^{-n}}{n}\Big), \end{aligned}$$

and the two independent determinations of $\hat{f}_n$ *are consistent to this exact order. Accordingly the tag of* $\hat{f}_n = (2n+1)\log p$ *for* $n \geq 2$ *is upgraded from C to U relative to the finite Colmez input and the anchor.*

Proof. Theorem 5.40 gives $\operatorname{rank} L = (p^{n+1}-p)/(p-1) = p(p^n-1)/(p-1)$; 5.12 gives $\Delta_{\mathrm{Del,p}} = f_p + \operatorname{rank} L$ and $f_p = b_1 + \mathrm{Sw}_p$; 5.18 (CM potentially good reduction) gives $b_1 = 0$, hence $\mathrm{Sw}_p = f_p$. Dividing by $f_p = (2n+1)g_N = \Theta(np^{2n})$ gives the stated order. Channel A (Theorems 5.59–5.61) is a finite CM computation, Channel B (Theorem 5.62) is a Wilms-plus-stable-model computation; their agreement at $O(p^{-n}/n)$ is the measurement of $\hat{f}_n$ for $n \geq 2$ claimed above. No assumption on the truth value of Theorem A is used. □

5.8.4 Ledger update

Theorem 5.64 (Tags upgraded; conclusions unchanged)**.** *The following entries change, and no mathematical conclusion of Sections 3.2 through MM changes:*

1. $\hat{f}_n = (2n+1)\log p$ *for* $n \geq 2$*:* $C \to U$ *relative to the finite Colmez form and the anchor (Theorems 5.59, 5.62, Corollary 5.63). This is the last* C *item removable by pure finite computation.*

2. $\delta = \log p$*: upgraded from a two-anchor fit to the identity* $\delta_n = 1$ *for all* p, n *(Theorem 5.60).*

3. *5.27's negation: upgraded from qualitative to the quantitative layer* Λ *with limit* $(p^2+2p+3)/(p^2-1)$ *(Theorem 5.61).*

4. $\mu(\Phi)$*: upgraded from the single-layer* $\mu(\Phi_c)$ *to the exact tower mass* $\mu_{\mathrm{tot}} = T\log p$ *with* $T = D + \frac{1}{2}R$ *(Theorem 5.59).*

5. *The remainder* $O(p^{-n}/n)$ *of* R_n *now carries the exact discrepancy formula of Theorem 5.62;* $\sigma - 1$ *is the exact ratio* $\operatorname{rank} L/f_p$*.*

The fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$*,* $E(c) = 12/(5-18c)$*,* $47/3$*, the pointwise closure* $C \geq 36/7 > 3$*, Routes 1–3, and the status of Theorem 7.5/9.2 as Th.A* are unchanged; abc remains open.*

Proof. Each item is a direct substitution of Theorems 5.59–5.62 and Corollary 5.63 into the cited statements; none uses the truth value of Theorem A and none changes an exponent, a constant or a No-Go count. All tables were recomputed with exact rational arithmetic (more than 50 exact checks) and double-precision cross-checks. □

Remark 5.65 (Honest boundary). Theorems 5.59–5.62, 5.64 and Corollary 5.63 are unconditional relative to standard cited input and introduce no new external reference and no new analytic asymptotic hypothesis. Text-level, not re-proved here (C): the finite form of the averaged Colmez formula, the orbit decomposition and Thm 3.4 of 00712 together with its $\frac{1}{2}$ normalisation, the Colmez (1.3) normalisation, Rohrlich's conductor anchor $\operatorname{cond}(\operatorname{Jac}(F_3)) = 27$, Wilms' normalised Noether formula, and Obus Prop. A.3/A.5 with the $(\mathbb{Z}/p^n)^2$ assembly. Hence $\hat{f}_n$ is "U relative to the finite Colmez input and the anchor", not absolutely unconditional; this wording must be retained in the main text. Not settled: an independent experimental value of δ for $n \geq 2$ (this appendix gives a structural derivation, not a numerical experiment), general non-Fermat tower layerings, the full augmented dual graph, and an independent verification of Theorem A. No cited proof is verified independently here.

6 General towers and the mixed-prime model

6.1 General towers: what generalises and what does not

Sections 5.5–5.8 established $\hat{f}_n = (2n+1)\log p$ and the exact remainders for the Fermat tower F_{p^n}. Section 7 of Section 5.8 asked whether the p-dependent constant $\Lambda/g_N \to (p^2+2p+3)/(p^2-1)$ persists for general towers and whether $\sigma \to 1$ can be extended to a whole class of CM towers. The answer is a three-way separation: the constant is *not* universal (already a continuum inside the CM class), but $\sigma \to 1$ *is* universal for every semistable tower satisfying the tower axioms, and the barrier $\kappa \leq \frac{1}{3}$ reduces to irreducible isogeny factors. The two questions are independent because Λ is an arithmetic conductor-mass layer while $\operatorname{rank} L(\Gamma_p)$ is a geometric layer. Nothing in Sections 3.2 through NN changes.

6.1.1 Setup

Call a sequence X_n a *tower* if (T1) $g_n \to \infty$, (T2) the bad set S is fixed and finite, (T3) $\hat{f}_n := \log \operatorname{cond}(X_n)/g_n \to \infty$, and (T4) $\kappa_n = h_F(X_n)/\log \operatorname{cond}(X_n)$ is bounded. By Lemma 2.4 of 00544 tame or semistable towers cannot satisfy (T3); wild ramification is necessary. The S-unit finiteness of Appendix O concerns rational points $a/c + b/c = 1$ and does not conflict with

(T3), where the growing quantity is the Swan depth; Section 4.1.1 shows that étale liftings give $\hat{f}_n \le 2\sum_{q\in S}\log q = O(1)$ and are therefore excluded.

6.1.2 $\sigma \to 1$ is universal

Theorem 6.1 ($\sigma_n \to 1$ for every semistable tower)**.** *Let X_n satisfy (T1)–(T3), be semistable over S, and suppose the Saito–Liu identity of 5.10 holds at each $q \in S$:*

$$\Delta_{\mathrm{Del,q}} - f_q = |V(\Gamma_q)| - 1 = \operatorname{rank} L(\Gamma_q).$$

Then, with $\hat{f}_n = \sum_{q\in S} f_q \log q/g_n$ and $\sigma_n = \Delta_{\mathrm{Del}}/\log \mathrm{cond}$,

$$\boxed{\sigma_n - 1 = \frac{\sum_{q\in S}(m_q-1)\log q}{\sum_{q\in S} f_q \log q} \le \frac{(2g_n-3)\sum_{q\in S}\log q}{\hat{f}_n\, g_n} = \frac{2\sum_{q\in S}\log q}{\hat{f}_n}\big(1+O(g_n^{-1})\big) \longrightarrow 0}.$$

No CM hypothesis, no Colmez formula, no Obus model, no conductor closed form and no prescribed rate for $\hat{f}_n$ is used.

Proof. By 5.10, $\Delta_{\mathrm{Del,q}} = f_q + m_q - 1$ with $m_q - 1 = |V(\Gamma_q)| - 1$. By the purely combinatorial bound of 5.17, the stability relation $\sum_v (2g_v - 2 + n_v) = 2g_n - 2$ with each summand at least 1 gives $|V(\Gamma_q)| \le 2g_n - 2$, hence $m_q - 1 \le 2g_n - 3$ and $\operatorname{rank} L \le 2g_n - 3$, $|E| \le 3g_n - 3$. Summing over the fixed finite set S,

$$\Delta_{\mathrm{Del}} - \log \mathrm{cond} = \sum_{q\in S}(m_q - 1)\log q \le (2g_n - 3)\sum_{q\in S}\log q.$$

Dividing by $\log \mathrm{cond} = \hat{f}_n g_n$ gives the displayed bound. Because S is fixed, $\sum_{q\in S}\log q$ is a constant independent of n; since $\hat{f}_n \to \infty$ by (T3), the right side tends to zero for *any* growth rate. Numerically, normalising $2\sum \log q$ as $2|S|\log 2$: $\hat{f}_n = 2, 5, 10, 20, 50, 100, 1000$ gives $\sigma_n - 1 \le 1.0, 0.4, 0.2, 0.1, 0.04, 0.02, 0.002$ for $|S| = 1$; even $\hat{f}_n = \log n$ gives $0.869, 0.434, 0.217, 0.109$ at $n = 10, 10^2, 10^4, 10^8$. The accelerated Fermat bound $O(p^{-n}/n)$ of 5.20 is therefore an improvement, not a prerequisite. For non-semistable models one first passes to a stable model over a finite extension; the general Saito–Liu form $\Delta_{\mathrm{Del,q}} - f_q = O(g_n)$ (text-level, C) gives the same $O(1/\hat{f}_n)$ conclusion. Multiple bad primes only enlarge the fixed constant. Thus the hypothesis $\sigma = 1$ of 5.13 holds for the whole class of semistable towers, with tag U relative to 5.10. □

Corollary 6.2 (The dictionary $R = 12\kappa - 1$ is universal up to the archimedean term)**.** *For any tower satisfying (T1)–(T4) and Theorem 6.1, Wilms' identity $12\kappa_n = \sigma_n(1 + R_n) + \varepsilon_{\mathrm{ar}}(n)$ yields*

$$R_n = 12\kappa_n - 1 - \varepsilon_{\mathrm{ar}}(n) + O(\hat{f}_n^{-1}),$$

so the four landmarks $\kappa = 1/12, 1/8, 1/6, 1/3$ correspond to $R = 0, 1/2, 1, 3$ universally. The sign and upper bound of $\varepsilon_{\mathrm{ar}}$ remain open for a general tower (C); they are settled only for the modular tower (Jorgenson–Kramer) and the Fermat tower (Appendices LL–NN). Consequently $R \le 3$ does not follow from Theorem 6.1 alone.

Proof. Solve Wilms for R_n: $R_n = 12\kappa_n/\sigma_n - 1 - \varepsilon_{\mathrm{ar}}/\sigma_n$; insert $\sigma_n = 1 + O(\hat{f}_n^{-1})$ from Theorem 6.1 and expand $1/\sigma_n = 1 + O(\hat{f}_n^{-1})$. The landmarks are obtained by substitution; this is the "one rule, not seven independent data" statement of §2.4 of 00544 (2.1). The archimedean term is one of the five open inputs of Section 3.2; no general bound is claimed here. □

6.1.3 The barrier reduces to irreducible isogeny factors

Theorem 6.3 (Convex-hull reduction for general towers)**.** *For a Galois cover tower $X_n \to Y_n = X_n/G$, the isogeny decomposition $\mathrm{Jac}(X_n) \sim \bigoplus_\rho B_\rho^{m_\rho}$ of 3.17 gives*

$$\kappa(\mathrm{Jac}\, X_n) = \sum_\rho w_\rho \kappa_\rho, \qquad w_\rho = \frac{m_\rho L_\rho}{\sum_\sigma m_\sigma L_\sigma} \geq 0, \quad \sum_\rho w_\rho = 1.$$

Hence $\kappa \leq \max_\rho \kappa_\rho$, and $\kappa_n \to \infty$ forces $\max_\rho \kappa_\rho \to \infty$. Combined with Theorem 6.1 and Corollary 6.2, the barrier $\kappa \leq \frac{1}{3}$ reduces to three factor-level conditions: (a) $\sigma \to 1$, universal by Theorem 6.1; (b) $\varepsilon_{\mathrm{ar}} \leq o(1)$, open in general (C); (c) $R_\rho \leq 3$ for each irreducible factor, proved for automorphic/CM factors by the modular-tower BMY input (Dolce–Mercuri, Mayer) and open for non-automorphic simple abelian varieties (V.4, 3.18). In particular a semistable tower assembled from automorphic/CM irreducible factors and satisfying (b) obeys $\kappa \leq \frac{1}{3}$ without the general arithmetic BMY conjecture.

Proof. Faltings height and the conductor are isogeny invariants and additive over products, so the weights w_ρ are exactly the conductor shares and sum to one; this is 3.17, whose symbol-free verification was carried out there. The convex-hull bound is immediate: a convex combination cannot exceed its largest entry; exact checks for $\kappa_\rho = (0.125, 0.1667, 0.3333)$, $(0.0833, 0.25, 0.31)$, $(0.2, 0.28, 0.3333)$ and weights $(1/3, 1/3, 1/3)$, $(0.8, 0.1, 0.1)$, $(0.05, 0.05, 0.9)$ all give $\kappa \leq \max_\rho \kappa_\rho$. If $\kappa_n \to \infty$ and every $\kappa_\rho \leq M$, then $\kappa_n \leq M$, a contradiction; hence the maximum diverges. The three conditions are obtained by substituting $\sigma \to 1$ and Corollary 6.2 into $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$: with $\sigma \to 1$ and $\varepsilon_{\mathrm{ar}} \leq o(1)$, $\kappa \leq \frac{1}{3}$ is equivalent to $R \leq 3 + o(1)$ at the factor level. Condition (c) is a theorem for the modular tower and an open conjecture for general factors; condition (b) is the archimedean gap of DD. □

6.1.4 The Λ constant is not universal

Theorem 6.4 (A continuum of discriminant-layer constants inside the CM class)**.** *For the Fermat tower, $\Lambda_q/g_N \to c(q) := (q^2+2q+3)/(q^2-1)$ (Theorem 5.61), with $c(3) = 9/4 = 2.25$, $c(5) = 19/12 \approx 1.583333$, $c(7) = 11/8$, $c(11) = 73/60$, $c(13) = 33/28$. Since $c(3) \neq c(5)$, no single constant is universal. Moreover, by the convex hull of Theorem 6.3, assembling F_{3^n} and F_{5^m} factors with weight θ produces the continuum*

$$\bar{c}(\theta) = \theta c(3) + (1-\theta)c(5) \in [19/12,\ 9/4], \qquad \theta \in [0,1].$$

For an explicit abelian-variety product tower with $\bar{\Lambda} = \sum_q \Lambda_q \log q / \sum_q g_q \log q$ and Λ_q the exact closed form of Theorem 5.61, varying n_3/n_5 from 1.0 to 1.8 sweeps w_3 from 0.0005 to 0.9965 and $\bar{\Lambda}$ from 1.5836 to 2.2477; the deviation from the linear interpolation $\bar{c}(w_3)$ decays to $\leq 4.1 \times 10^{-6}$ at $n_5 = 7$.

Proof. Theorem 5.61 gives the exact Λ_q and the limits $c(q)$; exact rational checks of Λ_q/g_q at $n = 3, 7, 12$ give, for $q = 3$, $2.046154, 2.247154, 2.249988$ and for $q = 5$, $1.550747, 1.583279, 1.583333$, converging to 2.25 and 19/12. The inequality $9/4 \neq 19/12$ is exact, so universality already fails within the CM class. By Theorem 6.3 any convex combination of the two conductor shares is realised at the level of isogeny factors, and the conductor-mass layer is additive across the direct sum, giving $\bar{c}(\theta)$. Sample values: $\bar{c}(0) = 1.583333$, $\bar{c}(0.25) = 1.750000$, $\bar{c}(0.5) = 1.916667$, $\bar{c}(0.75) = 2.083333$, $\bar{c}(1) = 2.250000$. For the explicit product tower, computing Λ_q by the closed form of Theorem 5.61 gives $(n_3, n_5) = (7,7) \mapsto \bar{\Lambda} = 1.583633$ with $w_3 = 0.000534$; $(10,7) \mapsto 1.770298$ with $w_3 = 0.280550$; $(13,7) \mapsto 2.247659$ with $w_3 = 0.996495$; the differences from $\bar{c}(w_3)$ are 5.6×10^{-5}, 6.9×10^{-5} and 4.1×10^{-6}, reflecting the $O(1/n_q)$ convergence of Λ_q/g_q. Exact recomputation of $f_p = 2D + R + \Lambda$, $\delta_n = 1$ and $T = D + \frac{1}{2}R$ holds for $q = 3, 5, 7, 11$ and $n = 1, 2, 3, 7, 12$ (all 20 cases). Whether such an abelian-variety product is realised as the

Prym direct summand of a curve covering is Problem 10.12 (4.5), tagged C; the non-universality is already strict at the abelian-variety level. This is the third independent occurrence of the pointwise-discrete / tower-continuous dichotomy of Section 2.1, the first two being the reachable κ-set and the $1/n$ coefficient of R_n. □

6.1.5 What does not generalise

Theorem 6.5 (Classification of Fermat-specific quantities)**.** *The following are consequences of the cyclotomic structure and do not extend to a general tower:*

1. $\delta_n = f_p/g_N - 2n = 1$ *(5.60): a general tower has no closed form* $f_p = (2n+1)g_N$*, so the constant term of the conductor may drift.*
2. $T = D + \frac{1}{2}R$ *and* $\mu(\Phi_c) = \log d_{\Phi_c} + \frac{1}{2}\log p$ *(5.53, 5.59): for a general CM type* Φ_ρ *the finite surplus* $\delta_{\Phi_\rho} = \mu(\Phi_\rho) - \log d_{\Phi_\rho}$ *is type-dependent, not* $\frac{1}{2}\log p$*.*
3. *The exact* $\varepsilon_c = \log p/[p^{c-1}(p-1)]$ *and the closed form for* u_n *(5.54–5.55): general factors admit no layer count, only an* $O(\hat{f}_n^{-1})$ *type remainder.*
4. $r = a + w$*, the coefficient* 3.914997 *for* $p = 3$*, the closed form for* R_n*, and the tightness of Corollary 3.11: these use 00712 Thm 3.4's cyclotomic height formula* C_p*, which has no general analogue.*

Accordingly the upgrade of $\hat{f}_n = (2n+1)\log p$ *in Section 5.8 is tightened to "U relative to the cyclotomic CM input and the anchor"; for a general tower only Theorem 6.1's combinatorial control is available.*

Proof. Each item is positive evidence restricted to the cyclotomic case: items (1)–(3) use the odd-Dirichlet-character conductor distribution of 5.52, which has no counterpart for a general Φ_ρ; item (4) uses $C_p = -6\log A - \frac{p\log p}{4(p^2-1)}$ from 00712 Thm 3.4. The 20-case exact verification of $f_p = 2D + R + \Lambda$, $\delta_n = 1$ and $T = D + \frac{1}{2}R$ confirms these identities inside the cyclotomic class but provides no general-tower identity. Tags are therefore tightened, not withdrawn. □

6.1.6 Ledger update

Theorem 6.6 (Tags updated; conclusions unchanged)**.** *The following entries change, and no mathematical conclusion of Sections 3.2 through NN changes:*

1. *The barrier hypothesis* $\sigma \to 1$ *(5.13) is extended from the Fermat tower with Obus acceleration to the whole class of semistable towers (Theorem 6.1); tag U relative to 5.10.*
2. *The reduction of* $\kappa \le \frac{1}{3}$ *is upgraded from 3.17 under a BMY hypothesis to Theorem 6.3: condition (a) universal, (b) and (c) factor-level.*
3. *The constant* Λ/g_N *is shown* non-universal *(Theorem 6.4); Section 7 of NN's question is answered in the negative, and the presumed dependence of* $\sigma \to 1$ *on that constant is withdrawn:* Λ *and* $\operatorname{rank} L$ *are different-order remainders.*
4. *The tag of* $\hat{f}_n = (2n+1)\log p$ *is tightened to "relative to cyclotomic CM plus anchor" (Theorem 6.5).*
5. *The dictionary* $R = 12\kappa - 1 + o(1)$ *is extended to general towers with the archimedean term still tagged C (Corollary 6.2).*

No item adds a new external reference; the remaining C layer is unchanged. The fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$*,* $E(c) = 12/(5-18c)$*,* $47/3$*, the pointwise closure* $C \ge 36/7 > 3$*, Routes 1–3, and the status of Theorem 7.5/9.2 as Th.A* are unchanged; abc remains open.*

Proof. Each item is a direct substitution of Theorems 6.1–6.5 and Corollary 6.2; none uses the truth value of Theorem A and none changes an exponent, a constant or a No-Go count. All tables were recomputed with exact rational arithmetic (more than 50 exact checks) and double-precision cross-checks. □

Remark 6.7 (Honest boundary). Theorems 6.1–6.5, 6.6 and Corollary 6.2 are unconditional relative to standard cited input and introduce no new external reference. Text-level, not re-proved here (C): the Saito–Liu identity, especially its $\Delta_{\mathrm{Del,q}} - f_q = O(g)$ form for non-semistable models; Wilms' normalised Noether formula; the isogeny invariance and $\overline{\mathbb{Q}}$-decomposition of 3.17; the archimedean bound (b) for general towers (Jorgenson–Kramer covers the modular tower); the modular-tower BMY input for (c) (Dolce–Mercuri, Mayer); and the curve-level realisation of product abelian varieties of Problem 10.12. Hence $\sigma \to 1$ is "U relative to 5.10", not absolutely unconditional, and Theorem 6.3 must not be abbreviated to "$\kappa \le \frac{1}{3}$ is proved for general towers": conditions (b) and (c) remain. Not settled: an independent verification of the general Saito–Liu form, R_ρ for non-automorphic simple factors, the Prym direct-summand question, a general bound on $\varepsilon_{\mathrm{ar}}$, and an independent verification of Theorem A. This appendix verifies no cited proof independently.

6.2 The next tower-side step: single-curve realisation and the mixed-prime reduction

Section 6.1 left three candidate continuations: drafting the referee report (N1), Problem 10.12 on single-curve realisation together with the curve-level status of the continuum of Section 6.1.3 (N2), and the non-semistable Saito–Liu form (N3). We carry out N2. The outcome is a three-way result: the lower half of Problem 10.12 is closed by a density theorem; the Torelli dimension shortcut is closed by an explicit codimension count; and the upper half of 10.12 is shown to be the *same* gap as the curve realisation of the Λ continuum, namely the stable model of mixed-prime wild three-point covers. Nothing in Sections 2.1 through OO changes.

6.2.1 A density theorem for single curves

Theorem 6.8 (Normalisation of the Fermat height)**.** *Let $A = 1.2824271291\ldots$ be the Glaisher–Kinkelin constant and $C_p = -6\log A - \frac{p\log p}{4(p^2-1)} < 0$ the constant of Theorem 3.4 of 00712 (used in 3.11, 5.44 and 5.56). Define*

$$\alpha_p := -\frac{C_p}{2\log p} = \frac{3\log A}{\log p} + \frac{p}{8(p^2-1)} > 0.$$

Then α_p is strictly decreasing, $\alpha_p \to 0$, and the Fermat height satisfies

$$\kappa(F_{p^m}) = \frac{1}{8} - \frac{\alpha_p}{m} + O(p^{-m}), \qquad \alpha_3 = 0.72615325,\ \alpha_5 = 0.48972121,\ \alpha_{293} = 0.13180704.$$

Proof. By 00712 Thm 3.4, $\kappa(F_{p^n}) = \frac{1}{8} + \frac{C_p}{2n\log p} + O(p^{-n})$; substituting $C_p = -2\alpha_p \log p$ gives the displayed form. Both terms in α_p are positive and strictly decreasing in $p \ge 3$ (the second is $\frac{p}{8(p^2-1)} = \frac{1}{8(p-1/p)}$), so α_p is strictly decreasing, and $3\log A/\log p \to 0$ while $p/(8(p^2-1)) \to 0$. The numerical values follow by direct substitution of $3\log A = 0.7462634310\ldots$. □

Lemma 6.9 (Relative density of the constants α_p)**.** *For consecutive odd primes $p_k < p_{k+1}$, $\alpha_{p_k}/\alpha_{p_{k+1}} \to 1$ as $k \to \infty$.*

Proof. The prime number theorem gives $p_{k+1}/p_k \to 1$, hence $\log p_{k+1}/\log p_k \to 1$. Since $\alpha_p = \frac{3\log A}{\log p}(1 + O(1/(p\log p)))$, the ratio of α's tends to $\log p_{k+1}/\log p_k \to 1$. Numerically the ratio is 1.006557 at $p = 127$ and 1.008340 at $p = 293$; the value 1.482789 at $p = 3$ is a low-end exception that does not affect the limit. □

Theorem 6.10 (Single-curve reachable κ is dense in $(0, \frac{1}{8}]$)**.** *Let $\mathcal{K}_{\text{curve}}$ denote the set of $\kappa(F_{p^m})$ with p an odd prime, $m \geq 1$ and $\kappa(F_{p^m}) > 0$. Then $\mathcal{K}_{\text{curve}}$ is dense in $(0, \frac{1}{8}]$, and $\frac{1}{8}$ is a limit point. Each element is realised by the single curve F_{p^m} itself (no product or convex combination).*

Proof. Fix $0 < a < b < \frac{1}{8}$ and put $x = \frac{1}{8} - \kappa^*$; as κ^* runs over (a, b), x runs over $(\frac{1}{8} - b, \frac{1}{8} - a) \subset (0, \frac{1}{8})$. By Theorem 6.8 we need $\alpha_p/m = x$ up to the exponentially small remainder $O(p^{-m})$. By Lemma 6.9, for every $\varepsilon > 0$ there is $P(\varepsilon)$ such that for consecutive primes $p > P$ every relative interval $(u, u(1+\varepsilon))$ inside $(0, \alpha_P)$ contains some α_p. Choose $m = \lceil \alpha_P/(2x) \rceil$, so that $mx \in [\alpha_P/2, \alpha_P/2 + x) \subset (0, \alpha_P)$; then the interval $(mx, mx(1+\varepsilon))$ contains some α_p, giving $|\alpha_p/m - x| \leq \varepsilon x$. Taking $\varepsilon \to 0$ (and hence $P \to \infty$, which also forces the $O(p^{-m})$ remainder to be exponentially small) yields density. Finally $m \to \infty$ at fixed p gives $\kappa(F_{p^m}) \to \frac{1}{8}$ from below because $C_p < 0$, so $\frac{1}{8}$ is a limit point; it is attained by the Fermat tower itself. Exact numerical evidence: with $m \leq 3000$, the largest gap in $[10^{-5}, 1/8)$ of the point set is 8.0573×10^{-3} for $p_{\max} \leq 293$, 2.7836×10^{-3} for $p_{\max} \leq 2003$, 1.9641×10^{-3} for $p_{\max} \leq 10007$, and 7.8885×10^{-4} for $p_{\max} \leq 19997$, the gap location migrating from $(0.003974, 0.012032)$ to $(0.054195, 0.054984)$. The convergence is step-like and governed by the relative gaps of $\{\alpha_p\}$, of order $1/(p \log^2 p)$, *not* by $1/(m \log p)$: increasing $m_{\max}$ from 200 to 3000 at fixed $p_{\max}$ leaves the maximal gap unchanged, because the low-κ gap is produced by the sparseness of $\{\alpha_p\}$ near α_3, where $\alpha_3 - \alpha_5 = 0.236$. An initial guess of a $O(1/(m \log p))$ rate is therefore withdrawn; the density conclusion is unaffected. □

6.2.2 What remains of Problem 10.12

Theorem 6.11 (Exact residual of the single-curve problem)**.** *With the universal dictionary $R = 12\kappa - 1 + O(\hat{f}^{-1})$ of Corollary 6.2, the five anchors are*

$$\tfrac{1}{12} \mapsto 0, \quad \tfrac{1}{8} \mapsto \tfrac{1}{2}, \quad \tfrac{1}{6} \mapsto 1, \quad \tfrac{5}{24} \mapsto \tfrac{3}{2}, \quad \tfrac{1}{3} \mapsto 3.$$

Consequently:

1. *On $(0, \frac{1}{8}]$ the single-curve realisation problem is closed by Theorem 6.10: the Fermat subcovers F_{p^m} give a countable dense set of single curves, all with $\kappa < \frac{1}{8}$ and converging upwards to $\frac{1}{8}$.*
2. *On $(\frac{1}{8}, \frac{1}{3})$ the realised single-curve values are only $\{\frac{1}{6}, \frac{5}{24}, \frac{1}{3}\}$, where $\frac{1}{6}$ is the function-field Mason value (tagged C on the number-field side); all other points remain open.*

Moreover the continuum of Section 6.1.3 (the Λ-constant interval $[19/12, 9/4]$) arises from isogeny weights of an abelian-variety product, whereas Theorem 6.10's dense set arises from the layer index m of a single curve: the two continuities have different sources. This is a fourth independent occurrence of the pointwise-discrete / tower-continuous dichotomy of Section 2.1, now inside *the tower side.*

Proof. Substituting the five anchor values into $R = 12\kappa - 1$ gives $0, \frac{1}{2}, 1, \frac{3}{2}, 3$; this is the "one rule, not seven independent data" line of §2.4 of 00544. Item (1) is Theorem 6.10 plus $C_p < 0$. Item (2) is the verbatim anchor table of 00544 together with 6.5; no further single-curve value in $(\frac{1}{8}, \frac{1}{3})$ appears in the audited record. The two sources of continuity are distinct by construction: Theorem 6.4 varies the convex weight θ of two isogeny factors, Theorem 6.10 varies m at fixed p (and p at fixed m). □

6.2.3 Torelli dimension count: the shortcut is closed

Theorem 6.12 (Coarse dimension counting cannot exclude product realisation)**.** *For principally polarised abelian varieties of dimension g,*

$$\dim \mathcal{A}_g = \frac{g(g+1)}{2}, \qquad \dim \mathcal{M}_g = 3g - 3, \qquad \operatorname{codim}_{\mathcal{A}_g} \mathcal{J}_g = \frac{(g-1)(g-2)}{2}.$$

A family assembled from $k \le g$ simple CM or automorphic factors spans a Shimura-type locus of dimension $O(k) = O(g)$, of the same order as $\dim \mathcal{J}_g = 3g-3 = O(g)$. Hence no coarse dimension argument can exclude the curve realisation of the abelian-variety products of Theorem 6.4; Problem 10.12 requires fine structure.

Proof. The dimensions are classical: $\mathcal{A}_g$ has dimension $g(g+1)/2$, $\mathcal{M}_g$ has dimension $3g-3$ for $g \ge 2$, and the Torelli locus $\mathcal{J}_g$ has dimension $3g-3$, so the codimension is $\frac{g(g+1)}{2} - (3g-3) = \frac{(g-1)(g-2)}{2}$. A CM factor contributes a zero-dimensional isogeny class and a GL_2-type automorphic factor a one-parameter Hecke family, so k simple factors give at most $O(k) \le O(g)$ moduli. Both $O(g)$ quantities are the same order, so there is no numerical contradiction. Explicit table: at $g = 6, 10, 100, 1000, 10^4$ the codimensions are $10, 36, 4851, 498501, 49985001$ while $\dim \mathcal{J}_g = 15, 27, 297, 2997, 29997$. Thus the tempting shortcut "the Torelli locus is too small" is closed, and the question is pushed to the mixed-prime geometry below. □

6.2.4 The mixed-prime reduction

Theorem 6.13 (The two open items are one gap)**.** *1. By 4.3 and 4.6, a curve-cover tower with fixed bad set S and a single dominant wild prime has Prym weight $w_P = \mathrm{Sw}_P/(\mathrm{Sw}_P + \mathrm{Sw}_Y) \to 1$, i.e. one dominant factor, while the number of branch points obeys $r_n \le R(S)$. The continuum of Theorem 6.4 needs two* comparable *weights, hence requires $|S| \ge 2$, $\mathrm{Sw}_{q_1} \asymp \mathrm{Sw}_{q_2}$, and a Galois group carrying higher wild subgroups at two distinct primes.*

2. *With $S = \{3,5\}$ and $\mathrm{Sw}_p \sim 2n_p g_{p^n}$, the weight $w_3 = \mathrm{Sw}_3/(\mathrm{Sw}_3 + \mathrm{Sw}_5)$ is indeed continuously tunable: for $(n_3, n_5) = (2,2), (3,2), (4,2), (4,3), (5,3), (6,4), (7,5), (8,6)$ it takes $0.0648, 0.5466, 0.9399, 0.2739, 0.8131, 0.5823, 0.3187, 0.1383$. Hence there is no combinatorial obstruction to curve realisation.*

3. *Obus (arXiv:0911.1103v3) describes the stable model only for three-point $\mathbb{Z}/p^n$-covers at a* single *prime; the mixed case needs $(\mathbb{Z}/3^n \times \mathbb{Z}/5^m)^2$-type groups, for which the dual graphs Γ_3, Γ_5, the ranks $\mathrm{rank}\, L(\Gamma_q)$, the conductors f_q and the discriminant layers Λ_q have no cited source. Therefore*

$$\begin{gathered} \textit{Problem 10.12 on } (\tfrac{1}{8}, \tfrac{1}{3}) \wedge \\ \textit{curve realisation of Theorem 6.3} \iff \\ \textit{stable model of mixed wild three-point covers,} \end{gathered}$$

a single concrete literature-level task (tag C), not a new theorem.

Proof. (1) is 4.3 together with the Evertse-type bound 4.6 and the convex-hull weights of 3.17; comparable weights $w, 1-w$ bounded away from 0 force two independent Swan depths. (2) is a direct substitution using $g_{p^n} = (p^n-1)(p^n-2)/2$; the eight displayed values were computed in exact integer arithmetic before taking the ratio, and they sweep $(0,1)$ without a gap, so no exclusion follows at the combinatorial level. (3) is the scope of Obus Prop. A.3/A.5 as recorded in 5.28: the chain of $n+1$ components is specific to $\mathbb{Z}/p^n$. The augmented dual graph for a group with two prime-to-p factors, and the resulting $m_q - 1 = \mathrm{rank}\, L(\Gamma_q)$ entering 5.10, are not supplied by that reference. By 5.10 and 5.61 these are exactly the inputs needed for $\sigma_n - 1$ and for Λ_q, so both open items reduce to the same missing model. This refines 's classification, where Problem 10.12 was listed as a separate "long-term geometric" item: the obstruction is not a new theorem but the multi-prime extension of an existing single-prime stable model. □

Remark 6.14 (Honest boundary). Theorems 6.8–6.13 and Lemma 6.9 are unconditional relative to standard cited input and introduce no new external reference. Text-level, not re-proved here (*C*): Theorem 3.4 of 00712 and its $O(p^{-n})$ remainder, the prime number theorem in the form

$p_{k+1}/p_k \to 1$, the Obus single-prime stable model, 4.3 and 4.6 and Evertse's bound, and the estimate $\mathrm{Sw}_p \sim 2n_p g_{p^n}$. Hence Theorem 6.10 is "U relative to 00712 plus the prime number theorem", not absolutely unconditional; the $\kappa(F_{p^m})$ values used are the asymptotic closed forms, and the $O(p^{-m})$ perturbations at small m were not measured pointwise — density uses only $m \to \infty$ or $p \to \infty$, where the perturbation is exponentially small. Not settled: the single-curve realisation of $(\frac{1}{8}, \frac{1}{3}) \setminus \{\frac{1}{6}, \frac{5}{24}, \frac{1}{3}\}$, the mixed-prime stable model, the number-field realisation of $\frac{1}{6}$, a general upper bound on $\varepsilon_{\mathrm{ar}}$, and an independent verification of Theorem A. No conclusion of Sections 2.1–6.1 and none of the fourteen No-Go reasons changes; $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5 - 18c)$, $47/3$, $C \geq 36/7 > 3$ and abc (still open) are unchanged. This appendix verifies no cited proof independently.

6.3 The mixed-prime model: decoupling wild inertia

Section 6.2.6 reduced the upper half of Problem 10.12 and the curve realisation of the Λ continuum of Section 6.1.3 to a single gap: the stable model of mixed-prime wild three-point covers, judged to require a multi-prime extension of Obus. We show that this is only half true. The geometric stable model is still missing, but every invariant needed for $\hat{f}$, the weights, $\sigma \to 1$ and the continuum can be obtained *without* it, by wild-inertia decoupling, Chevalley–Weil and Riemann–Hurwitz. A second, sharper finding is that a single mixed-prime curve and an abelian-variety product give *different* weight parametrisations. Nothing in Sections 2.1 through PP changes.

6.3.1 The construction

Let $p \neq q$ be odd primes, $a = p^n$, $b = q^m$, and let $K \supseteq \mathbb{Q}(\zeta_a, \zeta_b)$. Put

$$C_{a,b}: \; u^a + v^b = 1, \qquad \pi : C_{a,b} \to \mathbb{P}^1_t, \quad (u, v) \mapsto t = u^a = 1 - v^b,$$

on which $G = \mu_a \times \mu_b \cong \mathbb{Z}/a \oplus \mathbb{Z}/b$ acts by $(u, v) \mapsto (\zeta_a u, \zeta_b v)$, trivially on t. The covering is branched at exactly three points: $t = 0$ with inertia μ_a (b points of index a), $t = 1$ with inertia μ_b (a points of index b), and $t = \infty$ with inertia G and index ab (one point, since $\gcd(a, b) = 1$).

Theorem 6.15 (Genus and the H^1 character support)**.** *For $C_{a,b}$ one has*

$$g = \frac{(a-1)(b-1)}{2}, \qquad \dim H^1(C_{a,b})_{(i,j)} = 1 \;\; \textit{for all } 1 \leq i \leq a-1, \; 1 \leq j \leq b-1,$$

and $\dim H^1_{(i,j)} = 0$ whenever $i = 0$ or $j = 0$.

Proof. Riemann–Hurwitz with $|G| = ab$ and indices (a, b, ab) gives $2g - 2 = ab\big(-2 + (1 - \frac{1}{a}) + (1 - \frac{1}{b}) + (1 - \frac{1}{ab})\big) = ab - a - b - 1 = (a-1)(b-1) - 2$, hence $g = (a-1)(b-1)/2$. For the character decomposition, $H^{1,0}$ has basis $\omega_{i,j} = u^{i-1} v^{j-1}\, du$ for $1 \leq i \leq a-1$, $1 \leq j \leq b-1$ with $\frac{i}{a} + \frac{j}{b} < 1$; the G-action sends $\omega_{i,j}$ to $\zeta_a^i \zeta_b^j \omega_{i,j}$. The conjugate space $H^{0,1} = \overline{H^{1,0}}$ realises the characters $(-i, -j) = (a - i, b - j)$, for which $\frac{a-i}{a} + \frac{b-j}{b} = 2 - (\frac{i}{a} + \frac{j}{b}) > 1$. Because $\gcd(a, b) = 1$, the boundary $\frac{i}{a} + \frac{j}{b} = 1$ has no solution: $ib + ja = ab$ forces $a \mid ib$, hence $a \mid i$, contradicting $1 \leq i \leq a - 1$. Therefore every pair (i, j) with both entries non-trivial appears exactly once, for a total of $(a-1)(b-1) = 2g$ characters, matching $\dim H^1 = 2g$. Characters with $i = 0$ or $j = 0$ occur in neither half. □

6.3.2 Wild-inertia decoupling

Theorem 6.16 (Decoupling)**.** *Let $\mathfrak{p} \mid p$ and $\mathfrak{q} \mid q$, let $I_{\mathfrak{p}}$ be the inertia group with wild subgroup $P_{\mathfrak{p}}$ (pro-p), and let $\chi_{(i,j)} = \chi_a^i \otimes \chi_b^j$ be a character of G. Then*

$$f_{\mathfrak{p}}(\chi_{(i,j)}) = f_{\mathfrak{p}}(\chi_a^i), \qquad f_{\mathfrak{q}}(\chi_{(i,j)}) = f_{\mathfrak{q}}(\chi_b^j),$$

where f denotes the local Artin conductor. In particular the p-adic conductor is independent of j, and the q-adic conductor is independent of i.

Proof. Work with ℓ-adic cohomology, $\ell \neq p, q$. By Theorem 6.15, $V_{\chi_{(i,j)}}$ is one-dimensional, so for the lower ramification groups G_k one has $V^{G_k} = V$ if and only if $\chi_{(i,j)}|_{G_k} = 1$. For $k \geq 1$, G_k is a pro-p group, while the image of χ_b^j has order dividing $b = q^m$, coprime to p; hence $\chi_b^j|_{G_k} = 1$. For $k = 0$, since $\zeta_b \in K$ the inertia group acts trivially on μ_b, so again $\chi_b^j|_{G_0} = 1$. Consequently $\chi_{(i,j)}|_{G_k} = \chi_a^i|_{G_k}$ for every $k \geq 0$, and the Artin conductor

$$f_{\mathfrak{p}}(\chi) = \sum_{k \geq 0} \frac{|G_k|}{|G_0|} \dim\bigl(V/V^{G_k}\bigr)$$

coincides term by term with that of χ_a^i. The $\mathfrak{q}$ statement is symmetric, using that G_k is pro-q there and that $\zeta_a \in K$. □

Corollary 6.17 (Exact mixed conductor decomposition)**.** *Write* $F_p^{(a)} := \sum_{i=1}^{a-1} f_{\mathfrak{p}}(\chi_a^i)$ *and* $F_q^{(b)} := \sum_{j=1}^{b-1} f_{\mathfrak{q}}(\chi_b^j)$*; these are pure single-prime data, independent of* b *and* a *respectively. Then*

$$\log \mathrm{cond}_{\mathfrak{p}}(C_{a,b}) = (b-1)F_p^{(a)} \log p, \qquad \log \mathrm{cond}_{\mathfrak{q}}(C_{a,b}) = (a-1)F_q^{(b)} \log q.$$

Proof. Sum $f_{\mathfrak{p}}(\chi_{(i,j)}) = f_{\mathfrak{p}}(\chi_a^i)$ of Theorem 6.16 over all $(a-1)(b-1)$ characters of Theorem 6.15: for each fixed i the index j takes $b-1$ values and contributes the same conductor, giving $(b-1)\sum_i f_{\mathfrak{p}}(\chi_a^i)$. Multiply by $\log p$ to pass from the exponent to the logarithmic conductor; the $\mathfrak{q}$ case is symmetric. □

6.3.3 Density decomposition and the tower axioms

Theorem 6.18 (b-independence and additivity)**.** *With* $g = (a-1)(b-1)/2$ *of Theorem 6.15,*

$$\boxed{\hat{f}_p = \frac{2F_p^{(a)} \log p}{a-1}, \qquad \hat{f}_q = \frac{2F_q^{(b)} \log q}{b-1}, \qquad \hat{f}_{\mathrm{tot}} = \hat{f}_p + \hat{f}_q}.$$

In particular $\hat{f}_p$ *is exactly independent of* $b = q^m$*, and* $\hat{f}_q$ *is exactly independent of* $a = p^n$*. The tower axioms hold: (T1)* $g \to \infty$*, (T2)* $S = \{p, q\}$ *fixed, (T3)* $\hat{f}_{\mathrm{tot}} \to \infty$ *driven* linearly *by* $n, m \to \infty$ *rather than by the exponentials* p^n, q^m*.*

Proof. Divide Corollary 6.17 by $g = (a-1)(b-1)/2$: the factor $(b-1)$ cancels against the $(b-1)$ in g, leaving $2F_p^{(a)} \log p/(a-1)$, and similarly at $\mathfrak{q}$. Additivity is the additivity of log cond. For (T1) $g \sim ab/2 \to \infty$; (T2) is built into the construction; (T3) follows from $F_p^{(a)}/a \to \infty$ linearly in n (Theorem 6.19 below, or directly from any standard conductor distribution). Exact numerical check: fixing $(p, n) = (3, 4)$ and sweeping $q \in \{5, 7, 11\}$, $m = 1, 2, 3, 4, 8$, the genus ranges from 160 to 8.57×10^9 while $\hat{f}_p = 7.8001472495$ in all 15 cases, with maximal residual 8.88×10^{-16} (pure floating-point noise). □

Theorem 6.19 (Explicit single-prime sum, tag A)**.** *Under the Kummer conductor distribution* $f_{\mathfrak{p}}(\chi_a^i) = n - v_p(i)$ *(standard local class field theory, tag A),*

$$F_p^{(a)} = (p-1)\sum_{k=1}^{n} k\, p^{k-1} = \frac{p^n(np - n - 1) + 1}{p-1}, \qquad \frac{F_p^{(a)}}{a} = n - \frac{n+1}{p} + O(p^{-n}),$$

so $\hat{f}_p \sim 2n \log p$*, of the same order as the Fermat tower value* $\hat{f}_n = (2n+1)\log p$ *of Theorem 5.60.*

Proof. Characters with $v_p(i) = t$ have conductor exponent $n - t$, and there are $p^{n-t} - p^{n-t-1} = p^{n-t-1}(p-1)$ of them; setting $k = n - t$ gives $(p-1)\sum_{k=1}^{n} kp^{k-1}$, whose closed form is $\bigl(np^{n+1} - (n+1)p^n + 1\bigr)/(p-1)^2$, i.e. $\bigl(p^n(np - n - 1) + 1\bigr)/(p-1)$ after multiplying by $p - 1$. The quotient $F_p^{(a)}/a$ and the asymptotic $\hat{f}_p \sim 2n \log p$ follow immediately. All 15 pairs (p, n) with $p \in \{3, 5, 7, 11\}$ and $n \in \{1, \ldots, 5\}$ were checked by direct integer summation against the closed form. This is an independent order-of-magnitude cross-check of Theorem 5.60. Tag A: replacing the Kummer distribution changes the numerical value of $\hat{f}_p$ but not Theorems 6.16–6.18. □

6.3.4 Weights: a single curve is not an AV product

Theorem 6.20 (Two distinct weight parametrisations)**.** *For the single mixed curve,*

$$w_p = \frac{\hat{f}_p}{\hat{f}_p + \hat{f}_q}, \qquad \text{model } A: \quad w_p \approx \frac{n \log p}{n \log p + m \log q}\Big(1 + O\big(\tfrac{1}{n} + \tfrac{1}{m}\big)\Big),$$

a linear-fractional *function of* (n, m)*. By contrast the abelian-variety product of Section 6.2.5 uses* $w_3^{\mathrm{AV}} = \frac{n g_{3^n} \log 3}{n g_{3^n} \log 3 + m g_{5^m} \log 5}$ *with* $g_{p^n} = (p^n - 1)(p^n - 2)/2$*, which is* exponentially *sensitive. For the same* $(3^n, 5^m)$ *the two differ by up to* 0.4203*. Consequently the* Λ *continuum* $\bar{c}(w) = w\, c(p) + (1-w)c(q)$ *with* $c(p) = (p^2 + 2p + 3)/(p^2 - 1)$ *is realised on the single curves* $C_{3^n,5^m}$ *with* (n, m) *as parameter, e.g.* $(1, 50) \mapsto \bar{c} = 1.5924$ *and* $(50, 1) \mapsto \bar{c} = 2.2308$*, sweeping* $(19/12, 9/4)$*.*

Proof. The formula for w_p is the ratio of the two additive terms of Theorem 6.18; inserting $F_p^{(a)}/a \sim n$ and $F_q^{(b)}/b \sim m$ from Theorem 6.19 gives the linear-fractional approximation. The AV-product weight is the one computed in Theorem 6.13 from independent genera g_{3^n}, g_{5^m}. Sample comparisons: $(n, m) = (2, 2)$ gives 0.3945 versus 0.0648; $(4, 2)$ gives 0.5693 versus 0.9399; $(5, 5)$ gives 0.3937 versus 0.0041; $(7, 6)$ gives 0.4357 versus 0.0153, the largest gap 0.4203. Since w_p is continuously tunable in $(0, 1)$ by the 7×7 table of Theorem 6.18, $\bar{c}(w)$ sweeps the interval between $c(3) = 9/4$ and $c(5) = 19/12$. The *structure* of the continuum is unconditional; the exact value of $\Lambda^{(a)}$ for a pure μ_{p^n} cover (Theorem 5.61 was derived for the $(\mathbb{Z}/p^N)^2$ Fermat group) is tagged C. □

6.3.5 $\sigma \to 1$ and the stable-model question

Theorem 6.21 (6.1 holds on mixed towers; the geometric model is dispensable)**.** *1. (U) By the purely combinatorial bound* $\operatorname{rank} L \le 2g - 3$ *of Theorem 5.17 and 5.10, with* $S = \{p, q\}$ *fixed,*

$$\sigma_{n,m} - 1 \le \frac{2(\log p + \log q)}{\hat{f}_{\mathrm{tot}}}\big(1 + O(g^{-1})\big) \longrightarrow 0.$$

Numerically $(n, m) = (2, 2), (6, 5), (20, 15), (50, 40)$ *give* $0.5557, 0.1977, 0.05996, 0.02288$*. Hence the barrier hypothesis of Theorem 6.1 holds on mixed towers without any knowledge of* Γ_p, Γ_q*.*

2. *(C, structural)* $C_{a,b}/\mu_b \cong \mathbb{P}^1$ *via* $u \mapsto u^a$ *is a two-point* μ_a*-cover; since* $K \ni \zeta_a$ *its p-adic semistable model has a single-vertex dual graph (no cluster subdivision). The cover* $C_{a,b} \to C_{a,b}/\mu_b$ *is tame at* $\mathfrak{p}$ *because* $q \neq p$*, and Liu's tame semistable base change preserves tree-likeness; hence* $\operatorname{rank} L(\Gamma_p) = 0$ *and 5.10 would give* $\Delta_{\mathrm{Del,p}} = f_p$ exactly*, i.e.* $\sigma_{n,m} = 1$ *rather than* $1 + o(1)$*. The contrast with the Obus chain of the pure Fermat tower is explained by* $t = 1$*: there the inertia is* μ_{p^n} *(wild at p), here it is* μ_b *(tame at p).*

3. *Therefore 6.13's judgement is refined: the mixed stable model is still missing as a* geometric description*, but every invariant used by* $\hat{f}$*,* w_p*,* $\sigma \to 1$ *and the continuum is supplied by Theorems 6.16–6.20. The remaining C items are the explicit* $|V|, |E|, \operatorname{rank} L$ *of* Γ_p *and the exact constant* $\Lambda^{(a)}$ *of a pure* μ_{p^n} *cover. The upper half of Problem 10.12 is* not *closed:* κ *still needs the averaged Colmez height for the mixed CM type (C), and the realised values remain those of Theorem 6.11.*

Proof. (1) is Theorem 6.1 with the explicit $\hat{f}_{\mathrm{tot}}$ of Theorem 6.18; the numerical table is a direct substitution. (2) is the scope of Obus Prop. A.3/A.5 (single prime $\mathbb{Z}/p^n$, three fully ramified points) together with Liu's tame base-change theorem; both are cited at text level and not re-proved here, hence the tag C. (3) summarises Theorems 6.16–6.20 and records that Theorem 6.11's list of realised $\kappa \in (\frac{1}{8}, \frac{1}{3})$ is unchanged, since no height computation is performed in this appendix. □

Remark 6.22 (Honest boundary). Theorems 6.15–6.18, 6.20, 6.21(1) and Corollaries 6.17 are unconditional relative to standard Artin-conductor theory; Theorem 6.19 carries tag A (the Kummer conductor distribution), Theorem 6.21(2) carries tag C (Liu's tame base change and the two-point Kummer semistable model), and the exact $\Lambda^{(a)}$ constant and the mixed CM-type Colmez height remain C. In particular Theorem 6.21's second item is a structural inference, not a proved theorem of this appendix, and $\sigma = 1$ exactly is not claimed unconditionally. Fifteen integer checks of the closed form, fifteen b-independence checks, a 7×7 weight table, twelve $\bar{c}$ values and an eight-row σ table were recomputed exactly or in double precision. No new external reference is introduced and no cited proof is independently verified. The fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, Routes 1–3 and the status of Theorem 7.5/9.2 as $Th.A^*$ are unchanged; *abc* remains open.

6.4 Exact constants of the mixed-prime model

Section 6.3.6 used the Fermat-group constant $c(p) = (p^2+2p+3)/(p^2-1)$ of Theorem 5.61 as the endpoint of a continuum realised on the single mixed curve $C_{3^n,5^m}$, and listed "transporting the 5.61 derivation to a pure μ_{p^n} cover" as the next finite computation. We carry it out. The answer is negative in form and cleaner than expected: the mixed model carries *no* non-zero discriminant layer, and its own exact constant is $\lambda_p = -2/(p-1)$. We also remove tag A from Theorem 6.19. Nothing in Sections 2.1 through QQ changes; fourteen No-Go reasons remain intact.

6.4.1 The exact identity

Theorem 6.23 (Exact conductor-density identity)**.** *Let $a = p^n$ and $F_p^{(a)} = \sum_{i=1}^{a-1} f_{\mathfrak{p}}(\chi_a^i)$. Then*

$$F_p^{(a)} = \frac{p^n\big(n(p-1)-1\big)+1}{p-1}, \qquad \frac{2F_p^{(a)}}{a-1} = \frac{2an}{a-1} - \frac{2}{p-1},$$

$$\boxed{\hat{f}_p = \Big(2n + \frac{2n}{p^n-1} - \frac{2}{p-1}\Big)\log p},$$

with no $O(p^{-n})$ remainder. Consequently tag A of Theorem 6.19 is removed.

Proof. The elementary identity $a\big(n(p-1)-1\big)+1 = an(p-1)-(a-1)$, divided by $(a-1)$ and multiplied by $2/(p-1)$, gives $\frac{2F_p^{(a)}}{a-1} = \frac{2an}{a-1} - \frac{2}{p-1}$; and $\frac{2an}{a-1} = 2n + \frac{2n}{a-1}$. Sixteen pairs (p,n) with $p \in \{3,5,7,11\}$, $n \in \{1,\dots,4\}$ were checked by direct integer summation against the closed form and by exact rational arithmetic on $\frac{2F}{a-1}$: all differences are 0. For instance $p=3, n=4$ gives $F = 284$, $\frac{2F}{80} = 7.1$ and $\hat{f}_p = 7.1\log 3 = 7.8001472495$, matching the fifteen b-independence values of Theorem 6.18 verbatim.

For the removal of tag A: the character $\chi_{p^n}^i$ has kernel $\mu_{p^{n-v_p(i)}}$, hence factors through $\mathrm{Gal}\big(\mathbb{Q}_p(\zeta_{p^{n-v_p(i)}})/\mathbb{Q}_p\big)$; a cyclotomic character of level m has Artin conductor exponent exactly m. Therefore $f_{\mathfrak{p}}(\chi_{p^n}^i) = n - v_p(i)$ is a standard fact of local class field theory, not an adjustable input, and $F_p^{(a)}$ is uniquely determined. It remains U *relative to standard local class field theory*: a non-Kummer group would change $F_p^{(a)}$, but no freedom exists inside the model of Section 6.3. □

Theorem 6.24 (The exact constant)**.**

$$\lambda_p := \lim_{n\to\infty}\Big(\frac{\hat{f}_p}{\log p} - 2n\Big) = -\frac{2}{p-1},$$

with finite-n correction $2n/(p^n-1)$, also exact. Hence $\lambda_3 = -1$, $\lambda_5 = -\frac{1}{2}$, $\lambda_7 = -\frac{1}{3}$, $\lambda_{11} = -\frac{1}{5}$, $\lambda_{17} = -\frac{1}{8}$, $\lambda_{31} = -\frac{1}{15}$.

Proof. Take $n \to \infty$ in Theorem 6.23; the term $2n/(p^n-1) \to 0$ exponentially. Numerically $p=3, n=10$ gives $\hat{f}_p/\log 3 = 19.00033871 = 20-1+0.00033871$ and $p=5, n=10$ gives 19.50000205. For comparison, the Fermat tower of Theorem 5.60 has $\hat{f}_n = (2n+1)\log p$, i.e. the same slope $2\log p$ but constant $+1$ instead of $-2/(p-1)$. The two objects are therefore distinct at the level of the constant term, which is the first signal that the 5.61 layer is not transported. □

6.4.2 The discriminant layer vanishes

Theorem 6.25 ($\Lambda^{(a)}=0$)**.** *Define the discriminant layer of the mixed model by*

$$\Lambda^{(a)} := \log \mathrm{cond}_{\mathfrak{p}}(C_{a,b}) - \sum_{\chi \in H^1(C_{a,b})} f_{\mathfrak{p}}(\chi)\log p.$$

Then $\Lambda^{(a)}=0$ exactly, for all p,n,q,m. Equivalently $c'(p)=0$.

Proof. For a smooth proper curve, $\mathrm{cond}(X)=\mathrm{cond}(H^1(X))$ (H^0 and H^2 are trivial). By Theorem 6.15, $H^1(C_{a,b}) = \bigoplus_{1\le i<a,\,1\le j<b} V_{\chi_{(i,j)}}$, with every biply-nontrivial character occurring exactly once, with no parity screening and no omissions. The Artin conductor is additive on direct sums, and Theorem 6.16 gives $f_{\mathfrak{p}}(\chi_{(i,j)}) = f_{\mathfrak{p}}(\chi_a^i)$. Hence $\sum_\chi f_{\mathfrak{p}}(\chi)\log p = (b-1)F_p^{(a)}\log p$, which is precisely $\log \mathrm{cond}_{\mathfrak{p}}(C_{a,b})$ by Corollary 6.17. The decomposition of 6.17 is therefore the *definition*, leaving no residual layer.

Why Theorem 5.61 is non-zero: there $\Lambda = f_p - 2T$ with $T=\sum_c M_c S_c^-$ the mass of the *CM type* (odd characters only); $2T$ is the extrapolation that even-character mass equals odd-character mass. The orbit decomposition of the Fermat group $(\mathbb{Z}/p^n)^2$ separates characters by $\chi(-1)=\pm 1$, and the two masses differ, producing $\Lambda/g_N \to c(p)$. In the mixed model the Chevalley–Weil support of Theorem 6.15 is fully symmetric under conjugation $(i,j)\leftrightarrow(a-i,b-j)$ with no parity screening, so the asymmetry source disappears. Thus the task "transport 5.61" is better stated as: the mixed model *cannot* carry that layer; the C item is closed with answer $c'(p)=0$. □

6.4.3 Relation between the two constants

Theorem 6.26 (Exact relation and independent cross-check)**.**

$$c(p) = \frac{p^2+2p+3}{p^2-1}, \qquad \lambda_p = -\frac{2}{p-1}, \qquad \boxed{c(p) = 1-\lambda_p + \frac{2}{p^2-1} = \frac{p+1}{p-1} + \frac{2}{p^2-1}}.$$

Proof. $\frac{p+1}{p-1}+\frac{2}{p^2-1} = \frac{(p+1)^2+2}{p^2-1} = \frac{p^2+2p+3}{p^2-1} = c(p)$, and $1-\lambda_p = 1+\frac{2}{p-1} = \frac{p+1}{p-1}$. Exact rational checks for $p=3,5,7,11,13,17,31,101$ give $c(p)=9/4$, $19/12$, $11/8$, $73/60$, $33/28$, $163/144$, $171/160$, $5203/5100$ with differences $c-(1-\lambda_p) = 1/4, 1/12, 1/24, 1/60, 1/84, 1/144, 1/480, 1/5100$, i.e. $2/(p^2-1)$ in every case. Interpretation: 1 is the 5.60 Fermat anchor constant, $-\lambda_p = 2/(p-1)$ is the mixed-cover offset, and $2/(p^2-1)$ is the orbit correction of the $\frac{1}{2}\log p$ type of Theorem 5.53 (structural correspondence tagged C, not proved here). The identity itself is U and shows $c(p)$ and λ_p are not two approximations to one quantity but differ by two exactly separable sources. □

6.4.4 Endpoint correction to 6.20

Theorem 6.27 (The continuum survives, but not with the borrowed endpoints)**.** *1.*

***Withdrawn**: the claim of Theorem 6.20 that $\bar{c}(w) = w\,c(p)+(1-w)c(q)$ sweeps $(19/12, 9/4)$ on the single curves $C_{3^n,5^m}$. By Theorem 6.25, $\Lambda^{(a)}=0$ there, so the Λ continuum of Theorem 6.4 is* not *realised by mixed single curves; it remains the isogeny-weight continuum of an AV product.*

2. ***Preserved**: the conductor weight $w_p = \hat{f}_p/(\hat{f}_p + \hat{f}_q)$ is continuously tunable and linear-fractional in (n, m) (Theorem 6.20 first half), now with the exact $\hat{f}_p$ of Theorem 6.23. The endpoints $[19/12, 9/4]$ of Theorem 6.4 are unchanged.*

3. ***New**: the mixed model's own exact constant gives $\bar{\lambda}(w) = w\lambda_p + (1-w)\lambda_q$, sweeping $\big[-\frac{2}{p-1}, -\frac{2}{q-1}\big]$. For $p = 3, q = 5$ this is $[-1, -\frac{1}{2}]$, with $w = 0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1$ giving $-\frac{1}{2}, -\frac{5}{8}, -\frac{3}{4}, -\frac{7}{8}, -1$. This continuum is U relative to standard local class field theory and replaces the withdrawn one.*

Consequently 6.13 and 6.20's "curve realisation of the 6.3 continuum" is re-judged: weights *are continuously realisable on a single curve, the* Λ-constant *continuum is not. This is a fifth independent occurrence of the pointwise-discrete / tower-continuous dichotomy (Section 2.1) and of the separation of two tower-side continuities (Theorem 6.4 / Theorem 6.11), now with an exact boundary. The upper half of Problem 10.12 stays open: Theorem 6.11's three realised values are unchanged, since this appendix computes only conductor-layer quantities.*

Proof. (1) is Theorem 6.25 together with the definition of $c(p)$ in Theorem 5.61; a zero layer cannot produce a non-zero endpoint. (2) is Theorems 6.18 and 6.23: the ratio w_p depends only on the two additive terms, and replacing an asymptotic $\hat{f}_p$ by an exact one does not change its tunability. (3) is Theorem 6.24 plus linearity of $\bar{\lambda}$ in w; the seven listed values were computed in exact rational arithmetic. The last assertion records that no height computation is performed here, so Theorem 6.11 is untouched. □

6.4.5 Exact form of the σ bound

Theorem 6.28 (6.1 bound with exact constants)**.** *On the mixed tower,*

$$\sigma_{n,m} - 1 \le \frac{2(\log p + \log q)}{2n\log p + 2m\log q + \lambda_p \log p + \lambda_q \log q + \big(\frac{2n}{p^n-1} + \frac{2m}{q^m-1}\big)\log(pq)}\big(1 + O(g^{-1})\big) \longrightarrow 0.$$

Proof. Substitute the exact $\hat{f}_p, \hat{f}_q$ of Theorem 6.23 into the purely combinatorial bound of Theorem 6.1 (via Theorem 5.17), as in Theorem 6.21. For $p = 3, q = 5$ the right-hand side is 0.55570178 at $(n, m) = (2, 2)$, 0.19768481 at $(6, 5)$, 0.05996283 at $(20, 15)$, 0.02288046 at $(50, 40)$ and 0.00588468 at $(200, 150)$; these reproduce the eight rows of Theorem 6.21 verbatim, because the finite-n corrections of Theorem 6.23 were already contained in the exact $\hat{f}$ values used there. Hence the barrier hypothesis holds on mixed towers with exact constants and without knowledge of Γ_p, Γ_q. □

Remark 6.29 (Honest boundary). Theorems 6.23–6.28 are unconditional relative to standard local class field theory for cyclotomic characters and to the Artin-conductor additivity used in Theorem 6.25; the latter is a standard fact, cited not re-proved. Sixteen exact integer checks, sixteen exact rational identities, eight exact rational relations, seven λ_p values, seven $\bar{\lambda}$ values and eight σ rows were computed. Tag A of Theorem 6.19 is removed, but the result is *not* absolutely unconditional: a non-Kummer group would change $F_p^{(a)}$. Retained as C: the definition of T and the derivation of $c(p)$ in Theorem 5.61 (used here only for comparison), the structural correspondence between $2/(p^2-1)$ and the $\frac{1}{2}\log p$ term of Theorem 5.53, Liu's tame base change and the tree-like inference of Theorem 6.21, the averaged Colmez height for the mixed CM type, and the upper half of Problem 10.12. Withdrawn explicitly: the single-curve $\bar{c}$ endpoints of Theorem 6.20, tag A of Theorem 6.19, and 's formulation of the task as "transporting 5.61 to find a new $c'(p)$" (the answer is $c'(p) = 0$). Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \ge 36/7 > 3$, Routes 1–3, the status of Theorem 7.5/9.2 as $Th.A^*$ and the honest boundary "this does not constitute an independent verification of the proof of Theorem A" are unchanged; *abc* remains open. This appendix verifies no cited proof independently.

6.5 Geometric description of the mixed semistable model

Remark 6.30 (Erratum, Section 8.2). The symbol "rank L" in this appendix denotes $b_1(\Gamma)$, not the Laplacian rank of Theorem 5.10. By Theorem 8.8 the correct difference identity is $\Delta_{\mathrm{Del,p}} - f_p = |V(\Gamma_p)| - 1$, so Theorem 6.34's "$\sigma = 1$ exactly" is withdrawn and replaced by $\sigma - 1 = O(p^{-n}/n + q^{-m}/m) \to 0$. The tree-likeness results (Theorems 6.32–6.33) are unaffected.

Section 6.3.7(2) asserted that the two-point Kummer quotient $C_{a,b}/\mu_b$ has a single-vertex p-adic dual graph with no cluster subdivision, and invoked Liu's tame base change to infer tree-likeness at level C. We compute the graph. The assertion is wrong (the true graph is a depth-n cluster dendrogram), but the conclusion is not merely preserved: it becomes an unconditional structural result, and $\sigma_{n,m} = 1$ holds *exactly* rather than asymptotically. Liu's theorem is no longer needed.

6.5.1 Fibre-product structure

Theorem 6.31 (Fibre product)**.** *With $a = p^n$, $b = q^m$, $\gcd(a, b) = 1$,*

$$C_{a,b} = \{u^a + v^b = 1\} = Z \times_{\mathbb{P}^1_t} Z', \qquad Z = \{u^a = t\} \cong \mathbb{P}^1_u, \quad Z' = \{v^b = 1 - t\} \cong \mathbb{P}^1_v.$$

At $\mathfrak{p} \mid p$, $Z \to Y = \mathbb{P}^1_t$ is the wild μ_{p^n}-cover branched at $0, \infty$; $Z' \to Y$ is tame ($q \neq p$) and $Z' \cong \mathbb{P}^1_v$ is smooth with a single component over $\mathcal{O}_K$; and $X \to Z$ is the tame μ_b-cover $v^b = 1 - u^a$, branched at $D = \mu_{p^n} \cup \{\infty\}$. The three branch points $0, 1, \infty$ of Y remain distinct in the special fibre, so the base model is a single $\mathbb{P}^1_{\mathcal{O}_K}$.

Proof. The identity $u^a = t = 1 - v^b$ displays $C_{a,b}$ as the stated fibre product with group $\mu_a \times \mu_b$. Both Z and Z' are parametrised by u and v respectively, hence isomorphic to $\mathbb{P}^1$; the maps to Y are $t = u^a$ and $t = 1 - v^b$. Since $\gcd(b, p) = 1$, $v^b = 1 - t$ is tamely ramified at $\mathfrak{p}$, and $t = 1 - v^b$ is smooth over $\mathcal{O}_K$. The reduction of $0, 1, \infty$ modulo $\mathfrak{p}$ gives three distinct points of $\mathbb{P}^1_{\mathbb{F}_p}$, so no blow-up of the base is required. □

6.5.2 The cluster dendrogram

Theorem 6.32 (Explicit $\Gamma_p(Z)$)**.** *Let $K \supseteq \mathbb{Q}_p(\zeta_{p^n})$. Then the semistable model of $Z = \mathbb{P}^1_u$ marked by $D = \mu_{p^n} \cup \{\infty\}$ has dual graph the cluster dendrogram of μ_{p^n} with an extra handle to ∞, and*

$$\boxed{|V(\Gamma_p(Z))| = 1 + \frac{p^n - 1}{p - 1} = \frac{p^n + p - 2}{p - 1}, \quad |E| = |V| - 1, \quad \operatorname{rank} L = |E| - |V| + 1 = 0}.$$

Proof. The standard cyclotomic valuation gives $v_K(1 - \zeta_{p^n}^i) = p^{\,n - v_p(i)}$, so the p-adic clustering of μ_{p^n} is $i \sim_k j \iff i \equiv j \pmod{p^k}$ for $k = 0, \ldots, n$: level k has p^k clusters of size p^{n-k}, and $k = n$ gives singletons. Proper clusters (size ≥ 2) number $N_c = \sum_{k=0}^{n-1} p^k = (p^n - 1)/(p - 1)$. All p^n unit roots reduce to $u \equiv 1$, while ∞ is separate; separating them requires one blow-up per proper cluster, producing one $\mathbb{P}^1$ component per cluster plus the original $\mathbb{P}^1$ carrying ∞. Hence $|V| = N_c + 1$. The dual graph is the dendrogram (a rooted p-ary tree) with one extra leaf-edge to ∞, hence a tree, and $\operatorname{rank} L = 0$. Table: for $p = 3$ and $n = 1, \ldots, 8$, $|V| = 2, 5, 14, 41, 122, 365, 3281$ (seven of eight values shown); for $p = 5$: $2, 7, 32, 157, 782, 3907, 97657$; for $p = 7$: $2, 9, 58, 401, 2802, 19609, 960801$; for $p = 11$: $2, 13, 134, 1465, 16106, 177157, 21435889$. All four series satisfy $|E| = |V| - 1$ exactly.

For comparison, Obus Prop. A.3 describes the *three-point fully ramified* $\mathbb{Z}/p^n$ cover, whose graph is a chain with $n + 1$ vertices; here the marked set μ_{p^n} produces a depth-n rooted tree. Both are trees. The value $|V| = (p^{n+1} - 1)/(p - 1)$ of Section 5.5 belongs to the Fermat tower with group $(\mathbb{Z}/p^n)^2$ and must not be imported here. Consequently the three phrases "single-vertex dual graph", "no cluster subdivision" and "two-point Kummer has no cluster" in Theorem 6.21(2) are withdrawn. □

6.5.3 Tree-likeness transfers without Liu

Theorem 6.33 ($\Gamma_p(X) \cong \Gamma_p(Z)$)**.** *For every $p \neq q$ and all n, m,*

$$|V(\Gamma_p)| = \frac{p^n + p - 2}{p - 1}, \qquad |E(\Gamma_p)| = |V| - 1, \qquad \operatorname{rank} L(\Gamma_p) = 0,$$

and $\Gamma_p(X) \cong \Gamma_p(Z)$ as abstract graphs. Moreover $p \neq q$ is necessary: if $p = q$ the Kummer cover splits over some component.

Proof. On a component Z_v of the model of Theorem 6.32, the divisor of $1 - u^a = \prod_{i=1}^{a}(u - \zeta_a^i)$ has: zeros of multiplicity p^j $(0 \le j \le n)$ at the cluster carried by Z_v; multiplicity 0 at the nodes, where $1 - u^a$ is a unit; and a pole of order $a = p^n$ at ∞ on the root component. All multiplicities lie in $\{0\} \cup p^{\mathbb{Z}_{\ge 0}} \cup \{-p^n\}$. A Kummer cover $v^b = g$ has $\gcd(b, \{m_i\})$ irreducible components above Z_v; since $b = q^m$ with $q \neq p$, $\gcd(q^m, p^j) = 1$ for every j, so there is exactly one component above each Z_v. Likewise the thickness of every node is a p-power, so each node has exactly one node above it. Hence the dual graph is unchanged, proving the isomorphism and $\operatorname{rank} L = 0$.

Numerically, for $p = 3, 5, 7$, $n = 1, \dots, 4$ and $b \in \{5^3, 7^2\}$, all 24 legal pairs $p \neq q$ give $\gcd(b, \{p^j\}, a) = 1$; e.g. $p = 3, n = 4$ has multiplicities $\{81, 27, 9, 3\}$ and pole order 81, all coprime to 125. If illegally $p = q$, e.g. $(p, n, q, m) = (5, 1, 5, 3)$, then $\gcd(5^3, 5) = 5 > 1$ and the Kummer cover splits into 5 components, so tree-likeness fails. Thus $\gcd(a, b) = 1$ is not merely a convention but a necessary condition, giving a third occurrence of its role alongside Theorem 6.15's exclusion of the boundary $\frac{i}{a} + \frac{j}{b} = 1$ and Theorem 6.25's conductor additivity.

Liu's tame semistable base change would be a sufficient route to tree-likeness (it was cited in Theorem 6.21(2)), but the gcd argument above makes it unnecessary: that C item is deleted. What remains at text level is only the standard identification of the cluster dendrogram with the semistable dual graph and the Kummer component-count criterion. □

6.5.4 Symmetry and exact σ

Theorem 6.34 (Γ_q and $\sigma_{n,m} = 1$ exactly)**.** *By the symmetry $(u, a, p, n) \leftrightarrow (v, b, q, m)$, $|V(\Gamma_q)| = \frac{q^m + q - 2}{q - 1}$ and $\operatorname{rank} L(\Gamma_q) = 0$. Hence by 5.10 ($\Delta_{\mathrm{Del},\mathfrak{r}} - f_{\mathfrak{r}} = \operatorname{rank} L(\Gamma_{\mathfrak{r}})$),*

$$\Delta_{\mathrm{Del},\mathfrak{p}} = f_{\mathfrak{p}}, \qquad \Delta_{\mathrm{Del},\mathfrak{q}} = f_{\mathfrak{q}}, \qquad \boxed{\sigma_{n,m} = 1 \quad \textit{exactly, not } 1 + o(1)}.$$

Proof. At $\mathfrak{q}$ the wild direction is v (exponent $b = q^m$) and the tame direction is u (exponent $a = p^n$), so Theorems 6.32–6.33 apply verbatim with the roles swapped. 5.10 then gives $\Delta_{\mathrm{Del}} = f$ at both bad primes, and $\sigma = \sum \Delta / \sum f = 1$. Numerical table ($p = 3, q = 5$): $(n, m) = (2, 2)$ gives $|V_p|, |V_q| = (5, 7)$; $(4, 3) \to (41, 32)$; $(6, 5) \to (365, 782)$; $(12, 10) \to (265721, 2441407)$, all with $\operatorname{rank} L_p = \operatorname{rank} L_q = 0$ and $\sigma = 1$. The combinatorial bound of Theorem 6.1 remains valid but is now vacuous: $\sigma - 1 = 0 \le 2(\log 3 + \log 5)/\hat{f}_{\mathrm{tot}}$, whose right-hand side is 0.55570178 at $(2, 2)$ and 0.02288046 at $(50, 40)$. The barrier hypothesis $\sigma = 1$ is therefore obtained directly from the geometric description on mixed towers, without invoking the limit $\hat{f} \to \infty$ of axiom (T3); for a general semistable tower only the asymptotic version of Theorem 6.1 is available. □

6.5.5 Why the Fermat tower is not a tree

Theorem 6.35 (Structural contrast)**.** *The Fermat tower F_{p^n} has Galois group $(\mathbb{Z}/p^n)^2$, hence two wild directions at its single bad prime; its stable model is a fibre product of two Obus chains, contains cycles, and has $|V| = (p^{n+1} - 1)/(p - 1)$, $\operatorname{rank} L = (p^{n+1} - p)/(p - 1) = O(p^n)$, $\sigma - 1 = O(p^{-n}/n)$ (Sections 5.3, KK). The mixed model has exactly one wild direction at each bad prime, hence a tree, $\operatorname{rank} L = 0$ and $\sigma - 1 = 0$. This is the geometric reason for the exact equality of Theorem 6.34.*

Proof. Two wild $\mathbb{Z}/p^n$ directions give, by Obus Prop. A.3, two chains whose fibre product produces a grid with cycles; the counts recalled from Section 5.5 follow. One wild direction plus a tame connected cover gives Theorem 6.33. Explicitly, for $p = 3$ and $n = 2, 4, 6$ the Fermat values are $(|V|, \operatorname{rank} L) = (13, 12)$, $(121, 120)$, $(1093, 1092)$ while the mixed values are $(5, 0)$, $(41, 0)$, $(365, 0)$; for $p = 5$ and $n = 2, 4, 6$ they are $(31, 30)$, $(781, 780)$, $(19531, 19530)$ versus $(7, 0)$, $(157, 0)$, $(3907, 0)$. All six pairs were recomputed exactly. □

Remark 6.36 (Honest boundary). Theorems 6.31–6.35 are unconditional *relative to standard finite-cover and semistable-reduction theory*: the cyclotomic valuation, the identification of the cluster dendrogram with the dual graph of the semistable model, and the Kummer component-count criterion $\gcd(b, \{m_i\})$ are cited, not re-proved. The exact statement $\sigma_{n,m} = 1$ additionally uses 5.10 (Saito–Liu, tag C). Four vertex series (4×7 values), 24 gcd checks, six Fermat/mixed contrasts and eight σ rows were computed exactly. Withdrawn explicitly: the three 6.21(2) phrases, its dependence on Liu's theorem (that C item is removed, so Retained as C: 5.10, Obus's three-point chain (comparison only), the averaged Colmez height for the mixed CM type, the upper half of Problem 10.12, and the independent verification of Theorem A. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5 - 18c)$, $47/3$, $C \geq 36/7 > 3$, Routes 1–3, the status of Theorem 7.5/9.2 as $Th.A^*$ and the honest boundary "this does not constitute an independent verification of the proof of Theorem A" are unchanged; *abc* remains open. This appendix verifies no cited proof independently.

6.6 Finite decomposition of T and the exact constant $c(p)$

The remark of Section 6.4 retained "the definition of T and the derivation of $c(p)$ in Theorem 5.61" as tag C. We replace that derivation by finite integer identities. Given the finite definition $T := \sum_{c \leq n} M_c S_c^-$ (all data already U), the identity $T = D + \frac{1}{2}R$, an exact finite-n closed form for $\Lambda = f_p - 2T$, the limit $c(p)$ and its first error coefficient all follow without citing Theorem 5.61. Consequently $c(p)$ is upgraded from C to U (relative to the cited inputs), and the "structural correspondence" of Theorem 6.26 becomes an exact three-term decomposition.

6.6.1 Layer conductor sums without Colmez

Theorem 6.37 (Exact S_c^-)**.** *With* $n_1^- = (p-1)/2$ *and* $n_j^- = p^{j-2}(p-1)^2/2$ *for* $j \geq 2$ *(Theorem 5.52),*

$$\sum_{j \leq c} n_j^- = \frac{p^{c-1}(p-1)}{2} = \dim_c, \qquad \boxed{S_c^- = \sum_{j \leq c} j\, n_j^- = \frac{1 + p^{c-1}\big(c(p-1) - 1\big)}{2}}.$$

Proof. Using $\sum_{j=2}^{c} j p^{j-2} = \frac{cp^c - (c+1)p^{c-1} - p + 2}{(p-1)^2}$,

$$S_c^- = \frac{p-1}{2} + \frac{(p-1)^2}{2} \cdot \frac{cp^c - (c+1)p^{c-1} - p + 2}{(p-1)^2} = \frac{1 + p^{c-1}(c(p-1) - 1)}{2}.$$

Thirty exact rational checks (p, c) with $p \in \{3, 5, 7, 11, 13\}$ and $c \in \{1, 2, 3, 4, 7, 12\}$ give difference 0; e.g. $p = 3, c = 12$ gives $\dim_c = 177147$ and $S_c^- = 2037191$. Hence the S_c^- half of Theorem 5.53 no longer requires the Colmez asymptotic: the standard Dirichlet-character conductor count of Theorem 5.52 suffices. □

6.6.2 $T = D + \frac{1}{2}R$ is an algebraic identity

Theorem 6.38 (Exact decomposition)**.** *Let* $M_c = p^{c-1}(p+1) - 3$, $\dim_c = p^{c-1}(p-1)/2$, $A_c = M_c \dim_c$,

$$T = \sum_{c \leq n} M_c S_c^-, \quad D = \sum_{c \leq n} A_c \Big(c - \frac{1}{p-1}\Big), \quad R = \sum_{c \leq n} M_c.$$

Then $2(T-D)=R$, *i.e.* $T=D+\frac{1}{2}R$, *with*

$$R=\frac{(p+1)(p^n-1)}{p-1}-3n,$$

$$D=\frac{p^{2n}\big(np^2-n-p-2\big)+p+2}{2(p^2-1)}+\frac{3p^n\big(n+2-np\big)-6}{2(p-1)}.$$

Moreover $g_N=\sum_{c\le n}A_c=\frac{(p^n-1)(p^n-2)}{2}$.

Proof. Expanding $2M_cS_c^- - 2A_c(c-\frac{1}{p-1})$ layerwise with S_c^- from Theorem 6.37, $M_c=p^{c-1}(p+1)-3$ and $A_c=\frac{(p^2-1)p^{2c-2}-3(p-1)p^{c-1}}{2}$, all terms containing c cancel coefficient by coefficient in p^{2c-2} and p^{c-1}; what remains is $p^{c-1}(p+1)-3=M_c$. Hence $2(T-D)=\sum_c M_c=R$. The closed form for R is a geometric sum. For D, use $\sum_{c\le n}cp^{2c-2}=\frac{1-(n+1)p^{2n}+np^{2n+2}}{(p^2-1)^2}$ and $\sum_{c\le n}cp^{c-1}=\frac{1-(n+1)p^n+np^{n+1}}{(p-1)^2}$ together with $\frac{(p^2-1)}{p-1}=p+1$. The identity $g_N=\sum A_c=\frac{(p^n-1)(p^n-2)}{2}$ of Theorem 5.35 is recovered in the same way. Twenty-eight exact checks ($p\in\{3,5,7,11\}$, $n\in\{1,2,3,4,6,10,15\}$) confirm g_N, $D_{\text{dir}}=D_{\text{closed}}$, $R_{\text{dir}}=R_{\text{closed}}$ and $2(T-D)-R=0$ as exact integers/rationals. Thus the identity $T=D+\frac{1}{2}R$ of Theorem 5.59 is upgraded from a cited statement to a self-contained consequence. □

6.6.3 Exact finite-n closed form for Λ

Theorem 6.39 (Λ without O-terms)**.** *With* $f_p=(2n+1)g_N$ *(Theorem 5.60) and* $\Lambda=f_p-2T$ *(Theorem 5.61),*

$$\boxed{\Lambda=\frac{p^2+2p+3}{2(p^2-1)}\,p^{2n}-\Big(\frac{3}{2}+\frac{p+7}{p-1}\Big)p^n+5n+1+\frac{p^2+7p+5}{p^2-1}}.$$

Proof. Substitute $2T=2D+R$ (Theorem 6.38) and $g_N=\frac{p^{2n}-3p^n+2}{2}$. The p^{2n} coefficient is $\frac{2n+1}{2}-\frac{np^2-n-p-2}{p^2-1}$; its numerator $(2n+1)(p^2-1)-2(np^2-n-p-2)=p^2+2p+3$ is independent of n. The p^n coefficient is $-\frac{3(2n+1)}{2}+\frac{3n(p-1)-(p+7)}{p-1}=-\frac{3}{2}-\frac{p+7}{p-1}$. The constant is $5n+1-\frac{p+2}{p^2-1}+\frac{p+7}{p-1}=5n+1+\frac{p^2+7p+5}{p^2-1}$. Thirty-two exact checks ($p\in\{3,5,7,11\}$, $n\le 20$; the $p=11,n=20$ value has 41 digits) give $\Lambda_{\text{dir}}=\Lambda_{\text{closed}}$ with difference 0. This is finite-n information that neither Theorem 5.59 nor Theorem 5.61 supplied. □

6.6.4 $c(p)$ upgraded, with exact error coefficient

Theorem 6.40 (The limit and its first correction)**.**

$$\boxed{c(p)=\lim_{n\to\infty}\frac{\Lambda}{g_N}=\frac{p^2+2p+3}{p^2-1},\qquad \frac{\Lambda}{g_N}-c(p)=-\frac{2(p^2+5p+1)}{p^2-1}\,p^{-n}+O(p^{-2n})}.$$

Proof. With $\alpha=\frac{p^2+2p+3}{2(p^2-1)}$ and $\beta=-\frac{3}{2}-\frac{p+7}{p-1}$, Theorem 6.39 gives $\Lambda=\alpha p^{2n}+\beta p^n+O(n)$, while $g_N=\frac{1}{2}p^{2n}-\frac{3}{2}p^n+1$, hence $\frac{\Lambda}{g_N}=2\alpha+(2\beta+6\alpha)p^{-n}+O(p^{-2n})$ and $2\beta+6\alpha=-\frac{2(p^2+5p+1)}{p^2-1}$. For $p=3$ and $n=1,4,10,20$ the ratios are 1.000000, 2.177215, 2.249894, 2.2499999982, with scaled errors converging to $-6.25=-\frac{2(9+15+1)}{8}$; for $p=5,7,11$ the limits are -4.25, $-3.541\overline{6}$, -2.95, all matching $-\frac{2(p^2+5p+1)}{p^2-1}$ exactly. The eight values $c(3)=9/4$, $c(5)=19/12$, $c(7)=11/8$, $c(11)=73/60$, $c(13)=33/28$, $c(17)=163/144$, $c(31)=171/160$, $c(101)=5203/5100$ reproduce Theorem 5.61 verbatim.

Status. Given the finite integer definition of T, the identities of Theorems 6.38–6.40 no longer cite Theorem 5.61; its derivation of $c(p)$ is therefore upgraded from C to U. The U is *relative*

to Theorem 5.52 (Dirichlet-character conductor count), Theorem 5.1 (layer data), Theorem 5.35 (A_c identity, also re-proved here) and Theorem 5.60 ($f_p = (2n+1)g_N$, verified by two channels in); it is not absolutely unconditional. The C ledger decreases by one item. □

6.6.5 Exact three-term decomposition

Theorem 6.41 (Sources of the three terms of 6.26)**.**

$$\frac{D}{g_N} = n - \frac{1}{p^2-1} - \frac{1}{p-1} + O(p^{-n}), \qquad \frac{R}{g_N} = O(p^{-n}),$$

$$\boxed{c(p) = 1 + \frac{2}{p-1} + \frac{2}{p^2-1} = 1 - \lambda_p + \frac{2}{p^2-1}},$$

where the three terms come respectively from:

1. *the anchor constant $+1$ in $f_p = (2n+1)g_N$ relative to $2ng_N$ (Theorem 5.60; microscopically the $\frac{1}{2}\log p$ layer correction of Theorem 5.53);*
2. *the discriminant shift $-\frac{1}{p-1}$ inside $D = \sum A_c(c - \frac{1}{p-1})$ (Theorem 5.35's d_{K_c} formula), multiplied by -2;*
3. *the geometric-series tail $\sum_c cA_c/g_N \to n - \frac{1}{p^2-1}$ (since $A_c \sim p^{2c-2}$, ratio p^2), multiplied by -2.*

Proof. $\frac{\Lambda}{g_N} = 2n+1-\frac{2D}{g_N} - \frac{R}{g_N}$ by Theorem 6.39; insert the two asymptotics. Eight exact rational checks of $1+\frac{2}{p-1}+\frac{2}{p^2-1} = c(p)$ for $p = 3, 5, 7, 11, 13, 17, 31, 101$ give difference 0, and the deviation $D/g_N - (n - \frac{p+2}{p^2-1})$ at $n = 15$ is $\le 7.84 \times 10^{-8}$ with exponential decay. Thus the "structural correspondence" on the $c(p)$ side of Theorem 6.26 is upgraded from speculation to U; retained as an explanatory C item is only the claim that $\frac{2}{p^2-1}$ and the $\frac{1}{2}\log p$ of Theorem 5.53 are the same p^{-c}-tail mechanism (not proved here).

A curious but unforced observation: $-\lambda_p = \frac{2}{p-1}$ equals numerically the mixed model's $-\lambda_p$ of Theorem 6.24, but the origins are disjoint — here the discriminant shift of A_c, there the constant term of the Kummer conductor sum $\frac{2F}{a-1}$. It is recorded as two independent occurrences of one constant, not as a structural equivalence. □

6.6.6 Independent recovery of the 5.59 asymptotic

Theorem 6.42 (T/g_N)**.**

$$\frac{T}{g_N} = n - \frac{p+2}{p^2-1} + \frac{p^2+5p+1}{p^2-1}\,p^{-n} + O(p^{-2n}).$$

Proof. Since $T = D + \frac{1}{2}R$ (Theorem 6.38), the leading term follows from Theorem 6.41; the first correction is half of Theorem 6.40's coefficient with opposite sign, because $\Lambda = f_p - 2T$. For $p = 3, 5, 7, 11$ the scaled limits are $+3.125 = 25/8$, $+2.125 = 17/8$, $+1.77083\overline{3} = 85/48$, $+1.475 = 59/40$; e.g. $p = 3, n = 12$ gives $T/g_N = 11.37500588$ against the asymptotic $11.375 = n - \frac{5}{8}$, reproducing Theorem 5.59 verbatim. Twelve numerical rows were checked. □

Remark 6.43 (Honest boundary). Theorems 6.37–6.42 are exact finite computations: 30 rational checks for S_c^-, 28 integer identities for $T = D + \frac{1}{2}R$, 32 digit-exact checks of the Λ closed form, 8 rational values of $c(p)$, 8 three-term identities, and 12 asymptotic rows, all in exact rational arithmetic. The upgrade of $c(p)$ to U is *relative* to Theorem 5.52 (standard Dirichlet-character theory, cited not re-proved), Theorem 5.1, Theorem 5.35 and the Theorem 5.60 anchor; it is not absolute unconditionality. Retained as C: 5.10 (Saito–Liu); the standard input cluster (averaged

Colmez finite form, 00712 orbit decomposition, Wilms formula, Obus assembly, Rohrlich anchor); the averaged Colmez height for the mixed CM type and the upper half of Problem 10.12; the explanatory item of Theorem 6.41; and the independent verification of Theorem A. A first-pass error — writing $\Lambda = f_p - 2T - R$, which subtracts R twice since $-R$ is already inside $2T = 2D + R$ — was detected and corrected; all 32 closed-form checks pass after the correction and $c(p)$, the three-term decomposition and all earlier conclusions are unchanged. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, Routes 1–3, the status of Theorem 7.5/9.2 as $Th.A^*$ and the honest boundary "this does not constitute an independent verification of the proof of Theorem A" are unchanged; *abc* remains open. No cited proof is verified independently here.

7 The upper half of Problem 10.12 and mixed CM types

7.1 The upper half of Problem 10.12: exact CM census and the $R > 1/2$ threshold

Theorem 6.11 merged the upper half of Problem 10.12 and the curve realisation of the 6.3 continuum into a single gap: the mixed stable model plus the mixed CM-type Colmez height. Section 6.5 closed the stable model exactly ($\sigma_{n,m} = 1$), Theorem 6.25 gave $\Lambda^{(a)} = 0$, and Section 6.6 made $T = D + \frac{1}{2}R$ and $c(p)$ self-contained. Here we close the remaining finite arithmetic side and reduce what is left to one scalar.

7.1.1 Exact mixed CM type and isogeny decomposition

Theorem 7.1 (CM census)**.** *Let $a = p^n$, $b = q^m$, $\gcd(a,b) = 1$, and $C_{a,b} : u^a + v^b = 1$ with genus $g = \frac{(a-1)(b-1)}{2}$. The CM type is*

$$\Phi_{a,b} = \Big\{(i,j) :\ 1 \leq i \leq a-1,\ 1 \leq j \leq b-1,\ \tfrac{i}{a} + \tfrac{j}{b} < 1\Big\}, \qquad |\Phi_{a,b}| = g.$$

The involution $(-1,-1) \in (\mathbb{Z}/a)^\times \times (\mathbb{Z}/b)^\times$ pairs $\Phi_{a,b}$ with its complement in H^1. Grouping by $c = n - v_p(i)$, $d = m - v_q(j)$, each layer (c,d) consists of exactly one $\mathbb{Q}$-simple CM factor $A_{c,d}$ with field $K_{c,d} = \mathbb{Q}(\zeta_{p^c q^d})$ and

$$\dim_{c,d} = \frac{\varphi(p^c)\varphi(q^d)}{2}, \qquad \sum_{c \leq n,\ d \leq m} \dim_{c,d} = g,$$

so the Jacobian has exactly $n \cdot m$ simple factors. Every character in the layer has conductor exactly *$p^c q^d$.*

Proof. $(\mathbb{Z}/a)^\times \times (\mathbb{Z}/b)^\times$ acts on characters by $(i,j) \mapsto (ti, sj)$; on the set with $v_p(i) = n - c$, $v_q(j) = m - d$ its projection $(\mathbb{Z}/p^c)^\times \times (\mathbb{Z}/q^d)^\times$ is transitive, and the set has size $\varphi(p^c)\varphi(q^d)$. Hence one Galois orbit, i.e. one $\mathbb{Q}$-simple factor, per (c,d); summing $\sum_c \varphi(p^c) = a - 1$ and $\sum_d \varphi(q^d) = b - 1$ gives $\sum \dim_{c,d} = g$. The involution $(i,j) \mapsto (a-i, b-j)$ sends $\frac{i}{a} + \frac{j}{b}$ to $2 - (\frac{i}{a} + \frac{j}{b})$, and $\gcd(a,b) = 1$ excludes the boundary $\frac{i}{a} + \frac{j}{b} = 1$ (Theorem 6.15), so $|\Phi| = g$. Finally $v_p(a-i) = v_p(i)$ and $v_q(b-j) = v_q(j)$, so the pair has equal conductor; the conductor exponents are $n - v_p(i) = c$ and $m - v_q(j) = d$ by Theorem 6.16. Ten brute-force census instances (p,n,q,m) including $(3,3,5,2)$ with $g = 312$ verify $|\Phi| = g$, layer counts $= \dim_{c,d}$, that the pairing preserves each layer, and $\sum \dim = g$, all exactly. □

7.1.2 Conductor mass, CM discriminant, and a constant excess

Theorem 7.2 (Exact finite invariants). *With* $F_p^{(a)} = \sum_{c\le n} c\,\varphi(p^c) = \frac{p^n(np-n-1)+1}{p-1}$ *(Theorem 6.23, Theorem 6.38),*

$$\frac{\mu(\Phi_{c,d})}{\dim_{c,d}} = c\log p + d\log q, \qquad \frac{\log d_{\Phi_{c,d}}}{\dim_{c,d}} = \Big(c - \frac{1}{p-1}\Big)\log p + \Big(d - \frac{1}{q-1}\Big)\log q,$$

$$\boxed{\mu(\Phi_{a,b}) = \tfrac{1}{2}\big[(b-1)F_p^{(a)}\log p + (a-1)F_q^{(b)}\log q\big] = \tfrac{g}{2}\,\hat f_{\mathrm{tot}}},$$

$$\boxed{\mu(\Phi_{a,b}) - \log d_\Phi = g\,\varepsilon_{\mathrm{mix}}, \qquad \varepsilon_{\mathrm{mix}} := \frac{\log p}{p-1} + \frac{\log q}{q-1}}.$$

Consequently $\varepsilon_{c,d} := \frac{\mu - \log d_\Phi}{\dim} = \varepsilon_{\mathrm{mix}}$ *is* independent of (c,d), *exactly.*

Proof. All characters in the layer have conductor p^cq^d (Theorem 7.1), giving the first identity. For the discriminant use $d_{K_{c,d}} = d_{K_c}^{\varphi(q^d)} d_{K_d}^{\varphi(p^c)}$ together with $\log d_{K_c} = p^{c-1}((p-1)c-1)\log p$ and $\log d_\Phi = \frac{1}{2}\log d_K$ (Theorem 5.53); dividing by $\dim_{c,d} = \varphi(p^c)\varphi(q^d)/2$ gives the second identity. Summing with $\sum_c \varphi(p^c) = a-1$, $\sum_c c\varphi(p^c) = F_p^{(a)}$ yields the two boxes, since $\mu - \log d_\Phi = \sum_{c,d}\dim_{c,d}\big(\frac{\log p}{p-1} + \frac{\log q}{q-1}\big) = g\varepsilon_{\mathrm{mix}}$. The equality $\mu = \frac{g}{2}\hat f_{\mathrm{tot}}$ follows from $\hat f_p = 2F_p^{(a)}\log p/(a-1)$ (Theorem 6.23) and $g = \frac{(a-1)(b-1)}{2}$. Ten exact-integer instances and 24 cross-checks of $\hat f_p$ against $2F_p^{(a)}\log p/(a-1)$ give difference 0; e.g. $(p,n,q,m) = (3,3,5,2)$ gives $g = 312$, $\mathrm{coeff}_{\log p}\mu = 816$, $\mathrm{coeff}_{\log q}\mu = 572$.

Compare Theorem 5.54: for the Fermat tower $\varepsilon_c = \rho_c = \frac{\log p}{p^{c-1}(p-1)} \to 0$, because its CM type contains imprimitive characters of conductor p^j with $j \le c$. Mixed layers contain only primitive characters mod p^cq^d, so their excess does not decay. This is a structural difference between the two CM types, not a normalisation artefact. □

7.1.3 Exact archimedean terms and κ-neutrality

Theorem 7.3 (δ_{Fal}, $\varepsilon_{\mathrm{ar}}$). *Relative to the finite form of the averaged Colmez formula* $(\delta_{\mathrm{Fal}}(\Phi)/\dim = -L_2 + (\mu - \log d_\Phi)/\dim$, $L_2 = 2\log(2\pi^4)$, *Theorems 5.46, 5.54),*

$$\boxed{\frac{\delta_{\mathrm{Fal}}}{g} = -L_2 + \varepsilon_{\mathrm{mix}} \quad \textit{exactly}, \qquad \varepsilon_{\mathrm{ar}} = \frac{-L_2 + \varepsilon_{\mathrm{mix}}}{\hat f_{\mathrm{tot}}} \quad \textit{exactly}},$$

and $\varepsilon_{\mathrm{ar}} < 0$ *for all odd* $p \ne q$. *Moreover* $\varepsilon_{\mathrm{mix}}$ *is* κ-neutral*: writing* $\bar\omega^2/g = \frac{\hat f}{2} + 12C_{\mathrm{mix}} + L_2 - \varepsilon_{\mathrm{mix}}$, *Wilms' formula gives*

$$\boxed{\kappa = \frac{1}{8} + \frac{C_{\mathrm{mix}}}{\hat f_{\mathrm{tot}}}, \qquad C_{\mathrm{mix}} := \hat h - \frac{\hat f_{\mathrm{tot}}}{8}},$$

with L_2 *and* $\varepsilon_{\mathrm{mix}}$ *cancelling exactly.*

Proof. Weight $\delta_{\mathrm{Fal}}(\Phi_{c,d})/\dim_{c,d} = -L_2 + \varepsilon_{\mathrm{mix}}$ by $\dim_{c,d}/g$ and use $\sum\dim_{c,d} = g$; dividing by $\log\mathrm{cond} = g\hat f_{\mathrm{tot}}$ gives $\varepsilon_{\mathrm{ar}}$. Since $\varepsilon_{\mathrm{mix}} \le \frac{\log 3}{2} + \frac{\log 5}{4} = 0.951666 < L_2 = 10.544133$, the numerator is negative. For κ-neutrality, Wilms $12\hat h = \bar\omega^2/g + \Delta/g + \delta_{\mathrm{Fal}}/g$ with $\Delta/g = \hat f_{\mathrm{tot}}$ (Theorem 6.34, $\sigma = 1$ exactly) gives $12\kappa = 1 + R + \varepsilon_{\mathrm{ar}}$; substituting R from the stated form of $\bar\omega^2/g$ yields $12\kappa = \frac{3}{2} + 12C_{\mathrm{mix}}/\hat f$, the L_2 and $\varepsilon_{\mathrm{mix}}$ terms cancelling. Numerical table: $\varepsilon_{\mathrm{mix}} = 0.951665622$ and $\delta_{\mathrm{Fal}}/g = -9.592467825$ for $(p,q) = (3,5)$; $(3,7),(5,7),(3,11),(5,11),(7,11)$ give $0.873624503,\ldots,0.564107885$. Seven $\varepsilon_{\mathrm{ar}}$ rows from $(n,m) = (2,2)$ (-0.98420470) to $(50,40)$ (-0.04052363) were checked. This is the fourth occurrence of a κ-neutral census offset, after Theorem 5.3's 8 : 1 cancellation. □

7.1.4 The exact threshold for the upper half

Theorem 7.4 (Upper half $\iff R > 1/2$)**.** *With $\sigma = 1$ (Theorem 6.34),*

$$\boxed{12\kappa = 1 + R + \varepsilon_{\mathrm{ar}}}, \qquad R = \frac{\bar{\omega}^2}{\Delta_{\mathrm{Del}}} = \frac{\bar{\omega}^2}{g\hat{f}_{\mathrm{tot}}}.$$

Hence

$$\kappa > \tfrac{1}{8} \iff R > \tfrac{1}{2} - \varepsilon_{\mathrm{ar}} = \tfrac{1}{2} + |\varepsilon_{\mathrm{ar}}| \iff C_{\mathrm{mix}} > 0.$$

Unconditionally ($\bar{\omega}^2 \geq 0$, Ullmo–Zhang) $\kappa \geq \frac{1}{12} - \frac{L_2 - \varepsilon_{\mathrm{mix}}}{12\hat{f}_{\mathrm{tot}}}$; under arithmetic BMY ($R \leq 3$, tag C) $\kappa \leq \frac{1}{3} - \frac{L_2 - \varepsilon_{\mathrm{mix}}}{12\hat{f}_{\mathrm{tot}}} < \frac{1}{3}$. The three upper-half anchors require $R = 1, \frac{3}{2}, 3$ and $C_{\mathrm{mix}} = \hat{f}/24,\ \hat{f}/12,\ 5\hat{f}/24$ respectively, i.e. a slope-size *offset, not an $O(1)$ constant.*

Proof. Divide Wilms by $\log \mathrm{cond} = g\hat{f}_{\mathrm{tot}}$ and use $\sigma = \Delta/(g\hat{f}_{\mathrm{tot}}) = 1$: $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}} = 1 + R + \varepsilon_{\mathrm{ar}}$ (Theorem 3.1). Since $\varepsilon_{\mathrm{ar}} < 0$ (Theorem 7.3), $\kappa > \frac{1}{8}$ is equivalent to $R > \frac{1}{2} + |\varepsilon_{\mathrm{ar}}|$ and, by Theorem 7.3, to $C_{\mathrm{mix}} > 0$. The lower bound uses $R \geq 0$; the upper bound uses $R \leq 3$, and strictness follows from $\varepsilon_{\mathrm{ar}} < 0$. The $R = 12\kappa - 1$ dictionary (Theorem 2.1) gives $R = 1$ for Mason's $\kappa = 1/6$, $R = 3/2$ for $X(p^n)$ ($\kappa = 5/24$), $R = 3$ for $X_0(N)$ ($\kappa = 1/3$); inserting in $C_{\mathrm{mix}} = \hat{f}(\kappa - \frac{1}{8})$ gives the three offsets. Numerical table ($p = 3, q = 5$): $\kappa(R = \frac{1}{2})$ runs 0.04298294 at $(2,2)$ to 0.12162303 at $(50,40)$, always $< \frac{1}{8}$. □

7.1.5 Under the additive Colmez law the mixed tower stays below 1/8

Theorem 7.5 (Conditional negative result)**.** *Assume* (A)*: the affine law of 00712 Thm 3.4 extends additively to coprime conductors, i.e. $h_F(A_{c,d})/\dim_{c,d} = \frac{1}{4}(c\log p + d\log q) + \gamma_{p,q}$ and $\log \mathrm{cond}(A_{c,d})/\dim_{c,d} = 2(c\log p + d\log q) + b_{p,q}$ with $\gamma_{p,q} = \gamma_p + \gamma_q$, $b_{p,q} = b_p + b_q$. Then $8\gamma_{p,q} - b_{p,q} = 8(C_p + C_q)$ (Theorem 5.3), each factor has $\kappa_{c,d} < \frac{1}{8}$, and by the convexity of Theorem 3.17*

$$\boxed{\kappa_{n,m} = \frac{1}{8} + \frac{C_p + C_q}{\hat{f}_{\mathrm{tot}}} + O(\hat{f}^{-2}) < \frac{1}{8}, \qquad \sup_{n,m} \kappa_{n,m} = \frac{1}{8} \text{ (not attained)}}.$$

Moreover $\kappa > 0$ requires $\hat{f}_{\mathrm{tot}} > 8|C_p + C_q| = 25.374988$ for $(p,q) = (3,5)$.

Proof. Since $8\gamma_p - b_p = 8C_p$ (Theorem 5.3), $\kappa_{c,d} = \frac{1}{8} + \frac{C_p + C_q}{2(c\log p + d\log q) + b_{p,q}} + O(c^{-2})$ is below $\frac{1}{8}$ because $C_p, C_q < 0$. By isogeny additivity and Theorem 3.17, $\kappa_{n,m}$ is a convex combination of the $\kappa_{c,d}$, hence below $\frac{1}{8}$; summing the affine laws gives $\hat{h} = \hat{f}_{\mathrm{tot}}/8 + C_p + C_q$, i.e. the displayed closed form. With $C_3 = -1.595521764$, $C_5 = -1.576351753$, $C_3 + C_5 = -3.171873517$, the table gives $\kappa = -0.20044001$ at $(2,2)$, $+0.00922829$ at $(6,5)$, $+0.11160034$ at $(50,40)$ and $+0.11832700$ at $(100,80)$, all $< \frac{1}{8}$. Note also that even if $R \equiv \frac{1}{2}$ exactly, Theorem 7.4 gives $\kappa = \frac{1}{8} + \frac{\varepsilon_{\mathrm{ar}}}{12} < \frac{1}{8}$ strictly. Thus the mixed two-prime Kummer tower produces no point of the upper half, and its closure lies in the same $(0, \frac{1}{8}]$ as Theorem 6.8. □

7.1.6 The remaining gap is one scalar function

Theorem 7.6 (Refinement of 6.11)**.** *Of the three ingredients of Theorem 6.11, the stable model and Γ_p, Γ_q are closed by Theorem 6.32 and 6.33, $\sigma = 1$ exactly by Theorem 6.34, $\Lambda^{(a)} = 0$ by Theorem 6.25, and $T = D + \frac{1}{2}R$, $c(p)$ by Theorem 6.38 and 6.40. Consequently*

$$\boxed{\textit{upper half of Problem 10.12} \iff \exists (n,m):\ C_{\mathrm{mix}} > 0 \iff R > \tfrac{1}{2} + |\varepsilon_{\mathrm{ar}}|},$$

and the only missing input is the averaged Colmez height of the mixed CM type $\Phi_{c,d}$ (primitive characters mod $p^c q^d$ cut by $u/p^c + v/q^d < 1$), equivalently the spectral constant $\bar{\rho}_{c,d}$ of 00712

(4.1), equivalently $\bar{\omega}^2/g$ (the Bergman-kernel step of Rem. 4.6 of 00544). It cannot be supplied by extrapolating finite data: the factor-2 discrepancy between (1.3) and (4.1) of 00712 is unresolved (Remark 5.5). The three upper-half anchors $R = 1, \frac{3}{2}, 3$ all come from modular/Shimura curves, not Kummer covers; Kummer covers have an affine Colmez height $\hat{h} = \hat{f}/8 + C$ and hence $R \to \frac{1}{2}$ intrinsically.

Proof. The listed closures are verbatim from the cited theorems. Combining with Theorems 7.3–7.5 gives the boxed equivalence. The 6.3 Λ continuum remains an isogeny-weight statement only (Theorem 6.25); with $\Lambda^{(a)} = 0$ and the constant $\varepsilon_{\mathrm{mix}}$ it now has two exact witnesses for the separation of the two kinds of continuity recorded in Section 6.2, which is therefore unchanged. The observation about anchors follows from Theorem 7.4's table. □

Remark 7.7 (Honest boundary). Theorems 7.1–7.2 are exact: 10 brute-force CM censuses, 10 exact-integer invariant checks, 24 conductor cross-checks, all in integer/Fraction arithmetic. Theorems 7.3–7.4 are unconditional *relative to* the finite averaged-Colmez form, Wilms' formula and $\bar{\omega}^2 \geq 0$ (Ullmo–Zhang); BMY's $R \leq 3$ is tag C. Theorem 7.5 is conditional on the explicit hypothesis (A): additivity of the averaged Colmez height for coprime conductors is *not* proved here, so "mixed towers never reach the upper half" is not claimed unconditionally. Still open: the mixed CM-type averaged Colmez height, the (1.3)/(4.1) factor-2 ambiguity, the Bergman-kernel $\to \bar{\omega}^2$ step, the upper half of 10.12 beyond the three anchors, the number-field realisation of $\kappa = 1/6$, Problem 10.1 for general towers, and the independent verification of Theorem A. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, Routes 1–3, the status of Theorem 7.5/9.2 as $Th.A^*$ and the honest boundary "this does not constitute an independent verification of the proof of Theorem A" are unchanged; *abc* remains open. No cited proof is verified independently here, and

7.2 Exploring hypothesis (A): refutation and withdrawal of 7.5

Theorem 7.5 derived $\kappa_{n,m} < \frac{1}{8}$ and $\sup \kappa = \frac{1}{8}$ for the mixed two-prime Kummer tower from the explicit hypothesis (A): the affine law of 00712 Thm 3.4 extends additively to coprime conductors, so that $\gamma_{p,q} = \gamma_p + \gamma_q$. Here we test (A) against the verbatim Colmez–Lerch formulas of 00712. It is refuted; nevertheless the upper half of Problem 10.12 does not open, because the failure is an $O(1)$ constant, not a slope.

7.2.1 The four verbatim formulas

From 00712, formulas (1.3)–(1.6) and (4.1)–(4.2):

$$h_F(A) = \sum_{\chi} m(\Phi,\chi)\Big[-\tfrac{1}{2}\tfrac{L'}{L}(0,\chi)\Big] - \tfrac{1}{4}\log d_\Phi - \tfrac{g}{2}\log 2\pi, \tag{1.3}$$

$$m(\Phi,\chi) = \frac{1}{|G|}\sum_{g\in G}\Big(|\Phi \cap g\Phi| - \tfrac{g}{2}\Big)\bar{\chi}(g), \quad \text{"vanishes on even characters"}, \tag{1.4}$$

$$\frac{L'}{L}(0,\chi) = -\log f_\chi - \frac{1}{B_{1,\chi}}\sum_{a=1}^{f_\chi}\chi(a)\log\Gamma(a/f_\chi), \tag{1.5}$$

$$R_\chi = \frac{L'}{L}(0,\bar{\chi}), \qquad \rho_\chi = R_\chi + \log f_\chi, \tag{1.6}$$

$$h(\Phi) = \frac{\mu(\Phi)}{2} - \sum_{\chi} m(\Phi,\chi)\rho_\chi, \qquad \mu(\Phi) = \sum_{\chi} m(\Phi,\chi)\log f_\chi, \tag{4.1}$$

$$\bar{\rho} = 2.4194\ (\pm 0.04),\ k = 1,\dots,6, \quad \mathrm{sd} \to 0.57. \tag{4.2}$$

The anchor is reproduced digit by digit: $B_{1,\chi_{-3}} = -\frac{1}{3}$ and $\sum \chi \log\Gamma = \log\Gamma(\frac{1}{3}) - \log\Gamma(\frac{2}{3}) = 0.68227$ give $\rho_{\chi_{-3}} = 2.046811$ and $R = 0.948199$, as in 00712.

7.2.2 Even factors make the factorwise additive law meaningless

Theorem 7.8 (An elementary vanishing lemma)**.** *Let χ be an even primitive Dirichlet character mod $f > 2$ ($\chi(-1) = +1$). Then*

$$\boxed{B_{1,\chi} := \frac{1}{f}\sum_{a=1}^{f-1}\chi(a)a = 0},$$

hence $L(0,\chi) = -B_{1,\chi} = 0$, and $\frac{L'}{L}(0,\chi)$, R_χ and ρ_χ are undefined *in (1.5)–(1.6) (division by zero).*

Proof. Since $\chi(f-a) = \chi(-a) = \chi(a)$ and $\sum_{a=1}^{f-1}\chi(a) = 0$ for a nonprincipal character,

$$\sum_{a=1}^{f-1}\chi(a)a = \sum_{a=1}^{f-1}\chi(f-a)(f-a) = f\sum\chi(a) - \sum\chi(a)a = -\sum\chi(a)a,$$

so the sum, and therefore $B_{1,\chi}$, vanishes. In (1.5) the term $1/B_{1,\chi}$ is then infinite. Numerically $\max|B_{1,\chi}| \le 1.4\times10^{-14}$ over all even primitive characters mod $3^2, 5^2, 7, 11$. □

Theorem 7.9 (Refutation of the factorwise form of (A))**.** *Every character $\chi = \chi_1\chi_2$ occurring in the mixed CM type $\Phi_{c,d}$ of Theorem 7.1 is overall odd, $\chi(-1) = \chi_1(-1)\chi_2(-1) = -1$. Hence* exactly one *of χ_1, χ_2 is even, and by Theorem 7.8 its ρ is undefined. Consequently the identity $\rho_{\chi_1\chi_2} = \rho_{\chi_1} + \rho_{\chi_2}$ — the factorwise mechanism underlying hypothesis* (A) *— is not false but* meaningless. *The mixed character itself has $B_{1,\chi} \neq 0$ (measured $\min|B_1| \ge 1.035$), so $\rho_{\chi_1\chi_2}$ is finite and well defined; what fails is its decomposition.*

Proof. Since -1 has index $\varphi(p^c)/2$ and $\varphi(q^d)/2$ in the two cyclic factors, $\chi_i(-1) = (-1)^{j_i}$ with j_i the character index; overall odd means $j_1 + j_2$ odd, so exactly one index is even, i.e. exactly one χ_i is even. Exact enumeration of nine layers (p,c,q,d) gives occurrence counts $(oo, oe, eo, ee) = (0, *, *, 0)$ in every case, e.g. $(3,2,5,2)$: $(0,16,16,0)$; $(3,3,5,1)$: $(0,6,12,0)$; $(3,1,11,1)$: $(0,4,0,0)$. Apply Theorem 7.8. Finiteness of the overall ρ follows from $\chi(-1) = -1$ and the measured $\min|B_1|$. □

7.2.3 The layer mean is one universal constant, not a sum

Theorem 7.10 (Universality of $\bar\rho$ on mixed layers)**.** *Computing $\rho_\chi = -\overline{S_\chi}/\overline{B_{1,\chi}}$ with $S_\chi = \sum_a \chi(a)\log\Gamma(a/f)$ over all primitive overall-odd characters mod p^cq^d gives*

$$\boxed{\bar\rho_{c,d} \longrightarrow \bar\rho_\infty = 2.4194 \pm 0.04}$$

— the same *constant as on single-prime layers (Theorem 4.2 of 00712), not $\bar\rho_p + \bar\rho_q \approx 4.84$. Explicitly $\bar\rho_{\mathrm{mix}} - \bar\rho_p - \bar\rho_q \to -\bar\rho_\infty$.*

Proof. Exact evaluation (40–45 digits for the small layers, double precision for the large ones, cross-checked):

(p,c,q,d)	f	$\#\rho$	$\bar\rho_{\mathrm{mix}}$	sd
$(3,2,5,2)$	225	32	2.425347	0.766
$(3,3,5,2)$	675	96	2.401272	0.735
$(3,4,5,2)$	2025	288	2.417965	0.743
$(3,3,7,1)$	189	30	2.409866	0.799
$(5,2,7,2)$	1225	288	2.417720	0.788
$(3,4,7,2)$	3969	648	2.414675	0.758

Single-prime controls: $(3,3)$ gives 2.425205, $(5,3)$ gives 2.426421, $(7,3)$ gives 2.424335; the standard deviations 0.74–0.80 match the 0.57 concentration of (4.2) in order of magnitude. The additive comparison gives $\bar\rho_{\mathrm{mix}} - \bar\rho_p - \bar\rho_q = -2.5125$ at $(2,2)$ and -2.2118 at $(3,1)$, drifting to $-\bar\rho_\infty$ on large layers. The reason is structural: the proof skeleton of Thm 3.4 of 00712 states verbatim that "Lerch (1.5) collapse[s] each conductor layer by the Gauss multiplication theorem"; the resulting $\bar\rho$ is one Kronecker–Lerch spectral constant and does not double when the conductor contains two primes. Hypothesis (A) mistook a universal height constant for an additive quantity. □

7.2.4 The slope survives: the upper-half threshold is not met

Theorem 7.11 ($C_{\mathrm{mix}} = O(1)$)**.** *By (4.3),* $h(\Phi) = \frac{\mu(\Phi)}{2} - \bar\rho\frac{\dim}{2} + dev(\Phi)$*; with* $\mu_{c,d}/\dim_{c,d} = c\log p + d\log q$ *exactly (Theorem 7.2),* $\hat f_{c,d} = 2(c\log p + d\log q)$ *(Theorem 6.23), and bounded* $\bar\rho_{c,d}$ *(Theorem 7.10),*

$$\boxed{\hat h_{c,d} = \frac{c\log p + d\log q}{4} + \gamma_{p,q} + o(1), \qquad C_{\mathrm{mix}} = \hat h_{c,d} - \frac{\hat f_{c,d}}{8} = \gamma_{p,q} = O(1)}.$$

Hence the upper-half criterion of Theorem 7.4, $C_{\mathrm{mix}} \asymp \hat f_{\mathrm{tot}}$ *(the three anchors require* $\hat f/24$*,* $\hat f/12$*,* $5\hat f/24$*), is* still not met*: refuting* (A) *does not open the upper half of Problem 10.12.*

Proof. Divide (4.3) by dim; the $\mu/(2\dim)$ term equals $\frac14(c\log p + d\log q)$ because $\hat f = 2\mu/\dim$ (the CM type and its conjugate contribute equally), the $\bar\rho$ term is bounded by Theorem 7.10, and the $dev/\dim$ term is $o(1)$ by the concentration of (4.2). Thus $C_{\mathrm{mix}} = \hat h - \hat f/8$ is an $O(1)$ constant independent of (n,m) to leading order. Theorem 7.4 requires $C_{\mathrm{mix}} \asymp \hat f$ for any upper-half anchor; an $O(1)$ offset cannot supply it. □

7.2.5 Calibration failure and the honest boundary

Theorem 7.12 (Withdrawal of 7.5)**.** *Hypothesis* (A) *is refuted in both its factorwise (Theorem 7.9) and layer-averaged (Theorem 7.10) forms. Consequently Theorem 7.5 must be* withdrawn*:* $\kappa_{n,m} < \frac18$ *and* $\sup\kappa = \frac18$ *can no longer be quoted even conditionally. The upper half of Problem 10.12 returns to the single scalar criterion of Theorem 7.6, now sharper:*

$$\boxed{\textit{upper half} \iff C_{\mathrm{mix}} > 0, \qquad C_{\mathrm{mix}} = O(1) \textit{ determined by } \bar\rho_\infty \textit{ and } dev}.$$

The sign of C_{mix} *is* not *settled here.*

Proof. An attempt to calibrate the physical height directly from (1.3)+(1.4) with the CM types $\{u : 1 \le u < f/2\}$ and $\{(u,v) : u/p^c + v/q^d < 1\}$ fails: the offset $c_0 = \hat h^{\mathrm{calc}} - (\frac{c}{4}\log p + C_p + \frac{\log p}{8})$ drifts as $-0.110, -0.511, -0.820$ for $(p,c) = (3,1),(3,2),(3,3)$, -1.335 for $(5,3)$ and -1.664 for $(7,3)$, while the same normalisation gives $\bar\rho^{\mathrm{calc}}$ drifting from 2.047 to 2.884 instead of the constant 2.4194. Hence $\{u < f/2\}$ is *not* the CM type of 00712, which comes from the three-point Fermat orbits $(r,s,t) \in [1,N)^3$, $r+s+t \equiv 0$ (Thm 1.1); the spectrum of $m(\Phi,\chi)$ in (1.4) differs. This is a text-level alignment item, not a mathematical error, and it prevents us from promoting the sign of C_{mix} to a theorem.

What remains is data-level evidence: eighteen evaluations in two independent normalisations all give $C_{\mathrm{mix}} < 0$ (the (1.3)+(1.4) route gives -0.964 to -1.565 over nine layers; the semi-quantitative $\gamma_{p,q} - (\gamma_p + \gamma_q) \approx \frac{\bar\rho_\infty}{2} + \frac{\varepsilon_{\mathrm{mix}}}{4} \approx +1.45$ gives $C_{\mathrm{mix}} \approx C_p + C_q + 1.45 \approx -1.72$). Same sign, same order of magnitude, but neither is a proof. We therefore record (A) as refuted, Theorem 7.5 as withdrawn, and the upper half as still unopened with an undetermined $O(1)$ sign. □

Theorem 7.13 (Ledger)**.**

- *Hypothesis* (A)*: explicit assumption →* ***refuted*** *(strictly at factorwise level, numerically at layer-mean level); deleted.*

- *Theorem 7.5:* ***withdrawn****; the claim "Kummer towers never reach the upper half" is no longer asserted, conditionally or otherwise.*
- $\bar{\rho}_{c,d} \to \bar{\rho}_\infty = 2.4194 \pm 0.04$*: new, with* 648*-character layers;* $B_{1,\chi} = 0$ *for even primitive characters: new elementary lemma.*
- $C_{\mathrm{mix}} = O(1)$*: new, independent of* (A)*; the upper-half threshold of Theorem 7.4 remains unmet, so the refutation of* (A) *does not open the upper half.*
- *New text-level item* (C)*: align the three-point CM type of 00712 Thm 1.1 with (1.4) before calibrating physical heights.*

No new external reference, no new mathematical input to the C layer. Fourteen No-Go reasons are unchanged, none added; $(N/M_5)^{24+\varepsilon}$*,* $E(c) = 12/(5-18c)$*,* $47/3$*,* $C \geq 36/7 > 3$*, Routes 1–3, the status of Theorems 7.5/9.2 as Th.A* and Section 1.3's honest boundary are unchanged; abc remains open.*

Proof. Collect Theorems 7.8–7.12 and the ledger of Section 7.1. The 14 No-Go reasons concern the pointwise line and are unaffected by a tower-side hypothesis test; all displayed constants are quoted verbatim from the standing conclusions. □

Remark 7.14 (Honest boundary). Theorems 7.8–7.9 are strict elementary statements (4 numerical checks of $|B_{1,\chi}| \leq 1.4 \times 10^{-14}$, 9 exact parity censuses). Theorem 7.10 is unconditional *relative to* the Kronecker–Lerch form (1.5)–(1.6) of 00712; Thm 4.2 of 00712 is certified for layers $k \leq 6$ and stated as "expected for all primes". Theorems 7.11–7.13 use (4.1)/(4.3), 7.2 and 6.23 as inputs. The sign of C_{mix} is data-level only, because the CM-type calibration of Theorem 7.12 fails. Still open: that sign, the mixed CM-type averaged Colmez height, the factor-2 ambiguity between (1.3) and (4.1) of 00712 (Remark 5.5), the Bergman-kernel $\to \bar{\omega}^2$ step, the upper half of 10.12, the number-field realisation of $\kappa = 1/6$, Problem 10.1 for general towers, and the independent verification of Theorem A. No cited proof is verified independently here.

7.3 Textual alignment: the tower-level normalisation

Appendix 7.2 recorded a *calibration failure*: inserting the CM type $\{u : 1 \leq u < f/2\}$ into (1.3)+(1.4) produced an offset $c_0 = \hat{h}^{\mathrm{calc}} - \hat{h}^{(3.5)}$ drifting linearly from -0.110 to -1.664, and $\bar{\rho}^{\mathrm{calc}}$ drifting from 2.047 to 2.884. This was attributed to a text-level mismatch between the CM type of 00712 and $\{u < f/2\}$. A verbatim reading of 00712 refutes that diagnosis. The failure was a *single-orbit versus whole-tower* error, and the tower-level normalisation of Prop 2.5 closes it with no extra constant.

7.3.1 Three verbatim normalisations

Thm 1.1	$J(F_d) \sim \bigoplus_{O \subset Y_n} A_O$, $Y_n = \{(r,s,t) \in [1,N)^3 : r+s+t \equiv 0\}$, $\dim A_O = \varphi(d_0)/2$
Prop 2.5	$w_\chi^{\mathrm{intr}} = \sum_O m^{\mathrm{intr}}(\Phi_O, \chi) = c_{v_\chi} \lvert B_{1,\chi}\rvert^2$, $c_v = \frac{3p(p^{v+1}-1)}{(p-1)^2}$
(2.4)	$\sum_{\chi \text{ odd}} c_v \lvert B_{1,\chi}\rvert^2 = g_n/2$
line 200	$\hat{h}_n = h_F(J(F_{p^n}))/g_{p^n}$ **(whole tower)**
(3.5)	$\hat{h}_n = \frac{\log p}{4} n - 6\log A - \frac{p\log p}{4(p^2-1)} + \frac{\frac{9}{4}n\log p + b_p}{p^n} + o(p^{-n})$
(4.1)–(4.3)	$h(\Phi) = \frac{\mu(\Phi)}{2} - \sum_\chi m(\Phi,\chi)\rho_\chi = \frac{\mu}{2} - \bar{\rho}\frac{d}{2} + dev$, $\bar{\rho}$ **equally weighted**

Here j is odd, $1 \leq j < \varphi(p^n)$, $v = v_p(j)$ is the valuation of the *index* (not the conductor exponent), $f_{\chi_j} = p^{n-v}$, and the primitive character mod f has index $j_0 = j/p^v$.

7.3.2 Prop 2.5 verified digit by digit

Theorem 7.15 (Tower-level Bernoulli-moment identity)**.** *Computing* $B_{1,\chi} = \frac{1}{f}\sum_{a=1}^{f} \chi(a)a$ *and* $\rho_\chi = -\overline{S_\chi}/\overline{B_{1,\chi}}$ *from (1.5)–(1.6) and setting* $m(\chi_j) = c_{v_p(j)} \lvert B_{1,\chi_j}\rvert^2$ *gives, for* $p = 3, 5, 7, 11$ *and*

$n = 2, 3, 4, 5$,

$$\boxed{\sum_{\chi \text{ odd}} m(\chi) = \frac{g_n}{2} = \frac{(p^n - 1)(p^n - 2)}{4}}$$

with residual $\leq 3.8 \times 10^{-6}$ *(pure floating point at values up to* 6.5×10^9*). All seven certification values of 00712 are reproduced verbatim:* 14, 162.5, 1580 *(*$p = 3$*);* 138, 3813 *(*$p = 5$*);* 564, 29155.5 *(*$p = 7$*).*

Proof. Use one FFT per conductor layer $p^{c'}$ to evaluate $X_j = \sum_t w[t] e^{2\pi i j t/\varphi(p^{c'})}$ for $w[t] = a_t$ and $w[t] = \log \Gamma(a_t/p^{c'})$ simultaneously; then $B_{1,\chi_j} = X_j^B/f$, $S_{\chi_j} = X_j^S$. Assemble the tower m-vector by $j \mapsto (v, c' = n - v, j_0 = j/p^v)$ and sum. The table:

p, n	$\sum m$	$g_n/2$	residual
3, 2 / 3, 3 / 3, 4 / 3, 5	14 / 162.5 / 1580 / 14580.5	same	$\leq 3 \times 10^{-13}$
5, 2 / 5, 3 / 5, 5	138 / 3813 / 2439063	same	$\leq 5 \times 10^{-10}$
7, 2 / 7, 3 / 7, 5	564 / 29155.5 / 70606207.5	same	$\leq 3 \times 10^{-7}$
11, 2 / 11, 4	3570 / 53578740	same	$\leq 2 \times 10^{-8}$

This is the *tower-level* m-vector of (4.1); it is not the m-vector of any single orbit. □

7.3.3 The calibration does not fail: $c_0 = O(n/p^n)$

Theorem 7.16 (Withdrawal of the 7.11 calibration failure)**.** *Insert the tower-level* m*-vector of Theorem 7.15 into (4.1),* $h_n^{\text{tower}} = \frac{1}{2}\sum_j m_j \log f_j - \sum_j m_j \rho_{\chi_j}$*, and set* $\hat{h}_n^{\text{calc}} = h_n^{\text{tower}}/g_n$*. Then*

$$\boxed{c_0(n,p) := \hat{h}_n^{\text{calc}} - \left(\frac{n \log p}{4} + C_p\right) = O(n/p^n) \longrightarrow 0},$$

so (4.1) and (3.5) agree without any additional normalisation constant*. Appendix 7.2's diagnosis of a calibration failure is withdrawn.*

Proof. Exact evaluation gives

$p \backslash n$	2	3	4	5	6
3	0.347327	0.182874	0.087121	0.038515	0.016126
5	0.189749	0.063186	0.018184	0.004782	0.001187
7	0.121443	0.029157	0.005956	0.001110	0.000196
11	0.063495	0.009687	0.001247	—	—

Successive column ratios tend to $1/p$ in order of magnitude ($p = 5$: $0.333, 0.288, 0.263, 0.248$; $p = 7$: $0.240, 0.204, 0.186, 0.177$), i.e. c_0 has exactly the $\frac{\frac{9}{4} n \log p + b_p}{p^n}$ error order of (3.5). Hence the Colmez→Faltings $\frac{1}{2}\log 2\pi$ translation is already contained in (4.1), unifying the two-anchor normalisation of Section 5.6 and the "three normalisations differ by $c \cdot g$" of Section 3.2. □

7.3.4 $\{u < N/2\}$ is a legitimate three-point orbit

Theorem 7.17 ($\Phi_{(1,1,N-2)} = \{a \in G : a < N/2\}$)**.** *Let* $N \geq 3$*,* $G = (\mathbb{Z}/N)^\times$ *and* $\tau = (1, 1, N - 2) \in Y_N$ *(a primitive three-point orbit,* $d_0 = N$*). The CM type of Thm 1.1,* $\Phi_\tau = \{a \in G : \langle a \rangle_N + \langle a \rangle_N + \langle a(N-2) \rangle_N = N\}$*, equals* $\{a \in G : a < N/2\}$ *exactly. Consequently the claim in Appendix 7.2 that "*$\{u < f/2\}$ *is not the CM type of 00712" is* ***withdrawn****.*

Proof. For a unit a with $1 \leq a \leq N - 1$ we have $2a \not\equiv 0$, so $\langle -2a \rangle_N = N - 2a$ if $2a < N$ and $2N - 2a$ if $2a > N$. The sum is $2a + (N - 2a) = N$ in the first case and $2a + (2N - 2a) = 2N$ in the second. Exact enumeration at $N = 3, 5, 7, 9, 11, 15, 25, 27, 35, 45, 49, 63, 75, 81$ confirms equality and $|\Phi_\tau| = |G|/2$ in all fourteen cases, including the mixed $N = p^c q^d$. □

Theorem 7.18 (The real source of the linear drift)**.** *A single orbit has* $\sum_\chi m(\Phi_\tau,\chi)=\varphi(N)/4$*, whereas the whole tower has* $\sum_\chi m^{\text{tower}}(\chi)=g_n/2$ *(Theorem 7.15). Their ratio*

$$\frac{g_n/2}{\varphi(p^n)/4}=\frac{(p^n-1)(p^n-2)}{2\varphi(p^n)}\sim\frac{p^n}{2(p-1)}\longrightarrow\infty$$

is 1.000 *at* $(p,n)=(3,1)$ *(the unique point where the anchor reproduced),* 1089 *at* $(3,6)$*, and* 1.37×10^5 *at* $(7,6)$*. Comparing a single-orbit height with the whole-tower* $\hat h_n$ *of line 200 therefore produces an* $O(n)$ *discrepancy — precisely the linear drift* $-0.110\to-1.664$ *of Appendix 7.2.*

Proof. By (1.4), $\sum_\chi m(\Phi,\chi)=\frac{|\Phi|}{2}=\frac{|G|}{4}$ for a free orbit (the principal-character term is removed by the $-\frac{g}{2}$ subtraction). Theorem 7.15 gives the tower value $g_n/2$. Their quotient is displayed; the $O(n)$ drift follows because h^{tower}/g_n and $h^{\text{single}}/\varphi(N)/2$ have different linear coefficients in n. □

7.3.5 Two notions of $\bar\rho$; $C_{\text{mix}}=O(1)$ promoted

Theorem 7.19 (Equally weighted versus m-weighted)**.** *The* $\bar\rho=2.4194\pm0.04$ *of Theorem 7.10 is the* equally weighted layer mean *of 00712 Thm 4.2. The quantity actually entering (4.1) is the* m*-weighted mean* $\bar\rho_m=\sum m\rho/\sum m$*, which is* p*-dependent and bounded:*

$$\bar\rho_m\to2.8475\ (p=3,n=10),\quad 2.9179\ (p=5,n=8),\quad 2.9445\ (p=7,n=8).$$

Hence Theorem 7.10 must be read with the qualifier "equally weighted layer mean"; boundedness of $\bar\rho_m$ *still gives, unconditionally relative to (1.5)–(1.6),*

$C_{\text{mix}}=O(1)$ *promoted to* U, *the slope threshold* $C_{\text{mix}}\asymp\hat f_{\text{tot}}$ *of Theorem 7.4 remains unmet.*

Refuting hypothesis (A) *still does not open the upper half of Problem 10.12.*

Proof. Evaluate $\bar\rho_m$ with the tower m-vector of Theorem 7.15: $p=3$ gives $2.7373, 2.8269, 2.8444, 2.8472$ at $n=4,6,8,10$; $p=5$ gives $2.8941, 2.9164, 2.9179$ at $n=4,6,8$; $p=7$ gives $2.9366, 2.9443, 2.9445$. All are bounded by 2.95. By (4.3), $\hat h=\frac{\mu}{2\dim}-\frac{\bar\rho_m}{2}+o(1)$ with $\mu/\dim=c\log p+d\log q$ exact (Theorem 7.2) and $\hat f=2\mu/\dim$ (Theorem 6.23); hence $C_{\text{mix}}=\hat h-\hat f/8=\gamma_{p,q}=O(1)$. Theorem 7.4 requires $C_{\text{mix}}\asymp\hat f$ for any upper-half anchor. □

7.3.6 The mixed model is a different curve

Theorem 7.20 (Object mismatch and the status of the sign)**.** *The mixed Kummer curve* $C_{a,b}:u^a+v^b=1$ *(*$N=ab$*,* $\gcd(a,b)=1$*) of Sections 6.3–7.1 is* not *the Fermat curve* $F_N:x^N+y^N=1$ *of 00712:*

$$g(C_{a,b})=\frac{(a-1)(b-1)}{2},\quad \Phi_{a,b}=\{(i,j):i/a+j/b<1\}\subset\mathbb{Z}/a\times\mathbb{Z}/b,$$

whereas $g(F_N)=\frac{(N-1)(N-2)}{2}$*, the CM type is indexed by three-point orbits of* $(\mathbb{Z}/N)^\times$*, and* $|\Phi|=\varphi(d_0)/2$*. For* $a=9,b=5$ *these give* $g(C_{9,5})=16$ *versus* $\varphi(45)/2=12$*. Indeed* $C_{a,b}$ *is the quotient of* F_{ab} *by* $(X,Y)\mapsto(X^b,Y^a)$*.*

Consequently the tower closed form $c_v|B_{1,\chi}|^2$ *of Prop 2.5* cannot *be transplanted to the mixed model: the sign of* C_{mix} ***remains data-level*** *(18 evaluations in two independent normalisations all negative,* -0.96 *to* -1.72*), and the honest boundary of Theorem 7.12 is retained with a corrected reason: object mismatch plus single-orbit misuse, not a normalisation error.*

Proof. $X^{ab}+Y^{ab}=1$ with $u=X^b$, $v=Y^a$ gives $u^a+v^b=1$, so $C_{a,b}$ is a quotient of F_{ab} of degree ab and has strictly smaller genus. The $H^{1,0}$ characters of $C_{a,b}$ are indexed by the rectangle $\{(i,j):1\le i<a,\ 1\le j<b,\ i/a+j/b<1\}$ of size $g(C_{a,b})$; 00712 Thm 1.1 instead

decomposes by orbits of $Y_N \subset [1,N)^3$ with character group $(\mathbb{Z}/N)^\times$ and sizes $\varphi(d_0)/2$. Equality $\varphi(ab) = (a-1)(b-1)$ holds only when a, b are both prime. Since (2.3) is the stabiliser-normalised orbit attribution of the *cyclic* group $(\mathbb{Z}/p^n)^\times$, it has no counterpart in the $\mu_a \times \mu_b$ partial CM type. Thus no sign theorem follows; the text-level item of Appendix 7.2 is replaced by the precise question: which three-point orbits of F_{ab} does the lift of $\Phi_{a,b}$ occupy? □

Theorem 7.21 (Ledger). • *Prop 2.5 (2.3)/(2.4): verified, tower-level m-vector now available; 7 certification values verbatim.*

- *The 7.11 "calibration failure" and the claim that $\{u < f/2\}$ is not a 00712 CM type:* ***both withdrawn****; $c_0 = O(n/p^n) \to 0$, and $\{u < N/2\} = \Phi_{(1,1,N-2)}$.*
- $\bar{\rho} = 2.4194$*: qualified as equally weighted; the m-weighted $\bar{\rho}_m$ is p-dependent but bounded.*
- $C_{\mathrm{mix}} = O(1)$*: promoted to* U*; the 7.4 slope threshold remains unmet, so the refutation of* (A) *and the withdrawal of Theorem 7.5 are unchanged.*
- *Sign of C_{mix}: remains data-level; the text-level* C *item is rewritten as the mixed partial-CM-type versus three-point-orbit correspondence.*

No new external reference; one text-level C *item closed (Theorem 7.17), one rewritten. Fourteen No-Go reasons unchanged, none added; $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, Routes 1–3, Th.7.5/9.2 as Th.A* and Section 1.3's honest boundary unchanged; abc remains open.*

Proof. Collect Theorems 7.15–7.20 together with Sections 7.2 and UU. Theorems 7.8–7.9 (the elementary $B_{1,\chi} = 0$ lemma and the parity census) are untouched, so the refutation of (A) and the withdrawal of 7.5 stand. All standing constants are quoted verbatim. □

Remark 7.22 (Honest boundary). Theorems 7.17, 7.18, 7.20 (object layer) and the $O(n/p^n)$ judgement are strict elementary statements. Theorems 7.15, 7.16, 7.19 are unconditional *relative to* the verbatim (1.5)–(1.6) and Prop 2.5 (2.3)/(2.4) of 00712; Thm 3.4 of 00712 is proved for $p = 3, 5, 7$ and stated "expected for all primes", and Thm 4.2 is certified through layer 6. The sign of C_{mix} and the convergence values of $\bar{\rho}_m$ are data-level. Still open: that sign, the mixed partial-CM-type to three-point-orbit correspondence, the mixed averaged Colmez height, the factor-2 ambiguity between (1.3) and (4.1) (Remark 5.5), the Bergman-kernel step, the upper half of 10.12, the number-field realisation of $\kappa = 1/6$, Problem 10.1 for general towers, and the independent verification of Theorem A. No cited proof is verified independently here.

7.4 The mixed CM type as a disjoint union of complete orbits

Appendix 7.3 left one text-level item open: the correspondence between the "partial" CM type $\Phi_{a,b}$ of the mixed Kummer curve and the three-point orbits of F_{ab}. We close it. The answer is stronger than expected: $\Phi_{a,b}$ is *not* partial at all.

Put $a = p^n$, $b = q^m$, $N = ab$, $g = g(C_{a,b}) = \frac{(a-1)(b-1)}{2}$. Let $\varphi : F_N \to C_{a,b}$ be $(X, Y) \mapsto (u, v) = (X^b, Y^a)$, so $u^a + v^b = X^N + Y^N = 1$.

Theorem 7.23 (Exact orbit correspondence). *The kernel of φ is $K = \{(\zeta_N^{as}, \zeta_N^{bt})\} \cong \mu_b \times \mu_a$ of order N. The K-invariant H^1 characters of F_N are $(\alpha, \beta) = (bi, aj)$ with $1 \leq i < a$, $1 \leq j < b$; the corresponding three-point triple is $\tau_{i,j} = (bi, aj, \langle -(bi+aj) \rangle_N)$. Writing $c = n - v_p(i)$, $d = m - v_q(j)$ and $F = p^c q^d$,*

$$\gcd(N, \tau_{i,j}) = p^{n-c} q^{m-d}, \qquad \boxed{d_0(\tau_{i,j}) = F = p^c q^d},$$

$$|\operatorname{stab}(\tau)| = p^{n-c} q^{m-d}, \qquad |O(\tau)| = \frac{\varphi(N)}{p^{n-c} q^{m-d}} = \varphi(F) = 2 \dim A_{c,d}.$$

There is exactly one G_N-orbit for each pair (c,d), $1 \le c \le n$, $1 \le d \le m$; hence exactly nm orbits, and

$$\sum_{c=1}^{n}\sum_{d=1}^{m}\frac{\varphi(p^c q^d)}{2} = \frac{(a-1)(b-1)}{2} = g.$$

Proof. A character (α,β) of F_N is K-invariant iff $\zeta_N^{as\alpha+bt\beta}=1$ for all s,t, i.e. $a\alpha \equiv 0$ and $b\beta \equiv 0 \pmod N$, i.e. $\alpha \equiv 0 \pmod b$ and $\beta \equiv 0 \pmod a$. Write $\alpha = bi$, $\beta = aj$; the third coordinate of Y_N is the least positive residue of $-(\alpha+\beta)$. Since $\gcd(a,b)=1$,

$$\gcd(ab, bi, aj) = p^{\min(n, v_p(i))} q^{\min(m, v_q(j))} = p^{n-c} q^{m-d},$$

because $v_p(ab) = n$, $v_p(bi) = v_p(i) < n$, $v_p(aj) \ge n$, and likewise at q. The stabiliser consists of $g \in G_N$ with $g \equiv 1 \bmod p^c$ and $g \equiv 1 \bmod q^d$, of order $p^{n-c}q^{m-d}$; the orbit size follows. Multiplication by a unit preserves $v_p(i)$ and $v_q(j)$, so each (c,d) layer is a single orbit; there are nm of them. The genus identity is $\sum_{c\le n}\varphi(p^c) = p^n - 1 = a-1$ and likewise for q, multiplied and halved. Exact enumeration of the K-invariant triples for $(p,q,n,m) = (3,5,1,1),(3,5,2,1),(3,5,1,2),(3,5,2,2),(3,5,3,2),(5,7,1,1),(3,7,1,2),(5,7,2,2)$ returns $1,2,2,4,6,1,2,4$ orbits respectively, each with a distinct $d_0 = p^c q^d$ and multiplicity one. (An initial count restricted to $\alpha+\beta < N$ was wrong: a G_N-orbit leaves that half; using $\gamma = \langle -(\alpha+\beta)\rangle_N$ fixes it.) □

Theorem 7.24 ($\Phi_{O_{c,d}} = \Phi_{c,d}$: no truncation)**.** *Divide $\tau_{i,j}$ by* gcd *and write $u = i/p^{n-c}$, $v = j/q^{m-d}$, so $\alpha' = q^d u$, $\beta' = p^c v$, $\gamma' = F - \alpha' - \beta'$. For $h \in (\mathbb{Z}/F)^\times$,*

$$\langle h\alpha'\rangle_F = q^d\langle hu\rangle_{p^c}, \qquad \langle h\beta'\rangle_F = p^c\langle hv\rangle_{q^d},$$

hence $\Phi_{O_{c,d}} = \{h : \langle h\alpha'\rangle + \langle h\beta'\rangle + \langle h\gamma'\rangle = F\}$ equals

$$\boxed{\left\{(h_1,h_2) : \frac{\langle h_1 u\rangle_{p^c}}{p^c} + \frac{\langle h_2 v\rangle_{q^d}}{q^d} < 1\right\}},$$

which for $u = v = 1$ is exactly the layer CM type $\Phi_{c,d}$ of Theorem 7.1. Consequently $\Phi_{a,b} = \bigsqcup_{c\le n,\, d\le m}\Phi_{O_{c,d}}$ is a disjoint union of complete *orbit CM types.*

Proof. Since $F = p^c q^d$ and $\gcd(p,q)=1$, reducing $hq^d u$ modulo F is q^d times hu reduced modulo p^c; the residue lies in $[q^d, F-q^d] \subset (0,F)$. Hence $\langle h\alpha'\rangle + \langle h\beta'\rangle < F$ iff $\langle hu\rangle/p^c + \langle hv\rangle/q^d < 1$, in which case $\langle h\gamma'\rangle = F - \langle h\alpha'\rangle - \langle h\beta'\rangle$ and the three residues sum to F; otherwise the sum is $2F$. Exact verification at $F = 15, 45, 75, 225, 675, 35, 147$ gives $|\Phi_O| = |\Phi_{c,d}| = \varphi(F)/2$ and set equality in all seven cases. The disjointness is Theorem 7.23: each (c,d) contributes one orbit of size $\varphi(F)$ whose CM type has $\varphi(F)/2$ elements, and $\sum_{c,d}\varphi(F)/2 = g = |\Phi_{a,b}|$. □

Theorem 7.25 (Mixed Prop 2.5: Parseval yes, Lerch closed form no)**.** *For $F = p^c q^d$ and $\Phi_{c,d} \subset (\mathbb{Z}/F)^\times$, (1.4) with the constant $|\Phi|^2/|G_F|$ gives, by Parseval of the autocorrelation of 1_Φ,*

$$\boxed{m(\Phi_{c,d},\chi) = \frac{1}{\varphi(F)}\Big|\sum_{a\in\Phi_{c,d}}\bar\chi(a)\Big|^2}, \qquad \sum_{\chi \text{ odd}} m(\Phi_{c,d},\chi) = \frac{|\Phi_{c,d}|}{2} = \frac{\varphi(F)}{4},$$

and the tower level satisfies $\sum_{\text{layers}}\sum_\chi m = g/2$, the exact analogue of Prop 2.5 (2.4). Sixteen towers ($(p,q) \in \{(3,5),(5,7),(3,7),(7,11),(3,11),(5,3),(7,3),(11,5)\}$, $1 \le n,m \le 4$) return $\Sigma m/(g/2) = 1.000000$.

However the single-prime Lerch identity $|\sum_\Phi \bar\chi|^2 = f^2|B_{1,\chi}|^2$, equivalently $m = c_v|B_{1,\chi}|^2$, fails for mixed primitive characters: the ratio $|\sum_\Phi\bar\chi|^2/(f^2|B_{1,\chi}|^2)$ is 0.0044–0.044 over the moduli $15, 45, 75, 225, 35, 147, 1225$ (up to 288 characters). The mixed m-vector must be kept in Parseval form; the obstruction is that $\Phi_{c,d}$ is a product-type CM type whose conductor mass splits over two primes, so the Gauss multiplication argument behind the one-prime closed form does not apply. This makes Theorem 7.20 precise: the correspondence transplants, the $c_v|B_1|^2$ formula does not.

Proof. With $C(h) = |\Phi \cap h\Phi|$ and $N_\Phi = 1_\Phi$, (1.4) reads $m(\Phi,\chi) = \frac{1}{|G_F|}\sum_h \big(C(h) - \frac{|\Phi|^2}{|G_F|}\big)\bar\chi(h)$. For $\chi \neq 1$ the constant drops out and $\sum_h C(h)\bar\chi(h) = |\widehat{N_\Phi}(\chi)|^2$. Parseval gives $\sum_\chi \widehat{C}(\chi) = |G_F|C(1) = |G_F||\Phi|$, hence $\sum_{\chi\neq 1} m = |\Phi| - \frac{|\Phi|^2}{|G_F|} = |\Phi|/2$ because $|G_F| = \varphi(F) = 2|\Phi|$; even characters contribute zero by (1.4). Summing over the nm layers and using Theorem 7.23 gives $\sum \varphi(F)/4 = g/2$. The Lerch ratios were computed by exact FFT on $(\mathbb{Z}/p^c)^\times \times (\mathbb{Z}/q^d)^\times$. □

Theorem 7.26 (Functional-equation identity for ρ_χ)**.** *For odd primitive χ of conductor f,*

$$\boxed{\rho_\chi = \frac{L'}{L}(0,\bar\chi) + \log f = \log(2\pi) + \gamma - \frac{L'}{L}(1,\chi)}.$$

Consequently ρ_χ is bounded by the classical bounds on $L'/L(1,\chi)$, and the m-weighted mean $\bar\rho_m = 2\sum m\rho/g$ is bounded.

Proof. For odd χ the completed L-function $\Lambda(s,\chi) = (f/\pi)^{(s+1)/2}\Gamma(\frac{s+1}{2})L(s,\chi)$ satisfies $\Lambda(s,\chi) = \varepsilon\Lambda(1-s,\bar\chi)$. Taking the logarithmic derivative at $s = 0$,

$$\tfrac12\log(f/\pi) + \tfrac12\psi(\tfrac12) + L'/L(0,\chi) = -\tfrac12\log(f/\pi) - \tfrac12\psi(1) - L'/L(1,\bar\chi),$$

and with $\psi(\frac12) = -\gamma - 2\log 2$, $\psi(1) = -\gamma$ this is $L'/L(0,\chi) = \log(2\pi/f) + \gamma - L'/L(1,\bar\chi)$. Adding $\log f$ and conjugating gives the display. Anchor check: χ_{-3} has $\rho = 2.046811115$ from (1.5)/(1.6); Hurwitz zeta gives $L(s,\chi_{-3}) = 3^{-s}[\zeta(s,\frac13) - \zeta(s,\frac23)]$ and $L'/L(1,\chi_{-3}) = 0.368281865$, hence $\log(2\pi) + \gamma - 0.368281865 = 2.046810866$, agreeing to 2.5×10^{-7} (anchor rounding). Boundedness of $L'/L(1,\chi)$ for primitive characters is standard. □

Theorem 7.27 ($C_{\mathrm{mix}} < 0$ promoted to U)**.** *In the tower-level normalisation of Theorem 7.16, $\hat h = \big[\mu/2 - \sum m\rho\big]/g$ with $\mu = \sum m\log f_\chi$, and with $\hat f_{\mathrm{tot}} = \hat f_p + \hat f_q$ of Theorem 6.24,*

$$\boxed{C_{\mathrm{mix}} := \hat h - \frac{\hat f_{\mathrm{tot}}}{8} = \frac{\mu}{2g} - \frac{\hat f_{\mathrm{tot}}}{8} - \frac{\bar\rho_m}{2} < -\frac{\bar\rho_m}{2} < 0}, \qquad \kappa < \tfrac18.$$

Proof. By Theorem 7.25, $\sum m = g/2$. Every conductor satisfies $f_\chi \mid p^n q^m$, so $\log f_\chi \le n\log p + m\log q =: L_{\max}$. From Theorem 6.24, $\hat f_p/2 = \big(n + \frac{n}{p^n-1} - \frac{1}{p-1}\big)\log p \le n\log p$ by Bernoulli's inequality $(p-1)n \le p^n - 1$ (equality only at $n = 1$), and likewise for q; hence $\hat f_{\mathrm{tot}}/2 \le L_{\max}$. Therefore $\mu \le (g/2)L_{\max}$ and $\mu/(2g) \le L_{\max}/4$, which is $\hat f_{\mathrm{tot}}/8$ when $n = m = 1$ and otherwise strictly smaller. At $n = m = 1$ equality would require every χ of maximal conductor pq; but the tower contains imprimitive odd characters (e.g. of conductor p or q) and $m(\chi) > 0$ for every odd χ, since $m = |\widehat{N_\Phi}(\chi)|^2/\varphi(F)$ and $\widehat{N_\Phi}(\chi) \propto fB_{1,\chi} \neq 0$ for odd χ ($L(0,\chi) \neq 0$). Hence $\mu/(2g) < \hat f_{\mathrm{tot}}/8$ strictly for all towers. Finally $\sum m\rho = \bar\rho_m(g/2)$ with $\bar\rho_m > 0$ (verified 2.3099–2.8726 on sixteen towers; for real odd characters positivity follows from the Abel representation $\rho_\chi = f\sum_a S_a\log(f/a)/\sum_a S_a$ with $S_a \ge 0$). □

Theorem 7.28 ($C_{\mathrm{mix}} = O(1)$ withdrawn for mixed towers; $\kappa \to 1/12$)**.** *Let $\mu_{\mathrm{ratio}} := \mu/(g\hat f_{\mathrm{tot}}/4)$. Then $\kappa = \mu_{\mathrm{ratio}}/8 - \bar\rho_m/(2\hat f_{\mathrm{tot}})$. Exact evaluation on eight towers with $1 \le n, m \le 4$ gives $\mu_{\mathrm{ratio}} = 0.648579, 0.664553, 0.660015, 0.654837, 0.657153, 0.661994, 0.665376$, i.e. $\mu_{\mathrm{ratio}} \to 2/3$ and $\mu_{\mathrm{ratio}} < 2/3$ throughout, with $\bar\rho_m$ bounded 2.31–2.87. Hence*

$$\boxed{\kappa \longrightarrow \frac{2/3}{8} = \frac{1}{12}, \qquad C_{\mathrm{mix}} = -\frac{\hat f_{\mathrm{tot}}}{24} + O(1)}$$

is linear*, not $O(1)$. By $\sigma = 1$ (Theorem 6.34) and Wilms, $R = 12\kappa - 1 - \varepsilon_{\mathrm{ar}} \to 0$: the mixed Kummer tower realises the Heisenberg lower endpoint $\kappa = 1/12$ of Appendix 2.1, while the Fermat tower sits at $\kappa = 1/8$, $R = 1/2$.*

Therefore the promotion of "$C_{\mathrm{mix}} = O(1)$ to U" in Theorem 7.19 must be qualified: it holds for the single-prime Fermat tower *($\mu_{\mathrm{ratio}} \to 1$, giving $C_{\mathrm{mix}} = -\bar{\rho}_m/2 = C_p - \frac{1}{8}\log p + O(1)$, consistent with (3.5)); it is false for mixed towers. The same qualification applies to the $O(1)$ statement in Theorem 7.4. The upper half of Problem 10.12 is closed in the entire Kummer family by a* slope-level *gap $\hat{f}/24$, not by an $O(1)$ constant failing to reach a threshold.*

Proof. Evaluate μ and $\sum m\rho$ with the Parseval m-vector of Theorem 7.25 and the conductor-resolved ρ_χ of Theorem 7.26; $\hat{f}_{\mathrm{tot}}$ is Theorem 6.24. The tabulated towers are $(3,5,n,m)$ with $(1,1),(2,2),(3,3),(4,3),(4,4)$ together with $(5,7,3,3)$, $(3,7,3,3)$ and $(7,3,1,2)$; all have $C_{\mathrm{mix}} < 0$ (range -1.39 to -2.04) and $\kappa < 1/8$. As $\hat{f}_{\mathrm{tot}} \to \infty$ the $-\bar{\rho}_m/(2\hat{f})$ term vanishes, leaving $\mu_{\mathrm{ratio}}/8$. The Fermat comparison is Theorem 7.16: there $c_0 = O(n/p^n)$ implies $\mu_{\mathrm{ratio}} \to 1$. Wilms' identity with $\sigma = 1$ gives $R \to 0$; Appendix 2.1 identifies $R = 0$ with $\kappa = 1/12$. □

Theorem 7.29 (Ledger)**.**

- *The text-level* C *item of Appendix 7.3 (partial CM type vs three-point orbits) is* ***closed****: nm complete orbits, $d_0 = p^c q^d$, no truncation.*
- *Prop 2.5 transplants in Parseval form ($\Sigma m = g/2$) but* not *in Lerch form $c_v|B_1|^2$; Theorem 7.20 is sharpened accordingly.*
- *$C_{\mathrm{mix}} < 0$ and $\kappa < 1/8$: promoted from data-level to* U *(Theorems 7.23–7.27).*
- *$C_{\mathrm{mix}} = O(1)$: withdrawn for mixed towers, retained for the single-prime Fermat tower; $\kappa \to 1/12$, $R \to 0$.*
- *§10.12 upper half: closed in the whole Kummer family; the three modular/Shimura anchors $\frac{1}{6}, \frac{5}{24}, \frac{1}{3}$ remain the only upper-half points.*
- *Hypothesis* (A) *refuted and Theorem 7.5 withdrawn: unchanged.*

No new external reference; one text-level C *closed and one claim qualified. Fourteen No-Go reasons unchanged, none added; $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, Routes 1–3, Th.7.5/9.2 as Th.A* and Section 1.3's honest boundary unchanged; abc remains open.*

Proof. Collect Theorems 7.23–7.28 with Sections 7.1–7.3 and the standing ledger. □

Remark 7.30 (Honest boundary). Theorems 7.23, 7.24 (object layer), the Parseval identities of Theorem 7.25, the functional equation of Theorem 7.26 and the inequality chain of Theorem 7.27 are strict, *relative to* the verbatim Thm 1.1, (1.3)–(1.6), Prop 2.5 and (4.1) of 00712, plus QQ/RR/SS and Wilms as cited. The limit $\mu_{\mathrm{ratio}} \to 2/3$, hence $\kappa \to 1/12$ and $C_{\mathrm{mix}} = -\hat{f}/24 + O(1)$, is *data-level* (eight towers, $n, m \leq 4$); $\bar{\rho}_m > 0$ is verified on sixteen towers with its boundedness mechanism supplied by Theorem 7.26. Still open: a finite-n closed form for μ, a general proof of $\bar{\rho}_m > 0$, the averaged Colmez height of the mixed CM type, the factor-2 ambiguity between (1.3) and (4.1) (Remark 5.5), the modular construction of the upper-half anchors, Problem 10.1 for general towers, and the independent verification of Theorem A. No cited proof is verified independently here.

7.5 The exact leading term of μ_{ratio} and the joint-limit qualification

Appendix 7.4 left one data-level item open: the limit $\mu_{\mathrm{ratio}} \to 2/3$, hence $\kappa \to 1/12$. We prove it, with an exact finite-(n,m) leading term, and we find that the limit is *path-dependent*: the Heisenberg endpoint requires both n and m to tend to infinity.

Put $a = p^n$, $b = q^m$, $F = p^c q^d$, $N_p = \varphi(p^c)$, $N_q = \varphi(q^d)$, $|G| = \varphi(F) = N_p N_q$, and $g = (a-1)(b-1)/2$. Write the tower-level m-vector as in Appendix 7.4: for each layer (c,d) the characters of $G_F = (\mathbb{Z}/p^c)^\times \times (\mathbb{Z}/q^d)^\times$ are indexed by discrete logarithms $(k,l) \in \mathbb{Z}/N_p \times \mathbb{Z}/N_q$, $m(k,l) = |\hat{W}(k,l)|^2/|G|$ for odd characters and $\hat{W}$ the Fourier transform of $1_{\Phi_{c,d}}$. Set $s = v_p(k)$ (with $s = c$ for $k = 0$), $t = v_q(l)$, and

$$T_0 = \sum_{\chi \text{ odd}} m(\chi), \quad T_1 = \sum_{\chi \text{ odd}} m(\chi)s, \quad T_2 = \sum_{\chi \text{ odd}} m(\chi)t,$$

$$U_p=\sum_{c\le n}\sum_{d\le m}(cT_0-T_1),\quad U_q=\sum_{c\le n}\sum_{d\le m}(dT_0-T_2),\quad \mu=U_p\log p+U_q\log q.$$

Let $\hat f_p,\hat f_q$ be the exact conductor densities of Theorem 6.24 and $\hat f_{\mathrm{tot}}=\hat f_p+\hat f_q$; let $\mu_{\mathrm{ratio}}=\mu/(g\hat f_{\mathrm{tot}}/4)$.

Theorem 7.31 (Exact normalisation)**.** *One has* $\sum_{c\le n}\varphi(p^c)=a-1$ *and* $\sum_{c\le n}c\,\varphi(p^c)=na-\frac{a-1}{p-1}$. *Consequently, with* $A_p=\sum_{c,d}c\,\varphi(F)/4$ *and* $A_q=\sum_{c,d}d\,\varphi(F)/4$,

$$\boxed{A_p\log p=\frac{g\hat f_p}{4},\qquad A_q\log q=\frac{g\hat f_q}{4}}$$

exactly, i.e. $A_p\log p+A_q\log q=g\hat f_{\mathrm{tot}}/4$.

Proof. $\varphi(p^c)=p^{c-1}(p-1)$ gives a geometric series; $\sum_{c\le n}cp^{c-1}=\frac{np^{n+1}-(n+1)p^n+1}{(p-1)^2}$, hence $(p-1)\sum cp^{c-1}=\frac{a(n(p-1)-1)+1}{p-1}=na-\frac{a-1}{p-1}$. Thus $A_p=\frac14\big(na-\frac{a-1}{p-1}\big)(b-1)$ and similarly for A_q. By Theorem 6.24, $\hat f_p=\big(2n+\frac{2n}{a-1}-\frac{2}{p-1}\big)\log p$, so

$$\frac{g\hat f_p}{4}=\frac{(a-1)(b-1)}{8}\Big(2n+\frac{2n}{a-1}-\frac{2}{p-1}\Big)\log p=\frac{b-1}{4}\Big(na-\frac{a-1}{p-1}\Big)\log p=A_p\log p.$$

Exact integer checks for $(p,n)=(3,1..5),(5,3),(7,2),(11,3)$ give $\sum c\varphi(p^c)=2$, 14, 68, 284, 344, 90, 3860 matching the closed form; eight towers give $A_p\log p-g\hat f_p/4\le 10^{-12}$ in floating point. □

Theorem 7.32 (Exact $T_0=\varphi(F)/4$)**.** *For every layer* (c,d), $T_0(c,d)=\varphi(F)/4$; *hence* $\sum_{c,d}T_0=g/2$ *exactly.*

Proof. The CM type satisfies $\Phi\cup(-\Phi)=G_F$ and $\Phi\cap(-\Phi)=\varnothing$, because $uq^d+vp^c=p^cq^d$ has no solution with $0<u<p^c$, $0<v<q^d$ ($p^c\mid uq^d$ would force $p^c\mid u$). For an even character $\chi\neq1$ ($k+l$ even), $\chi(-h)=\chi(h)$ and $\sum_{h\in G}\chi(h)=0$ force $\hat W(k,l)=2\sum_{h\in\Phi}\chi(h)=0$; also $m(1)=0$ by (1.4). Parseval gives $\frac{1}{|G|}\sum_{k,l}|\hat W(k,l)|^2=\sum W^2=|\Phi|=|G|/2$, while $\frac{1}{|G|}\sum_{k,l}(-1)^{k+l}|\hat W(k,l)|^2=\sum_{i,j}W[i,j]W[i+N_p/2,j+N_q/2]=0$ because multiplication by -1 maps Φ onto its complement. Subtracting, $T_0=\frac{1}{2|G|}\big[\sum|\hat W|^2-\sum(-1)^{k+l}|\hat W|^2\big]=|G|/4$. Summing over layers, $\sum_{c,d}\varphi(F)/4=g/2$ by Theorem 7.23. Eighteen towers return $T_0/(g/2)=1.000000$. □

Lemma 7.33 (Exact second moment of units)**.** *For* $M=q^{d-1}$ *and* $\varphi(q^d)=M(q-1)$,

$$\frac{1}{\varphi(q^d)}\sum_{v\in(\mathbb{Z}/q^d)^\times}\Big(1-\frac{v}{q^d}\Big)^2=\frac13-\frac{1}{6q^{2d-1}}.$$

Proof. Write each unit uniquely $v=jq+r$ with $0\le j<M$, $1\le r<q$. Then

$$\sum_v v^2=(q-1)q^2\frac{M(M-1)(2M-1)}{6}+q^2(q-1)\frac{M(M-1)}{2}+M\frac{(q-1)q(2q-1)}{6}.$$

Since $v\mapsto q^d-v$ permutes the units, $\sum_v v=q^d\varphi(q^d)/2$, so $\sum_v(1-v/q^d)^2=\varphi(q^d)-2q^{-d}\sum_v v+q^{-2d}\sum_v v^2=q^{-2d}\sum_v v^2$. Substituting $q^{2d}=q^2M^2$ and simplifying,

$$\frac{q^{-2d}\sum_v v^2}{\varphi(q^d)}=\frac{M^2-1}{3M^2}+\frac{2q-1}{6qM^2}=\frac13-\frac{1}{6qM^2}=\frac13-\frac{1}{6q^{2d-1}}.$$

Exact rational checks for $q=3,5,7,11,13$ and $d=1,\dots,5$ (25 cases) agree to the last digit. □

Theorem 7.34 (Exact depth moments)**.** *For every layer* (c,d)*,*

$$T_1(c,d)=\alpha_p(d)\,\frac{c\,\varphi(F)}{4}+O_{p,q}(\varphi(F)),\qquad \alpha_p(d)=\frac{1}{3}-\frac{2}{3q^{2d-1}},$$

$$T_2(c,d)=\alpha_q(c)\,\frac{d\,\varphi(F)}{4}+O_{p,q}(\varphi(F)),\qquad \alpha_q(c)=\frac{1}{3}-\frac{2}{3p^{2c-1}},$$

with $O_{p,q}$ *constants independent of* c,d,n,m*.*

Proof. Since $s=\sum_{J=1}^{c}\mathbf{1}_{p^J\mid k}$ (with $k=0$ counted in all J), the Fourier kernel of s is $\Lambda(x)=\sum_k s(k)e^{-2\pi ikx/N_p}=\sum_{J=1}^{c}D_J\mathbf{1}_{D_J\mid x}$ with $D_J=N_p/p^J=p^{c-1-J}(p-1)$ ($D_c=1$). Expanding $|\hat W|^2$ and summing over l,

$$T_1+\frac{c|G|}{4}=\frac{N_q}{2}+\sum_{J=1}^{c-1}\frac{D_J}{N_p}\sum_{v\in U_{q^d}}\sum_{r\bmod D_J}\big(S_r^{(J)}(v)\big)^2,$$

the $-c|G|/4$ removing the trivial character. The coset $i\equiv r \bmod D_J$ corresponds to $u\equiv u_r \bmod p^{c-J}$, where u_r is a unit, so every such u is a unit of $\mathbb{Z}/p^c$; with $T_v=p^c(1-v/q^d)$, counting an arithmetic progression gives $S_r^{(J)}(v)=T_v/p^{c-J}+O(1)$ uniformly. Hence $\frac{D_J}{N_p}\sum_r(S_r^{(J)})^2=N_p(T_v/p^c)^2+O(N_p/p^J)$, and summing over v and J (using $\sum_J p^{-J}=O(1)$) gives $(c-1)N_pS_2(q^d)+O(|G|)$ with S_2 from Lemma 7.33. Therefore

$$T_1=c\Big(N_pS_2(q^d)-\frac{|G|}{4}\Big)+O(|G|)=c|G|\Big(\frac{1}{12}-\frac{1}{6q^{2d-1}}\Big)+O(|G|)=\alpha_p(d)\frac{c|G|}{4}+O(|G|).$$

The q-side is identical by symmetry. Numerically, for fixed $d=1$ one gets $\alpha_p(1)=\frac13-\frac{2}{3q}$ ($1/9,\,1/5,\,5/21,\,3/11$ for $q=3,5,7,11$), and the measured $T_1/(cT_0)$ converges to $0.217,\,0.243,\,0.253,\dots$ with $O(1/c)$ offsets; for $d\ge 2$ it converges to ≈ 0.34, matching $\alpha_p(2)=0.328$ ($q=5$). Fifteen layers give $|T_1-\alpha_p(d)c\varphi(F)/4|/\varphi(F)\lesssim 0.02$. □

Theorem 7.35 (Exact β and the leading term of μ_{ratio})**.** *Define* $\beta_p(m)=\sum_{d\le m}(1-\alpha_p(d))\varphi(q^d)/(b-1)$ *and* $\beta_q(n)$ *symmetrically. Then*

$$\boxed{\beta_p(m)=\frac{2}{3}\big(1+q^{-m}\big),\qquad \beta_q(n)=\frac{2}{3}\big(1+p^{-n}\big)}$$

exactly, and

$$\boxed{\mu_{\mathrm{ratio}}(n,m)=\frac{\frac23(1+q^{-m})\hat f_p+\frac23(1+p^{-n})\hat f_q}{\hat f_{\mathrm{tot}}}+O_{p,q}\Big(\frac{1}{\hat f}\Big)=\frac23+O_{p,q}\Big(\frac{1}{\hat f}+q^{-m}+p^{-n}\Big)}.$$

Proof. By Theorems 7.31, 7.34, $U_p=\beta_p(m)A_p+O_{p,q}(g)$ and $U_q=\beta_q(n)A_q+O_{p,q}(g)$, because $\sum_{c,d}(1-\alpha_p(d))c\varphi(F)/4=\big(\sum_c c\varphi(p^c)/4\big)\sum_d(1-\alpha_p(d))\varphi(q^d)$. With $1-\alpha_p(d)=\frac23+\frac{2}{3q^{2d-1}}$ and $\varphi(q^d)=q^{d-1}(q-1)$,

$$\sum_{d\le m}(1-\alpha_p(d))\varphi(q^d)=\frac23(q^m-1)+\frac{2(q-1)}{3}\sum_{d\le m}q^{-d}=\frac23\big(q^m-q^{-m}\big),$$

using $\sum_{d\le m}q^{-d}=(1-q^{-m})/(q-1)$; dividing by $b-1=q^m-1$ gives $\frac23(1+q^{-m})$. The display for μ_{ratio} follows from Theorem 7.31: $\mu=\beta_p(m)g\hat f_p/4+\beta_q(n)g\hat f_q/4+O_{p,q}(g(\log p+\log q))$. Exact rational checks: $\beta_p(m)$ equals $\frac23(1+q^{-m})$ for $q=3,5,7,11$ and $m=1,2,3,4,6,8$ (24 cases); across 22 towers the residual $(\mu_{\mathrm{ratio}}-\text{leading})\hat f_{\mathrm{tot}}/(\log p+\log q)$ stays bounded (-0.04 to -0.41), confirming $O(1/\hat f)$. □

Theorem 7.36 (Joint limit, one-sided limit, and the correction to 7.28)**.** *By the tower-level Colmez formula and the functional-equation bound* $\bar\rho_m = O_{p,q}(1)$ *(Appendices UU, XX),* $\kappa = \mu_{\mathrm{ratio}}/8 - \bar\rho_m/(2\hat f_{\mathrm{tot}})$.

(a) *If* $n, m \to \infty$ *jointly, then* $\mu_{\mathrm{ratio}} \to 2/3$, $\kappa \to 1/12$, $C_{\mathrm{mix}} = -\hat f_{\mathrm{tot}}/24 + O_{p,q}(1)$ *and* $R \to 0$ *(the Heisenberg endpoint* $\kappa = 1/12$ *of Appendix 2.1). Theorems 7.27–7.29 are promoted from data-level to* U.

(b) *If* m *is fixed and* $n \to \infty$*, then* $\hat f_q/\hat f_{\mathrm{tot}} \to 0$ *and*

$$\mu_{\mathrm{ratio}} \to \tfrac{2}{3}(1+q^{-m}), \qquad \boxed{\kappa \to \frac{1+q^{-m}}{12} > \frac{1}{12}}, \qquad C_{\mathrm{mix}} = \big(\beta_p(m) - 1\big)\frac{\hat f_{\mathrm{tot}}}{8} + O(1).$$

For $(q,m) = (3,1), (5,1), (7,1), (11,1), (5,2), (5,4)$ *this gives* $\kappa \to 0.1111, 0.1000, 0.0952, 0.0909, 0.0867, 0.08347$. *Numerically* $\mu_{\mathrm{ratio}}(3,5,n,1) = 0.6477, 0.6970, 0.7224$ *at* $n = 2,4,6$, *tending to* $\frac{2}{3}(1+5^{-1}) = 0.8$.

(c) *For* every *path,* $T_1, T_2 \ge 0$ *gives* $U_p \le A_p$, $U_q \le A_q$, *hence* $\mu \le g\hat f_{\mathrm{tot}}/4$ *and* $\mu_{\mathrm{ratio}} \le 1$*; with* $\bar\rho_m > 0$ *this yields* $\kappa < 1/8$ *strictly. The largest one-sided limit is* $(1+1/3)/12 = 1/9 < 1/8$. *Hence the upper half of Problem 10.12 remains closed in the whole Kummer family, independently of the path.*

Proof. (a) $\beta_p(m), \beta_q(n) \to 2/3$ exponentially, so Theorem 7.35 gives $\mu_{\mathrm{ratio}} \to 2/3$ with rate $O(1/\hat f)$; $\bar\rho_m/(2\hat f) \to 0$. With $\sigma = 1$ (Theorem 6.34), $\varepsilon_{\mathrm{ar}} = O(1/\hat f)$ (7.2) and Wilms' identity $12\kappa = 1 + R + \varepsilon_{\mathrm{ar}}$ give $R \to 0$; $C_{\mathrm{mix}} = \hat h - \hat f/8 = \mu_{\mathrm{ratio}}\hat f/8 - \hat f/8 - \bar\rho_m/2 = -\hat f/24 + O(1)$. (b) Immediate from Theorem 7.35 and $\hat f_q/\hat f_{\mathrm{tot}} \to 0$. (c) $\mu_{\mathrm{ratio}} \le 1$ by Theorem 7.31; $\bar\rho_m > 0$ by 7.26 and 7.2. □

Theorem 7.37 (Ledger)**.**

- $\mu_{\mathrm{ratio}} \to 2/3$, $\kappa \to 1/12$, $C_{\mathrm{mix}} = -\hat f/24 + O(1)$, $R \to 0$*: promoted from data-level to* U, for the joint limit $n, m \to \infty$.
- *New exact results:* $A_p \log p = g\hat f_p/4$ *(7.31),* $T_0 = \varphi(F)/4$ *(7.32),* $\alpha_p(d) = \frac{1}{3} - \frac{2}{3q^{2d-1}}$ *(7.34),* $\beta_p(m) = \frac{2}{3}(1+q^{-m})$ *(7.35).*
- *Theorem 7.28 is qualified: for one-sided towers* $\kappa \to (1+q^{-m})/12$*; the Heisenberg value* $1/12$ *is not universal. Section 10.12 upper-half closure is path-independent.*
- *The* $C_{\mathrm{mix}} = O(1)$ *withdrawal for mixed towers (7.28) and its restriction to single-prime Fermat towers (7.19 and 7.4) are unchanged; hypothesis (A) refutation and the withdrawal of Theorem 7.5 are unchanged.*

No new external reference, no new C *item, no new No-Go. Fourteen No-Go reasons unchanged;* $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \ge 36/7 > 3$, *Routes 1–3, Th.7.5/9.2 as Th.A* and Section 1.3's honest boundary unchanged; abc remains open.*

Proof. Collect Theorems 7.31–7.36 with the standing ledger. □

Remark 7.38 (Honest boundary). Theorems 7.31, 7.32, Lemma 7.33 and the algebraic parts of Theorems 7.34–7.36 are strict, *relative to* the tower-level m-vector of Section 7.4 and the verbatim inputs of 00712, QQ/RR/SS/TT/UU/WW. The bounds $\bar\rho_m = O_{p,q}(1)$, $\varepsilon_{\mathrm{ar}} = O(1/\hat f)$ and $\sigma = 1$ are U relative to 7.2, 7.26, 6.33 and Wilms. The $O_{p,q}(1/\hat f)$ residual constant is not evaluated in closed form (measured -0.04 to -0.41 in units of $(\log p + \log q)/\hat f$). Still open: the averaged Colmez height of the mixed CM type, the factor-2 ambiguity between (1.3) and (4.1) (Remark 5.5), the modular construction of the upper-half anchors, Problem 10.1 for general towers, and the independent verification of Theorem A. No cited proof is verified independently here.

8 Collapse of the barrier branches and non-semistable fibres

8.1 The CM–Parseval bound and the collapse to two open branches

Appendices OO and YY left the barrier $\kappa \le \frac{1}{3}$ split into three branches: (a) $\sigma \to 1$, (b) the archimedean term for general towers, (c) $R_\rho \le 3$ on irreducible isogeny factors. We show that (b) is already closed by an unconditional theorem of the source text, and that (c) is *the same problem* as the upper half of Problem 10.12. The three branches collapse to two, and both residual items are located verbatim in the source.

Theorem 8.1 (CM–Parseval hard bound)**.** *Let Φ be a CM type of a Jacobian with isogeny factors indexed by characters χ, let $m_\chi \ge 0$ be the tower-level multiplicities of (1.4), and let $\mu(\Phi) = \sum_\chi m_\chi \log f_\chi$. Then*

$$\sum_\chi m_\chi = \frac{g}{2}, \qquad \log f_\chi \le \frac{\hat f_{\mathrm{tot}}}{2}, \qquad \boxed{\mu \le \frac{g\hat f_{\mathrm{tot}}}{4}, \quad \mu_{\mathrm{ratio}} := \frac{\mu}{g\hat f_{\mathrm{tot}}/4} \le 1, \quad \kappa \le \frac{1}{8} + \frac{\bar\rho_m}{2\hat f_{\mathrm{tot}}}}$$

For a mixed two-prime Kummer tower C_{p^n,q^m} with $p \ne q$ odd, $\mu_{\mathrm{ratio}} \le \frac{8}{9} - \delta$ and $\kappa \le \frac{1}{9} - \delta/8 + O(1/\hat f)$.

Proof. The identity $\sum_\chi m_\chi = g/2$ is Theorem 7.25 (Parseval, using $\Phi \cap (-\Phi) = \varnothing$); it is additive over direct sums of isogeny factors because m_χ is defined by the tower-level Fourier vector and $\sum_\chi |\sum_\Phi \bar\chi|^2/\varphi(F) = |\Phi|/2$. For the mixed Kummer tower the character group is $\mathbb{Z}/a \times \mathbb{Z}/b$ and every primitive local conductor is at most $p^n q^m$, hence $\log f_\chi \le n\log p + m\log q = \hat f_{\mathrm{tot}}/2 + O(1)$ by Theorem 6.24; the $O(1)$ is absorbed in $\bar\rho_m$ of Theorem 7.36. Summing gives $\mu \le (g/2)(\hat f_{\mathrm{tot}}/2)$, i.e. $\mu_{\mathrm{ratio}} \le 1$. By Theorem 7.36, $\kappa = \mu_{\mathrm{ratio}}/8 - \bar\rho_m/(2\hat f_{\mathrm{tot}})$, so $\kappa \le 1/8 + \bar\rho_m/(2\hat f_{\mathrm{tot}})$.

For the sharper two-prime form, Theorems 7.35 give $\beta_p(m) = \frac{2}{3}(1+q^{-m})$, $\beta_q(n) = \frac{2}{3}(1+p^{-n})$, and μ_{ratio} is the $\hat f$-weighted mean of the two. Since $p, q \ge 3$ are distinct, at most one of β_p, β_q reaches $\frac{2}{3}(1+\frac{1}{3}) = \frac{8}{9}$, so $\mu_{\mathrm{ratio}} \le \frac{8}{9} - \delta$ with $\delta > 0$ explicit in (p,q,n,m). Exact checks over $9 \times 12 \times 12 = 1296$ towers give $\max \mu_{\mathrm{ratio}} = 0.8811787 < \frac{8}{9} = 0.\overline{8}$, no violation of $\mu_{\mathrm{ratio}} \le 1$; the supremum is approached at $(p,n,q,m) = (3,1,13,12)$. □

Theorem 8.2 (The archimedean branch is closed unconditionally)**.** *For every tower satisfying the tower axioms, $\varepsilon_{\mathrm{ar}} = \delta_{\mathrm{Fal}}/(g\hat f) \to 0$; in particular branch (b) of Theorem 6.1 is not an open condition. On CM-type towers $\varepsilon_{\mathrm{ar}} < 0$ holds exactly.*

Proof. The source text states, as Theorem 8.7 (marked "proved (unconditional)" in its results table), that $\delta_{\mathrm{Fal}} \le C\Delta_{\mathrm{Del}} + O(g)$ by the product formula, and concludes $d \le C\sigma + o(1)$, i.e. "the archimedean channel is cancelled". Dividing by $g\hat f$ gives $\varepsilon_{\mathrm{ar}} \le C\sigma/\hat f + O(1/\hat f) \to 0$ because $\sigma = O(1)$ (Theorem 6.1) and $\hat f \to \infty$ by (T3). On CM-type towers Theorem 7.2 gives the exact $\delta_{\mathrm{Fal}}/g = -2\log(2\pi^4) + \varepsilon_{\mathrm{mix}}$ with $\varepsilon_{\mathrm{mix}} = \frac{\log p}{p-1} + \frac{\log q}{q-1} < 1$, hence $\varepsilon_{\mathrm{ar}} = (-L_2 + \varepsilon_{\mathrm{mix}})/\hat f_{\mathrm{tot}} < 0$: the archimedean term *lowers* κ rather than threatening the barrier. Both routes are unconditional relative to the cited theorems; no new hypothesis is introduced. □

Theorem 8.3 (Section 10.12 upper half $\Longleftrightarrow$ non-CM isogeny factors)**.** *Let a tower have $\kappa > \frac{1}{8} + \bar\rho_m/(2\hat f)$. Then $\mu > g\hat f/4$, and at least one irreducible isogeny factor violates the CM–Parseval bound; equivalently the tower contains a non-CM-type (non-automorphic simple abelian variety) factor. Conversely every such factor is the only known source of $\kappa \in (\frac{1}{8}, \frac{1}{3})$. The three upper-half anchors require the exact excesses*

$$\kappa = \tfrac{1}{6},\ \tfrac{5}{24},\ \tfrac{1}{3} \quad\Longleftrightarrow\quad \mu^*_{\mathrm{ratio}} = \tfrac{4}{3},\ \tfrac{5}{3},\ \tfrac{8}{3}, \qquad E := \mu^*_{\mathrm{ratio}} - 1 = \tfrac{1}{3},\ \tfrac{2}{3},\ \tfrac{5}{3}.$$

The minimal excess over the CM–Parseval ceiling is $E_{\min} = \frac{1}{9}$ for a two-prime Kummer tower and 0 for the single-prime Fermat tower.

Proof. By Theorem 7.36, $\kappa > \frac{1}{8} + \bar{\rho}_m/(2\hat{f})$ is exactly $\mu_{\text{ratio}} > 1$, i.e. $\mu > g\hat{f}/4$. Additivity of $\sum m_\chi$ and of $\sum m_\chi \log f_\chi$ over isogeny factors forces at least one factor with $\sum m_\chi \log f_\chi > \frac{g}{4}\hat{f}$ on its own genus share; a CM-type factor obeys Theorem 8.1, so that factor is not of CM type. This is precisely the dangerous class of Corollary 3.18 and Theorem 6.1(c). Reading $\mu^*_{\text{ratio}} = 8\kappa^*$ at $\bar{\rho}_m = 0$ gives the three stated values; E is the excess above the Parseval ceiling 1. The two-prime ceiling is 8/9 (Theorem 8.1), so $E_{\min} = 1 - 8/9 = 1/9$; the single-prime Fermat tower has $\mu_{\text{ratio}} \to 1$, so $E_{\min} = 0$ and $\kappa \to 1/8$ is its exact boundary. □

Theorem 8.4 (The $R > \frac{1}{2}$ threshold has a Parseval origin)**.** *For any CM-type tower, $R \le \frac{1}{2} + O(1/\hat{f})$. Consequently the threshold $R > \frac{1}{2}$ of Theorem 7.4 is exactly the CM–Parseval boundary, and the upper half of Problem 10.12 cannot be reached inside the whole Kummer family.*

Proof. Theorem 2.4 gives $R = 12\kappa - 1 - \varepsilon_{\text{ar}} + O(1/\hat{f})$. Insert $\kappa \le \frac{1}{8} + \bar{\rho}_m/(2\hat{f})$ from Theorem 8.1 and $\varepsilon_{\text{ar}} \le 0$ from Theorem 8.2 to obtain $R \le \frac{1}{2} + O(1/\hat{f})$. This matches $R \to 0$ (Theorem 7.36) and $R_n \to \frac{1}{2}$ for the Fermat tower, and turns the previously observed threshold into a structural consequence. □

Theorem 8.5 (Finite-level barrier slack)**.** *With $C_S = 2\sum_{q\in S} \log q$ and C the constant of Theorem 8.7,*

$$\kappa \le \frac{1}{3} + \frac{8C_S + C}{12\hat{f}} + o(\hat{f}^{-1}).$$

Hence a finite-level value slightly above 1/3 is not a counterexample.

Proof. Wilms' identity gives $12\kappa = \sigma(1 + R) + \varepsilon_{\text{ar}}$. Theorem 6.1 supplies $\sigma \le 1 + 2C_S/\hat{f}$, Theorem 8.2 supplies $\varepsilon_{\text{ar}} \le C\sigma/\hat{f}$, and branch (c) supplies $R \le 3$. Multiplying, $12\kappa \le 4 + (8C_S + C)/\hat{f} + o(\hat{f}^{-1})$. Numerical illustration with $C = 0$ (the explicit part): for $S = \{3, 5\}$ the slack is 0.666667 at $(n, m) = (1, 1)$ (trivial bound $\kappa \le 1$), 0.143126 at $(5, 5)$, 0.057232 at $(12, 12)$ and 0.015254 at $(50, 40)$, giving $\kappa \le 0.348587$ at the last level; the four pairs $\{3, 5\}, \{3, 7\}, \{5, 7\}, \{3, 11\}$ agree to within 5×10^{-4}. □

Remark 8.6 (Ledger: three branches to two). Branch (b) is withdrawn from the open list: the source text itself marks Theorem 8.7 as proved unconditionally (results table: $\delta_{\text{Fal}} \le C\Delta_{\text{Del}} + O(g)$, archimedean channel cancelled, marked "proved (unconditional)"), so the earlier ledger entry "general-tower archimedean bound remains C" is a misreading and is hereby corrected. Branch (c) and the upper half of Problem 10.12 are identified by Theorem 8.3. The two residual open items are therefore:

(i) an upper bound $R_\rho \le 3$ (equivalently a bound on the excess conductor mass $\mu - g\hat{f}/4 > 0$) for non-automorphic simple isogeny factors — identical to the curve realisation of $\kappa \in (\frac{1}{8}, \frac{1}{3})$;

(ii) the general (non-semistable) conductor–discriminant identity (L1): the source states in Problem 10.11 that "Liu 1994 Prop. 1 is commonly cited for arbitrary genus, but no explicit independent proof has been found", and in §10.1 that the equivalence is published only for the Mumford-graph case. This is the sole condition still attached to $\sigma \to 1$ outside the semistable class.

All fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5 - 18c)$, 47/3, $C \ge 36/7 > 3$, Routes 1–3, $Th.A^*$ and *abc* open are unchanged. The U statements above are relative to the cited theorems of 00544v2, 00712v2 and Appendices BB–YY; no cited proof is verified independently.

8.2 The non-semistable (L1): exact scope and the thickness criterion

Remark 8.6 listed "the general (non-semistable) conductor–discriminant identity (L1)" as an open conditional input. Reading the source text verbatim shows this is a misdescription. We correct it, and in doing so uncover a notational conflict between Appendix 5.2 and Appendix 6.5 that changes one claimed equality into an exponentially small limit.

Notation. Throughout $\mathcal{L}_S := \sum_{p\in S} \log p$; this differs from the constant $C_S = 2\mathcal{L}_S$ used in Theorem 8.5. We write m_p for the number of irreducible components of the geometric special fibre of the minimal regular model, Γ_p for its dual graph, $\mathrm{rk}_{\mathrm{Lap}}\, L(\Gamma)$ for the rank of its Laplacian matrix $D - A$ (equal to $|V| - 1$ on any connected graph), and $b_1(\Gamma) = |E| - |V| + 1$ for the first Betti number.

Theorem 8.7 (Exact scope of (L1))**.** *The identity*

$$\boxed{\Delta_{\mathrm{Del,p}} = f_p + m_p - 1, \qquad \sigma = 1 + \frac{\sum_{p\in S}(m_p - 1)\log p}{\sum_{p\in S} f_p \log p} \geq 1}$$

is unconditional for every genus, every residue characteristic, and without any semistability or hyperellipticity hypothesis*. It is the conjunction of Saito's Theorem 1 ($\delta_p = -\mathrm{Art}_p$, Bloch's conjecture in relative dimension one) and Liu's Proposition 1 ($-\mathrm{Art} = m - 1 + \phi(X/K)$), plus $m_p \geq 1$.*

Proof. The source states verbatim, in Remark 2.5, "Scope: any genus, any residue characteristic, no semistability, no hyperellipticity required" and calls (L1) "an identity plus a trivial inequality, containing no conjecture". Section 10.3 moreover gives the non-semistable example $X_1(p)$, $p > 7$: by Conrad–Edixhoven–Stein the minimal regular model has geometrically integral special fibre, hence $m_p = 1$, $\Delta_{\mathrm{Del,p}} = f_p > 0$ and $\sigma = 1$ exactly, while a semistable thickness reading would incorrectly give $\nu_p = 0$. Thus (L1) is precisely the identity that is valid *outside* semistability; the semistable formulae $f_p = b_1 + \mathrm{Sw}_p$ and $\nu_p = \sum_i N_i$ are the ones that are not interchangeable with it. Problem 10.11 asks only for an "independent first-principles proof" of Liu's Proposition 1 for arbitrary genus, i.e. a completeness-of- citation item; it is not a mathematical hypothesis and its resolution changes no U/C status in this volume. □

Theorem 8.8 (Unified difference identity; correction to SS)**.** *For* any *regular model, without semistability,*

$$\boxed{\Delta_{\mathrm{Del,p}} - f_p = m_p - 1 = |V(\Gamma_p)| - 1 = \mathrm{rk}_{\mathrm{Lap}}\, L(\Gamma_p)}.$$

The decomposition $\Delta_{\mathrm{Del,p}} = |E(\Gamma_p)| + \mathrm{Sw}_p$ and $f_p = b_1(\Gamma_p) + \mathrm{Sw}_p$ of Theorem 5.10 requires semistability; the difference *identity does not. Consequently Theorem 6.34's "$\sigma = 1$ exactly" is withdrawn: the tree graphs of Theorems 6.32–6.33 have $b_1 = 0$ but $\mathrm{rk}_{\mathrm{Lap}}\, L = |V| - 1 = O(p^n)$, not* 0*.*

Proof. The first equality is Theorem 8.7. The second is pure graph theory: the vertices of Γ_p are the irreducible components of the geometric special fibre, so $|V| = m_p$; Zariski connectedness makes Γ_p connected, hence $\ker(D - A)$ is one-dimensional and $\mathrm{rk}_{\mathrm{Lap}}\, L = |V| - 1 = m_p - 1$ for a graph with or without cycles. Appendix 5.2 uses exactly this convention (P_6: $|V| = 6$, $\mathrm{rk}_{\mathrm{Lap}}\, L = 5$, $b_1 = 0$). Appendix 6.5 instead writes $\mathrm{rank}\, L = |E| - |V| + 1 = 0$, which is b_1, not the Laplacian rank; it then feeds 0 into Theorem 5.10. With the correct convention the mixed Kummer tower gives, for $a = p^n, b = q^m$,

$$\sigma - 1 = \frac{(|V(\Gamma_p)| - 1)\log p + (|V(\Gamma_q)| - 1)\log q}{g\, \hat{f}_{\mathrm{tot}}}, \qquad |V(\Gamma_p)| = \frac{p^n + p - 2}{p - 1},$$

by Theorems 6.32 and 6.24. Exact rational checks: $(p, n, q, m) = (3, 12, 5, 12)$ gives $|V(\Gamma_3)| = 265721$, $g = 64873046609280$, $\hat{f}_{\mathrm{tot}} = 63.0899$ and $\sigma - 1 = 2.4072 \times 10^{-8}$; $(3, 50, 5, 40)$ gives

4.7359×10^{-27}. So $\sigma \to 1$ survives, exponentially, but not by equality. The coarser $O(1/n)$ bound of Theorem 6.1 remains valid. □

Theorem 8.9 (Thickness decomposition and the exact criterion)**.** *Let $N_{p,i}$ be the thicknesses of the nodes of the stable model at p, so that resolving $xy = p^{N_{p,i}}$ inserts $N_{p,i}-1$ rational components. Then*

$$m_p - 1 = \underbrace{(m_p^{\mathrm{st}} - 1)}_{\le 2g-3} + \sum_i (N_{p,i} - 1), \qquad \sigma - 1 = \frac{2\vartheta_n \mathcal{L}_S}{\hat{f}} + \tau_n,$$
$$\tau_n := \frac{\sum_{p,i}(N_{p,i} - 1)\log p}{g\hat{f}}, \quad \sigma \to 1 \iff \tau_n \to 0$$

with $0 \le \vartheta_n \le 1$. The stable part tends to 0 by (T3) alone; the thickness density τ_n is the only possible obstruction.

Proof. The thickness formula and $m = m^{\mathrm{st}} + \sum_i (N_i - 1)$ are stated verbatim in Section 10.3; Theorem 5.17 gives $m^{\mathrm{st}} - 1 \le 2g - 3$ for the stable fibre. Insert into Theorem 8.7 and divide by $g\hat{f} = \sum_p f_p \log p$. The bounded factor ϑ_n records that $m_p^{\mathrm{st}} - 1$ may be smaller than $2g - 3$. Section 10.3 states "channel (D) is the higher-genus Szpiro conjecture in thickness language" and Section 10.1 tabulates "$\sigma = O(1)$ (number field): missing = Szpiro/abc". Hence the general-tower case of $\tau_n \to 0$ is Szpiro/*abc*-grade; it is not a new independent obstruction but the tower-side avatar of the pointwise-side classification of Appendix 3.1. □

Theorem 8.10 (Closure without semistability under bounded thickness)**.** *If the thicknesses are uniformly bounded, $N_{p,i} \le N_0$ along the tower, then $\sigma \to 1$ unconditionally relative to (L1) and (T3); no semistability, no Mestre bound and no BMY input are used.*

Proof. The stable fibre has at most $2g-2$ nodes (stability: each component contributes at least 1 to $\sum_v (2g_v - 2 + n_v) = 2g-2$). Hence $\sum_i (N_{p,i} - 1) \le (2g-2)(N_0 - 1)$ and $\tau_n \le 2N_0 \mathcal{L}_S / \hat{f} \to 0$. The stable part likewise tends to 0. This extends Theorem 6.1 from semistable towers to all towers of bounded thickness; precise checks for $N_0 = 1, 2, 3, 5, 10$ give $\mathrm{rk}_{\mathrm{Lap}}\, L \le 2g - 3 + (2g-2)(N_0 - 1) = O(N_0 g)$. □

Theorem 8.11 (All explicit towers of this volume have $\tau_n \to 0$)**.** *(a) $X_1(p)$, $p > 7$ (non-semistable): $m_p = 1$, $\sigma = 1$ exactly.*

(b) $X_0(p)$: $m_p = 2$, $f_p = g$, $\sigma = 1 + 1/g$.

(c) Fermat and mixed Kummer towers: $|V| = O(\sqrt{g})$ and $\hat{f}_{\mathrm{tot}} \asymp n \log p + m \log q$, so $\sigma - 1 = O(p^{-n}/n + q^{-m}/m) \to 0$.

(d) The $\mathbb{P}^1$-cactus of Section 10.3 has arbitrary thickness, $\sigma = \nu/g$ unbounded, but $\hat{f} = O(1)$; it violates (T3) and is outside the tower world.

Proof. (a),(b) are the two anchors of Section 2.4 and Section 5.2 ($g = 2, 3, 11, 101, 10^4$ give $\sigma = 1.5, 1.333, 1.091, 1.010, 1.0001$). (c) is the exact rational computation of Theorem 8.8; nine towers $(3,1,5,1)$ through $(3,50,5,40)$ give $\sigma - 1$ from 1.25×10^{-1} down to 4.74×10^{-27}. (d) is the source's own counterexample: g self-nodes of thicknesses N_i give $f = b_1 = g$, $\nu = \sum N_i$, $\sigma = \nu/g$ arbitrary, while $\hat{f} = \log p = O(1)$. Thus thickness *alone* is insufficient, as Section 10.3 states. □

Theorem 8.12 (Discriminant conventions and Problem 10.11)**.** *Three distinct objects must not be interchanged: $\Delta_{\mathrm{Del}} = -\mathrm{Art} = f + m - 1$ (used throughout this volume), $\Delta_{\min}$ (minimal Weierstrass discriminant, used by Szpiro), and the stable-model singularity count $e_2 \le 2g - 2$. One has $\nu(\Delta_{\min}) \ge \Delta_{\mathrm{Del}}$, equality for $g = 1$ (Ogg–Saito), an exact formula for $g = 2$ (Liu, Théorème 1), and an inequality in residue characteristic 2 (Obus–Srinivasan); the exact difference for arbitrary genus is open. This affects only the Szpiro/abc translation, never σ, κ, R. Problem 10.11 (an*

independent first-principles proof of Liu's Proposition 1 for arbitrary genus) reduces, by Saito's Theorem 1, to the ℓ-adic Euler characteristic decomposition in which the H^0, H^2 terms contribute exactly $-(m-1)$; in the semistable case this is the graph identity of Theorem 5.10. We record the reduction but do not supply the non-semistable normalisation argument (C).

Proof. Verbatim from Remark 2.5 (terminological warning) and Section 10.1. The reduction statement follows because H^0, H^2 of a curve are $\mathbb{Q}_\ell$ with trivial inertia action, so the entire surface-level Artin conductor is carried by H^1 plus the special-fibre component defect $m-1$; the semistable case is realised explicitly by $\Delta - f = |V| - 1$ (Theorem 8.8). No consequence for the standing conclusions depends on the missing normalisation. □

Remark 8.13 (Ledger). Withdrawn: Remark 8.6(ii) as an "open (L1) for non-semistable towers" (Theorem 8.7); the equality "$\sigma = 1$ exactly" of Theorem 6.34 (Theorem 8.8, replaced by an exponentially small limit); and the semistability restriction on the *difference* identity of Theorem 5.10. Promoted: $\sigma \to 1$ to all towers of bounded thickness (Theorem 8.10). Reclassified: Problem 10.11 from a conditional input to a completeness-of-citation note; the $\Delta_{\min}$ gap to a Szpiro-translation interface; general-tower $\sigma \to 1$ to the thickness form of higher-genus Szpiro, i.e. *abc*-grade and already tabulated in Section 10.1. Appendix 9.1 shows that this label conflates three distinct classes: the upper-half anchors are already realised by automorphic non-CM GL_2 modular factors, while the missing input $R \le 3$ belongs to general GL_2 and non-GL_2 towers. The remaining items are therefore two gaps (Problem 10.12 single-curve density, and general arithmetic BMY) together with the convexity warning of Section 7.6.1 of the source text. All fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \ge 36/7 > 3$, Routes 1–3, $Th.A^*$ and *abc* open are unchanged. All U statements are relative to the cited theorems of 00544v2, 00712v2 and Appendices BB–BBB; no cited proof is verified independently here.

---

9 Isogeny factors: the two-gap verdict

9.1 Non-automorphic simple isogeny factors: category correction and the two-gap verdict

Remark 8.6(i) and the closing ledger of Appendix 8.2 name "the non-automorphic simple isogeny factor with $R_\rho \le 3$" as the single remaining mathematical open item, and identify it with the upper half of Problem 10.12. Reading the source text verbatim shows that three distinct classes of isogeny factors were conflated. We separate them, derive an exact linear dictionary between the conductor mass ratio and the curvature ratio, quantify the convexity constraint, and restate the open items.

Theorem 9.1 (Category correction)**.** *The inference of Theorem 8.3, namely that $\kappa > \frac{1}{8}$ forces an isogeny factor which is not of CM type, is valid. Its identification with a* non-automorphic *simple factor is not. The three anchor values $\kappa = \frac{5}{24}, \frac{1}{4}, \frac{1}{3}$ are realised by the modular towers $X(p^n)$, $X(N)$, $X_0(N)/X_1(N)$, which are GL_2-type (automorphic) and non-CM, and whose R-values are theorems of Dolce–Mercuri, Mayer and Grados–von Pippich.*

Proof. The anchor table of §2.4 of the source text lists $R \le \frac{3}{2}$ for $X(p^n)$, $R \to 2$ for $X(N)$ and $R \to 3$ for $X_0(N), X_1(N)$, with $\sigma \to 1$ throughout; Proposition 2.3 then gives $\kappa = (1+R)/12$, i.e. $\frac{5}{24}$, $\frac{1}{4}$, $\frac{1}{3}$. §4 of that text is titled "The one-third barrier for GL_2-type towers" and Theorem 4.5 states the barrier for exactly this class; Theorem 4.7 records $B = 1$ for GL_2-type towers and $B = 2$ for CM towers. The table of Problem 10.1 lists these modular families as "theorem". Hence the factors providing $\kappa > \frac{1}{8}$ are automorphic and non-CM. The Parseval identity $\sum_\chi m_\chi = g/2$ of

Appendix 7.4 is a statement about CM types (character m-vectors of three-point orbits); a GL_2 newform factor has no such vector, so "not of CM type" is the correct conclusion of Theorem 8.3. Since "non-CM" does not imply "non-automorphic", the label in Remark 8.6(i) is a category error: it is a valid inference followed by an invalid strengthening. □

Theorem 9.2 (The exact μ_{ratio}–R line)**.** *Assume $\sigma \to 1$ and archimedean absorption as in Appendix 8.1. Then*

$$R = \tfrac{3}{2}\mu_{\mathrm{ratio}} - 1, \qquad \mu_{\mathrm{ratio}} = \tfrac{2}{3}(1+R), \qquad \kappa = \tfrac{1}{8}\mu_{\mathrm{ratio}}.$$

Consequently arithmetic BMY, $R \le 3$, Parseval's bound $\mu_{\mathrm{ratio}} \le 1$, and the barrier $\kappa \le \frac{1}{3}$ are three readings of the same line:

$$R \le 3 \iff \mu_{\mathrm{ratio}} \le \tfrac{8}{3} \iff \kappa \le \tfrac{1}{3}, \qquad \mu_{\mathrm{ratio}} \le 1 \iff R \le \tfrac{1}{2} \iff \kappa \le \tfrac{1}{8}.$$

Proof. Theorem 8.1 gives $\mu_{\mathrm{ratio}} = 8\kappa + O(\hat{f}^{-1})$ in the limit, and the dictionary of Appendix 2.1 gives $R = 12\kappa - 1 - \varepsilon_{\mathrm{ar}} + O(\hat{f}^{-1})$; with $\varepsilon_{\mathrm{ar}} \to 0$ (Theorem 8.2) we obtain $R = 12\kappa - 1 = \frac{3}{2}(8\kappa) - 1 = \frac{3}{2}\mu_{\mathrm{ratio}} - 1$. Solving for μ_{ratio} and for κ gives the other two forms. The three equivalences follow by substituting $R = 3$, $\mu_{\mathrm{ratio}} = 1$ and $\kappa = \frac{1}{3}$. All relations are exact rational identities in the limit; the $O(\hat{f}^{-1})$ terms are the finite-layer remainders controlled in Appendices 7.5 and 8.1. □

Corollary 9.3 (Seven-anchor calibration)**.** *For the seven known tower values $\kappa = \frac{1}{12}, \frac{1}{9}, \frac{1}{8}, \frac{1}{6}, \frac{5}{24}, \frac{1}{4}, \frac{1}{3}$, the triples $(\kappa, R, \mu_{\mathrm{ratio}})$ are respectively $(\frac{1}{12}, 0, \frac{2}{3})$, $(\frac{1}{9}, \frac{1}{3}, \frac{8}{9})$, $(\frac{1}{8}, \frac{1}{2}, 1)$, $(\frac{1}{6}, 1, \frac{4}{3})$, $(\frac{5}{24}, \frac{3}{2}, \frac{5}{3})$, $(\frac{1}{4}, 2, 2)$, $(\frac{1}{3}, 3, \frac{8}{3})$; in each case $R = \frac{3}{2}\mu_{\mathrm{ratio}} - 1$ holds exactly, and the Parseval excess $E = \mu_{\mathrm{ratio}} - 1$ equals $-\frac{1}{3}, -\frac{1}{9}, 0, \frac{1}{3}, \frac{2}{3}, 1, \frac{5}{3}$.*

Proof. Substitute each rational κ into $R = 12\kappa - 1$ and $\mu_{\mathrm{ratio}} = 8\kappa$ and reduce in $\mathbb{Q}$; the identity $R - \frac{3}{2}\mu_{\mathrm{ratio}} + 1 = 12\kappa - 1 - 12\kappa + 1 = 0$ is a polynomial identity, hence holds for every value in the list. The row $\kappa = \frac{1}{4}$ is the reading of $R \to 2$ for $X(N)$ from §5.4 of the source text through the dictionary of Appendix 2.1; the row $\kappa = \frac{1}{9}$ is the single-sided mixed Kummer ceiling of Appendix 7.5. □

Theorem 9.4 (The three-fold classification of the source text)**.** *The source text distinguishes exactly three classes with respect to [Arith-BMY]: (a) the modular towers $X_0(N)$, $X_0(p^n)$, $X_1(N)$, $X(N)$, for which BMY is a theorem; (b) general GL_2-type towers, listed as open in the table of Problem 10.1; (c) arbitrary non-GL_2 high-dimensional towers, of which §5.4 says verbatim that they "remain covered only by the conjectural BMY". Consequently the upper half of $(\frac{1}{8}, \frac{1}{3})$ already has theorem-level realisations, whereas the input still missing is $R \le 3$ on classes (b) and (c).*

Proof. This is a verbatim reading of the Problem 10.1 table, Theorem 4.5 with its Remark 4.6, and the closing sentence of §5.4. The §5.4 table is headed "R for known non-automorphic towers" yet lists three modular families that are automorphic; we therefore classify by the explicit per-row attributions and by the Problem 10.1 table, not by the heading. Remark 8.6(i) attributed $\kappa > \frac{1}{8}$ to a new analytic input on non-automorphic factors; Theorem 9.1 shows the realised upper-half anchors require no such input, so the missing input belongs to (b) and (c). □

Proposition 9.5 (Convexity: scope and quantitative consequence)**.** *§7.6.1 of the source text states that convexity may not be used in $\kappa \to \infty$ arguments. The isogeny-factor convex hull of Appendix 3.3 applies to the isogeny decomposition of a single curve Jacobian (h_F and the conductor are isogeny invariants), hence it is legitimate for bounded averages but not for Problem 10.2. If all factors other than one satisfy the CM–Parseval bound $\kappa_{\mathrm{bg}} \le \frac{1}{8}$ and the distinguished factor has weight w, then*

$$\kappa_{\max} \;\ge\; \kappa^* + \frac{(1-w)(\kappa^* - \frac{1}{8})}{w}, \qquad R_{\max} = 12\kappa_{\max} - 1 = 8\mu_{\mathrm{ratio,max}} - 1.$$

In particular $w = 1$ and $\kappa^* = \frac{1}{3}$ gives $R_{\max} = 3$ (BMY equality), while $w = \frac{1}{2}$ and $\kappa^* = \frac{1}{3}$ already requires $R_{\max} = \frac{11}{2} > 3$.

Proof. By the convex hull, $\kappa^* \le (1-w)\kappa_{\mathrm{bg}} + w\kappa_{\max}$; solve for $\kappa_{\max}$ using $\kappa_{\mathrm{bg}} = \frac{1}{8}$. The formula for $R_{\max}$ is Theorem 9.2. With $w = \frac{1}{2}$, $\kappa^* = \frac{1}{3}$ one gets $\kappa_{\max} = \frac{1}{3} + \frac{1/2\,(1/3-1/8)}{1/2} = \frac{13}{24}$, hence $R_{\max} = \frac{11}{2}$. The full rational table for $\kappa^* \in \{\frac{1}{6}, \frac{5}{24}, \frac{1}{4}, \frac{1}{3}\}$ and $w \in \{1, \frac{1}{2}, \frac{1}{4}, \frac{1}{10}\}$ shows that every interior realisation with $w < 1$ and $\kappa^* \ge \frac{1}{4}$ forces $R_{\max} > 3$, and that for $\kappa^* = \frac{1}{3}$ any $w < 1$ forces $R_{\max} > 3$. The hypothesis $\kappa_{\mathrm{bg}} \le \frac{1}{8}$ is essential: if the remaining factors contain a GL_2 modular factor, κ_{bg} can reach $\frac{1}{3}$ and the bound is void. □

Theorem 9.6 (Two levels of the curvature input)**.** *Closing Problem 10.2 (no $\kappa \to \infty$) requires only $R = O(1)$, whereas the barrier $\kappa \le \frac{1}{3}$ requires $R \le 3$. The weaker statement is strictly weaker than arithmetic BMY, is true in the function-field case by Deligne semipositivity, and its number-field reduction is Theorem 9.1 and Corollary 9.2 of the source text ($\bar{\omega}^2 \le 12h_F + 2\log(2\pi^4)g$; Faltings–Shafarevich $h_F = O_S(g) \Rightarrow \bar{\omega}^2 = O(g)$), whose required tools the text itself records as missing (Remark 9.2).*

Proof. By Proposition 2.3, $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$; with $\sigma = O(1)$ and $\varepsilon_{\mathrm{ar}}$ absorbed (Theorem 8.7), $\kappa \to \infty$ is equivalent to $R \to \infty$. Hence $R = O(1)$ suffices for boundedness, while $\kappa \le \frac{1}{3}$ needs the sharp constant $R \le 3$ at $\sigma \to 1$. §9.5 of the source text states verbatim that arithmetic Bogomolov $R = O(1)$ is strictly weaker than BMY, that refuting $\kappa \to \infty$ does not require full BMY, and that Deligne semipositivity proves the function-field analogue. Remark 9.2 records that the tools for the reverse direction are still missing. Therefore Remark 8.6(i)'s single item "$R_\rho \le 3$" must be split into a weak level ($R = O(1)$) and a strong level ($R \le 3$). □

Corollary 9.7 (The two-gap verdict)**.** *The tower side has two remaining mathematical open items and one methodological warning.*
(I) Construction*: the dense realisation of interior values of $(\frac{1}{12}, \frac{1}{3})$ by single-curve towers (Problem 10.12). The anchors $\frac{5}{24}, \frac{1}{4}, \frac{1}{3}$ are already realised by modular towers and $(0, \frac{1}{8}]$ is dense by Appendix 6.2; the convexity warning of §7.6.1 forbids lifting the product-tower filling directly.*
(II) Input*: $R \le 3$ on general GL_2-type and non-GL_2 towers (Problem 10.1), equivalent to general arithmetic BMY and to a component of generalized Szpiro; the weak level $R = O(1)$ closes $\kappa \to \infty$.*
(III) Warning*: convexity is an abelian-variety statement and may not be used in $\kappa \to \infty$ arguments.*
Together with the bounded-thickness translation of $\sigma \to 1$ to the thickness form of higher-genus Szpiro (Appendix 8.2), the two abc*-grade inputs lie one in the discriminant channel and one in the curvature channel.*

Proof. Theorem 9.1 removes the identification of (I) with a non-automorphic factor, Theorem 9.4 supplies the three classes, Proposition 9.5 quantifies the convexity constraint, and Theorem 9.6 separates the weak and strong curvature inputs. The remaining statements are the table of Problem 10.12 and the bounded-thickness result of Appendix 8.2. □

Remark 9.8 (External literature ledger and honest boundary). Ribet (Thm 4.4) together with Khare–Wintenberger gives: every GL_2-type abelian variety over $\mathbb{Q}$ is modular. Tsimerman (*Ann. of Math.* **176** (2012), 637–650) proves that for $g > 3$ there exists an abelian variety over $\bar{\mathbb{Q}}$ isogenous to no Jacobian, so non-automorphic simple varieties are not in short supply; the obstruction is their Torelli realisation. The Coleman–Oort conjecture as recorded by Moonen–Oort predicts no positive-dimensional special subvariety generically inside the Torelli locus for large g, while the counterexamples for $g \in \{4, 5, 6, 7\}$ all arise from abelian covers of $\mathbb{P}^1$ with fixed group and local monodromy — precisely the Kummer family of Appendices 6.3–7.5, which satisfies $\kappa \le \frac{1}{8}$. Lu–Zuo prove the hyperelliptic case for $g > 7$. Lu (arXiv:1211.0451) records generalized Szpiro as a conjecture and reduces it to the semistable case via uniform base-change conductors. Finally $\dim \mathcal{M}_g = 3g - 3$ and $\dim \mathcal{A}_g = g(g+1)/2$ give codimension $0, 0, 1, 3, 6, 10, 15, 21, 28, 36, 45$ for $g = 2, \ldots, 12$; codimension becomes positive exactly at $g = 4$, the Tsimerman threshold $g > 3$.

All six external items are *text-level checks*: no cited proof is verified independently. The structural reading "towers satisfying (T3) and non-automorphy are mutually exclusive" is conditional on Coleman–Oort and is labelled C, not U. Theorem 9.1–Corollary 9.7 are U relative to the cited sections of 00544v2, 00712v2 and Appendices 2.1–8.2. All U statements are relative to standard citations, not absolutely unconditional. This appendix does not constitute an independent verification of the proof of Theorem A. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, Routes 1–3, $Th.A^*$ and *abc* open are unchanged.

10 The construction layer: the weight ruler and the cover

10.1 The construction layer of gap (I): the weight ruler and the curve-level reduction

Appendix 9.1 closes with the verdict that the upper half of Problem 10.12 (dense single-curve realisation of interior κ values) and general arithmetic BMY are the two remaining open items, and records the warning that convexity may not be used in $\kappa \to \infty$ arguments. The present appendix explores the *construction layer* of gap (I): not whether a desired κ is consistent, but which explicit object has to be built. The result is again a correction: on the isogeny (AV) side the interior is already closed by an exact linear ruler, so 9.7(I) must be narrowed; on the curve side the only missing object is a covering realisation of a single $R = 3$ factor.

Theorem 10.1 (Weight ruler)**.** *Let a semistable tower have an isogeny decomposition whose factors split into a CM/Kummer class with $\kappa_{\mathrm{CM}} \leq 1/8$ (the CM–Parseval bound of Appendix 8.1) and a GL_2-type (modular) class with $\kappa_{\mathrm{G}} = 1/3 + O(\hat{f}_{\mathrm{G}}^{-1})$. Put*

$$w = \frac{g_{\mathrm{G}}\hat{f}_{\mathrm{G}}}{g_{\mathrm{G}}\hat{f}_{\mathrm{G}} + g_{\mathrm{CM}}\hat{f}_{\mathrm{CM}}} \in [0,1].$$

Then, up to $O(\hat{f}^{-1})$ corrections,

$$\kappa = \frac{1}{8} + \frac{5}{24}w, \qquad w = \frac{24\kappa - 3}{5}, \qquad R = \frac{3}{2}(8\kappa) - 1 = \frac{w}{2} + \frac{1}{2}.$$

The anchor values are exact rationals on this ruler: $\kappa = 1/8, 1/6, 5/24, 1/4, 1/3$ correspond to $w = 0, 1/5, 2/5, 3/5, 1$; the Heisenberg endpoint $\kappa = 1/12$ gives $w = -1/5$ and therefore requires a background factor with $\kappa_{\mathrm{bg}} < 1/8$ (equivalently $R = 0$), not a CM background.

Proof. Stable Faltings height and conductor are isogeny invariants and are additive under products (Appendix 3.3); hence κ is the conductor-mass weighted mean of the factor values, with weights $w_i = g_i\hat{f}_i / \sum_j g_j\hat{f}_j$. With two classes this gives $\kappa = w\kappa_{\mathrm{G}} + (1-w)\kappa_{\mathrm{CM}}$. Substituting the two endpoint values gives $\kappa = 1/8 + w(1/3 - 1/8) = 1/8 + 5w/24$. Inverting is elementary arithmetic: $w = (24\kappa - 3)/5$. The R formula is the exact line $R = \frac{3}{2}\mu_{\mathrm{ratio}} - 1$ with $\mu_{\mathrm{ratio}} = 8\kappa$ (Appendix 9.1), i.e. $R = 12\kappa - 1 = w/2 + 1/2$. The six rational evaluations are verified with exact fractions; for $\kappa = 1/12$ one gets $w = -1/5 < 0$, which is impossible for a probability weight, so that endpoint cannot be generated from a CM background. □

Theorem 10.2 (Convexity identity; correction to 9.5)**.** *Let $\kappa^*(w) = 1/8 + 5w/24$ be an interior value on the ruler and let $\kappa_{\mathrm{bg}} \leq 1/8$. The convex-hull requirement of Appendix 9.1,*

$$\kappa_{\max} \;\geq\; \kappa^* + \frac{(1-w)(\kappa^* - \kappa_{\mathrm{bg}})}{w},$$

is, for $\kappa_{\mathrm{bg}} = 1/8$, the exact identity

$$\kappa_{\max} \equiv \frac{1}{3}, \qquad R_{\max} \equiv 3 \quad \textit{for every } 0 < w \leq 1.$$

Consequently no factor with $R > 3$ is needed to realise any interior value on the ruler: one $R = 3$ (BMY-equality, modular) factor suffices. The value $R_{\max} = 11/2$ quoted in 9.5 arises only in the different problem where $\kappa^* \equiv 1/3$ is held fixed while $w = 1/2$: then $\kappa_{\max} = 1/3 + (1/2)(1/3 - 1/8)/(1/2) = 13/24$ and $R_{\max} = 11/2$.

Proof. On the ruler $\kappa^* - 1/8 = 5w/24$, so $(1-w)(\kappa^* - 1/8)/w = 5(1-w)/24$, and $\kappa_{\max} = 1/8 + 5w/24 + 5(1-w)/24 = 1/8 + 5/24 = 1/3$, independent of w; then $R_{\max} = 12 \cdot (1/3) - 1 = 3$. With $\kappa^* = 1/3$ and $w = 1/2$ the same formula gives $1/3 + 5/24 = 13/24$ and $R_{\max} = 13/2 - 1 = 11/2$. All computations are exact fractions ($w = 1/5, 2/5, 3/5, 1/2, 1$ checked). Hence the reading "interior realisation requires a factor with $R > 3$" in 9.7 is withdrawn: it confuses a fixed-target problem with the ruler-varying problem. □

Theorem 10.3 (AV-side closure of gap (I))**.** *In the category of abelian varieties up to isogeny, the interior $(1/8, 1/3)$ is densely realised along the ruler of Theorem 10.1. If the two conductor masses are governed by two multiplicatively independent primes $p \neq q$, i.e. $\log(g_{\mathrm{G}}\hat{f}_{\mathrm{G}}) - \log(g_{\mathrm{CM}}\hat{f}_{\mathrm{CM}}) = n \log p - m \log q + O(1)$, then the weights $w_{n,m} = 1/(1 + \exp(-(n \log p - m \log q + O(1))))$ are dense in $(0, 1)$ by Kronecker's theorem ($\log p / \log q \notin \mathbb{Q}$). Consequently 9.7(I), which states that a dense realisation of arbitrary interior values is missing, must be narrowed: the density is already present on the AV side; what is missing is its* curve *realisation.*

Proof. In characteristic zero the group algebra of the (finite) Galois group of a cover is semisimple (Maschke), so the Chevalley–Weil decomposition of H^1 splits $J(X)$ into a direct sum of isogeny factors, with no extension ambiguity. Products of abelian varieties therefore realise the weighted mean exactly. For density, the set $\{n \log p - m \log q : n, m \in \mathbb{N}\}$ is dense in $\mathbb{R}$ modulo the $O(1)$ normalisation constants because $\log p / \log q$ is irrational; the logistic map is continuous and monotone from $\mathbb{R}$ onto $(0, 1)$. A finite enumeration with $p = 3, q = 5$ and $1 \leq n, m \leq M$ gives maximal gaps 0.088763, 0.031493, 0.019237, 0.006975 for $M = 10, 20, 40, 80$; the decay is data-level, the density itself is a theorem. □

Proposition 10.4 (Finite-layer strictness)**.** *Let the GL_2 factor satisfy $\kappa_{\mathrm{G}}(n) = 1/3 - C_{\mathrm{G}}/\hat{f}_{\mathrm{G}} + O(\hat{f}_{\mathrm{G}}^{-2})$ with $C_{\mathrm{G}} > 0$ (convergence to the BMY endpoint from below). Then for $w > 0$*

$$\kappa \;\leq\; \frac{1}{8} + \frac{5w}{24} - \frac{wC_{\mathrm{G}}}{\hat{f}_{\mathrm{G}}} + O(\hat{f}^{-2}) \;<\; \frac{1}{8} + \frac{5w}{24}.$$

Hence every anchor of Theorem 10.1 is an asymptotic point, not a finite-layer value; $\kappa \to 1/3$ requires $w \to 1$ and $\hat{f}_{\mathrm{G}} \to \infty$ simultaneously.

Proof. Substitute $\kappa_{\mathrm{G}}(n)$ into the weighted mean and use $\kappa_{\mathrm{CM}} \leq 1/8$; the coefficient of w is $1/3 - 1/8 - C_{\mathrm{G}}/\hat{f}_{\mathrm{G}} = 5/24 - C_{\mathrm{G}}/\hat{f}_{\mathrm{G}}$. Since $C_{\mathrm{G}} > 0$ and $w > 0$ the strict inequality follows. For $w = 1$ the correction is maximal, so the BMY endpoint is approached but never attained at finite level; this is consistent with the from-below convergence recorded in Appendices 5.6 and 5.7. □

Theorem 10.5 (Curve-level reduction)**.** *A curve realisation of an interior value $\kappa^* \in (1/8, 1/3)$ along the ruler requires (i) an isogeny factor with $R = 3$ (a BMY-equality GL_2/modular factor), and (ii) the expression of that factor as a Chevalley–Weil irreducible summand of a* curve *cover, i.e. a Prym–Torelli realisation. Condition (ii) is exactly Problem 10.12. Moreover the mixed Kummer family $C_{a,b} : u^{p^n} + v^{q^m} = 1$ cannot supply (i): by the exact orbit correspondence of Appendix 7.4 all its irreducible summands are CM types, hence $\kappa \leq 1/8$ strictly (Appendices 7.5, 8.1), and its weights are locked to the Swan share (Appendices 4.1, 5.3). The minimum weight required by a target is*

$$w_{\min}(\kappa^*) = \frac{24\kappa^* - 3}{5}.$$

Proof. Theorem 10.2 gives (i) with the precise value $R = 3$ and shows no stronger factor is needed. Since every AV product used in Theorem 10.3 must be replaced by a single curve, the decomposition has to come from a Galois cover of a curve; this is the Prym realisation question of Appendix 4.1 and Problem 10.12. For the mixed family, Appendix 7.4 proves $\Phi_{a,b} = \bigsqcup_{c,d} \Phi_{O_{c,d}}$ is a disjoint union of complete three-point-orbit CM types, so every summand is CM; Appendix 7.5 gives $\mu_{\text{ratio}} \le 8/9 - \delta$ and $\kappa \le 1/9$ in the joint limit, and Appendix 8.1 gives the strict $\kappa < 1/8$ for the whole Kummer class. Inverting the ruler gives $w_{\min}$. □

Corollary 10.6 (Restated ledger)**.** *The two gaps of Appendix 9.1 are restated as: (I′)* curve *realisation of one $R = 3$ GL_2 factor as a Chevalley–Weil summand (a Prym–Torelli problem), with the weight constraint $w \ge w_{\min}(\kappa^*)$; the AV-side interior density is closed by Theorem 10.3. (II) general GL_2/non-GL_2 towers with $R \le 3$ (general arithmetic BMY / generalised Szpiro), unchanged. The convexity warning of 00544v2 § 7.6.1 is unchanged and applies only to $\kappa \to \infty$ arguments. The three upper-half anchors $1/6, 5/24, 1/3$ require $w = 1/5, 2/5, 1$ respectively.*

Proof. Immediate from Theorems 10.2, 10.3, 10.5 and the exact rational evaluations of Theorem 10.1. Gap (II) is untouched by this appendix; the convexity warning forbids applying the weighted mean to divergent κ, not to bounded convex combinations. □

Remark 10.7 (Status and boundaries). 10.1–10.6 are U relative to the isogeny additivity of Appendix 3.3, the CM–Parseval bound of Appendix 8.1, the exact μ–R line and the modular endpoints of Appendix 9.1, the orbit correspondence of Appendix 7.4, and the Swan/thickness results of Appendices 4.1, 8.2. They are not absolutely unconditional. The density rate in Theorem 10.3 and the positivity of C_{G} in Proposition 10.4 are data-level / assumptions. The Prym–Torelli realisation (I′) is subsumed under the already recorded C item of Problem 10.12: no new C input and The readings "interior needs $R > 3$" (9.7) and "dense interior realisation is missing" (9.7(I)) are withdrawn as stated and replaced by Corollary 10.6. This appendix does not constitute an independent verification of the proof of Theorem A. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5 - 18c)$, $47/3$, $C \ge 36/7 > 3$, Routes 1–3, $Th.A^*$ and *abc* open are unchanged; no new No-Go is added.

10.2 Doing (I′): the reindexing ruler, the product obstruction, and the two halves of the curve realisation

Appendix 10.1 reduces the upper half of Problem 10.12 to (I′): expressing a single $R = 3$ (equivalently $\kappa = 1/3$) GL_2 factor as a Chevalley–Weil irreducible summand of a curve cover, with weight $w \ge w_{\min}(\kappa^*) = (24\kappa^* - 3)/5$. Reading the source text verbatim shows that the weight ruler of Theorem 10.1 is already the Reindexing formula of 00544v2, that exact attainment is only countable-dense, and that (I′) splits into a half that is already closed and a half that remains open. No new external source is needed.

Theorem 10.8 (Reindexing = weight ruler)**.** *00544v2 Theorem 7.4 and Convention 7.5 give, for the tower $T_k^{(a,b)} = J(F_{p^k})^{\oplus a} \times J_0(p^{2k})^{\oplus b}$ with the* fixed *bad set $\{p\}$,*

$$\kappa(T_k^{(a,b)}) \longrightarrow \frac{a\rho\,\kappa_F + b\,\kappa_X}{a\rho + b}, \qquad \rho = \frac{6p}{p+1}, \quad \kappa_F = \frac{1}{8}, \quad \kappa_X = \frac{1}{3}.$$

With $t = b/(a\rho)$ and $w = t/(1+t) = b/(a\rho + b)$ this is exactly

$$\kappa = \frac{1}{8} + \frac{5}{24}w, \qquad w = \frac{24\kappa - 3}{5},$$

i.e. the ruler of Theorem 10.1. In particular Convention 7.5's condition "$\kappa = 1/4$ at $b/(a\rho) = 3/2$, i.e. $b/a = 9p/(p+1)$" is precisely $w = 3/5$ of the CCC ruler. Consequently 10.1 is an equivalent rewriting of the published Reindexing formula, not a new result; 10.2 (the convexity identity) and the boundary analysis remain new.

Proof. Substitute $\kappa_F = 1/8$, $\kappa_X = 1/3$: $\kappa = (1/8 + t/3)/(1+t) = 1/8 + t(1/3 - 1/8)/(1+t) = 1/8 + (5/24)\cdot t/(1+t)$. Hence $w = t/(1+t)$ and inversion gives $w = (24\kappa - 3)/5$. For $t = 3/2$, $w = (3/2)/(5/2) = 3/5$ and $\kappa = 1/8 + (5/24)(3/5) = 1/4$; $b/a = t\rho = (3/2)(6p/(p+1)) = 9p/(p+1)$. Exact-fraction checks at $t = 0, 1/2, 3/2, 2, 3, 10^6$ give $w = w_{\text{full}}(\kappa)$ identically, and $w = 1/5, 2/5, 3/5, 1$ reproduce $\kappa_{\max} \equiv 1/3$ of Theorem 10.2. The extraction "$b/(a\rho) = 2/3$" in the plain-text version of Convention 7.5 is an OCR reading of $3/2$; the value $3/2$ is forced by both the $9p/(p+1)$ companion and the $\kappa = 1/4$ target. □

Theorem 10.9 (Countable-dense correction to 10.3)**.** *00544v2 Theorem 7.9 states that exact attainment of an interior value holds* precisely *for v in the countable set $\{\frac{a\rho\kappa_F + b\kappa_X}{a\rho + b} : a, b \in \mathbb{Z}_{>0}\}$ and that "the attainable set is not closed"; its closure is $[1/12, 1/3]$. Hence the weights $w = b/(a\rho + b)$ realisable exactly are countable and dense in $[0,1]$, with closure $[0,1]$. The Kronecker/logistic argument of Theorem 10.3 establishes density of the* closure*; it must not be read as exact attainment of every real w. With $\rho = 9/2$ ($p = 3$) and $1 \le a, b \le M$ the maximal gaps are* 0.070588, 0.045113, 0.020022, 0.007203, 0.003301 *for $M = 5, 10, 20, 40, 80$.*

Proof. For fixed $\rho \in \mathbb{Q}_{>0}$ the map $(a,b) \mapsto b/(a\rho + b)$ has countable image; because $\mathbb{Q}_{>0}$ is dense and the map is continuous in $t = b/a$, the image is dense in $(0,1)$. The closure is $[0,1]$ (limits $t \to 0, \infty$). The five gap values are exact finite enumerations. Theorem 7.9 also records that the endpoints $1/12$ (Heisenberg) and $1/3$ (modular) are attained, so the closure statement of Appendix 10.1 is correct while the word "dense" must be qualified by "countable-dense, closure dense". This is parallel to the single-curve density of Appendix 6.2 on $(0, 1/8]$: the two densitites cover the lower and the full interval respectively. □

Theorem 10.10 ((I′) splits: single factor closed, weighted curve open)**.** *(I′) decomposes as follows.*

(I′a) Single-factor realisation. *A GL_2 modular factor A_f with $\kappa = 1/3 + O(\hat{f}^{-1})$ is an isogeny summand of the Jacobian of a* single *curve: the Eichler–Shimura / Atkin–Lehner newform decomposition gives $J(X_0(N)) \sim \bigoplus_f A_f^{m_f}$; taking $N = p^{2k}$ keeps the bad set fixed at $\{p\}$, exactly as Theorem 7.4 demands. This half is closed. The 10.5 condition (i) is therefore already satisfied and the assertion that "(I′) needs a new Prym construction" is withdrawn.*

(I′b) Weighted single-curve realisation. *Realising an interior κ^* with $w \ge w_{\min}(\kappa^*)$ on one curve still requires both CM and GL_2 summands with a tunable conductor-mass ratio. 00544v2 Theorem 7.8 excludes the product shortcut: the theta divisor of $J(C_1) \times J(C_2)$ is reducible whereas the theta divisor of any smooth Jacobian is irreducible, so no product ppav is a Jacobian. Hence the decomposition must come from a genuine cover. This half remains open and is subsumed under the existing* C *item of Problem 10.12; no new* C *is created.*

Proof. For (I′a), Eichler–Shimura identifies $H^1(X_0(N), \mathbb{C})$ with $S_2(\Gamma_0(N)) \oplus \overline{S_2}$, and Atkin–Lehner theory splits S_2 into newform subspaces of levels $d \mid N$; strong multiplicity one gives the isogeny decomposition of $J(X_0(N))$ into the abelian varieties A_f attached to newforms, which are GL_2-type and have conductor dividing N. With $N = p^{2k}$ all bad reduction is at p and $\hat{f} \to \infty$ as $k \to \infty$, so this is a legitimate tower in the sense of Definition 1.1 and matches Theorem 7.4. For (I′b), Theorem 7.8 is verbatim in 00544v2 §7.6.1 with the irreducible/reducible theta-divisor proof; consequently the convexity construction of Theorem 7.4 lives in the AV category and cannot be transported to a single curve without a cover. The category warning of §7.6.1 is unchanged. □

Proposition 10.11 (Section 6 Heisenberg table recomputed; non-CM ≠ high R)**.** *For the non-abelian Heisenberg cover $H_N \to F_N$ of 00544v2 §6.1 the three published rows are reproduced by the exact identities*

$$g_F = \frac{(N-1)(N-2)}{2}, \quad g_X = \frac{N^2(N-3)}{2} + 1, \quad g_X - 1 = N(g_F - 1), \quad \dim \mathrm{Prym} = g_X - g_F,$$

and the "non-CM fraction" is exactly $\dim \mathrm{Prym}/g_X$.

$$(g_F, g_X, \dim \mathrm{Prym}, \textit{fraction}) = (6, 26, 20, 0.769231),\ (45, 485, 440, 0.907216),$$

$$(4950, 499850, 494900, 0.990097),$$

matching 0.769/0.907/0.990. *Nevertheless §6.2 gives* $\kappa_X = \frac{c+1}{12c} \to \frac{1}{12}$ *with* $c = F_X/(ng_X) \geq 2$, *hence* $R = 12\kappa_X - 1 = 1/c \to 0$*: the Prym is* strongly non-CM *but has* $R \to 0$. *Therefore non-CM is insufficient for (I′); this is the second separation of the kind first recorded in Appendix 9.1 (non-CM ≠ non-automorphic).*

Proof. Riemann–Hurwitz for the étale cyclic N-cover gives $g_X - 1 = N(g_F - 1)$; solving with $g_F = (N-1)(N-2)/2$ yields the stated g_X. $\dim \mathrm{Prym} = g_X - g_F$ is the Jacobian decomposition; dividing by g_X reproduces the three printed fractions to six decimals. For $\kappa_X = (c+1)/(12c)$, $R = (c+1)/c - 1 = 1/c$, which tends to 0 for $c \to \infty$; at $c = 2$ one gets $\kappa = 1/8$, $R = 1/2$. The Prym is strongly non-CM because $\mathrm{End}(\mathrm{Prym}) \supseteq M_N(\mathbb{Q}(\zeta_N))$ (§6.1), yet its height/conductor ratio is the *floor*. All evaluations are exact fractions. □

Theorem 10.12 (ρ and the weight dictionary)**.** *The conductor-mass ratio of Theorem 7.4 is*

$$\rho = \frac{6p}{p+1} = 4.5,\ 5,\ 5.25,\ 5.5,\ 5.941176 \quad \textit{for } p = 3, 5, 7, 11, 101, \qquad \rho \to 6,$$

reproducing the numerical table of §7.4. The dictionaries

$$t = \frac{b}{a\rho}, \quad w = \frac{t}{1+t}, \quad \kappa = \frac{1}{8} + \frac{5w}{24}, \quad R = \frac{w}{2} + \frac{1}{2}, \quad b/a = t\rho$$

are mutually consistent; in particular $\kappa = 1/4 \iff t = 3/2 \iff w = 3/5 \iff b/a = 9p/(p+1)$.

Proof. From the §7.4 derivation, $\rho_k = \frac{p^{2k} k \log p}{p^{2k-1}(p+1)(2k)\log p/12} = \frac{12p}{2(p+1)} = \frac{6p}{p+1}$, with the stated exact values. The remaining identities are elementary substitutions already verified in Theorem 10.8; all are exact rational numbers. □

Corollary 10.13 (Restated ledger)**.** *The open items are now: (I′b) weighted single-curve interior realisation (≡ Problem 10.12, existing* C*); (II) general* GL_2*/non-*GL_2 *towers with* $R \leq 3$ *(general arithmetic BMY / generalised Szpiro, unchanged); and the convexity warning of §7.6.1. (I′a) is closed by the newform decomposition of* $J(X_0(p^{2k}))$. *10.1 is downgraded to an equivalent rewriting of Theorem 7.4; 10.3 is qualified to countable-dense exact attainment; 10.5(i) is closed; the convexity identity 10.2, the finite-layer strictness 10.4 and the* $w_{\min}$ *bound are unchanged.*

Proof. Immediate from Theorems 10.8–10.12 and Theorem 10.5. No statement of the pointwise side is touched. □

Remark 10.14 (Status and boundaries). 10.8–10.13 are U relative to 00544v2 §6–§7 (Thm 7.4, 7.8, 7.9, Convention 7.5, Appendix B E19/E27), the Eichler–Shimura / Atkin–Lehner newform decomposition, and Appendices 2.1–10.1. They are not absolutely unconditional. The theta-divisor argument of Theorem 7.8 and the non-CM criterion of §6.1 are text-level checks, not re-proofs. The OCR reading "$b/(a\rho) = 2/3$" is corrected to 3/2 on exact-fraction grounds; if the published PDF differs, only the label of t changes and every numerical conclusion above is unaffected. No new external citation, no new C input and no new No-Go is introduced; (I′b) is absorbed into the already recorded Problem 10.12 item. This appendix does not constitute an independent verification of the proof of Theorem A. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, 47/3, $C \geq 36/7 > 3$, Routes 1–3, $Th.A^*$ and *abc* open are unchanged.

10.3 Attempting the cover: three explicit constructions, multiplicity–conductor cancellation, and the Prym dilution bound

Appendix 10.2 splits (I′) into (I′a), closed by the newform decomposition of $J(X_0(p^{2k}))$, and (I′b), the single-curve interior realisation with a tunable conductor-mass weight. This appendix attempts the construction requested for (I′b): a nonabelian Galois cover whose Chevalley–Weil decomposition contains both CM and GL_2 summands with adjustable mass ratio. Three constructions are computed exactly; the outcome is a genuine reduction but not a solution.

Theorem 10.15 (Construction A: swapping the Fermat coordinates gives no new curve)**.** *Let $F_N : x^N + y^N = 1$ and $K = (\mathbb{Z}/N)^2$ act by $(x, y) \mapsto (\zeta_N^i x, \zeta_N^j y)$. The involution $\sigma : (x, y) \mapsto (y, x)$ normalises K, so $G = K \rtimes S_2$ (order $2N^2$, nonabelian for $N \geq 3$) acts on F_N and permutes $\widehat{K}$ by $(i, j) \mapsto (j, i)$. Mackey's little-group method gives*

$$\begin{gathered}\mathrm{Irr}(G) = \{\, 2N \text{ one-dimensional}\,\} \cup \\ \{\, \tfrac{N^2-N}{2} \text{ two-dimensional } \rho_{[i,j]} = \mathrm{Ind}_K^G \chi_{i,j} \,\}, \\ \sum_{\rho} (\dim \rho)^2 = 2N^2 = |G|.\end{gathered}$$

The Kummer character set of the three-point cover is $\Phi = \{(i, j) : i, j \neq 0,\ i + j \not\equiv 0\ (N)\}$, $|\Phi| = (N - 1)(N - 2) = 2g$, each character with Chevalley–Weil multiplicity 1. As a G-module

$$\begin{gathered}H^1(F_N, \mathbb{C}) = \bigoplus_{i:\, 2i \neq 0} \chi_{i,i}^{\pm} \oplus \\ \bigoplus_{[i,j],\, i \neq j} \rho_{[i,j]}, \\ \#\{\text{diagonal}\} = N - \gcd(2, N) + 1, \\ \#\{\text{two-orbits}\} = \tfrac{(N-1)(N-\gcd(2,N)-1)}{2}.\end{gathered}$$

Every summand is monomial (induced from a one-dimensional character of K), hence of CM type; therefore by the CM–Parseval bound of Appendix 8.1, $\kappa \leq 1/8$ and $w_{\mathrm{GL}_2} = 0$. Moreover $X = F_N$ is unchanged: changing the group only regroups the same Jacobian. A Chevalley–Weil construction for (I′b) must change the curve, not merely the group.

Proof. The three inertia groups of $F_N \to \mathbb{P}^1$ are the kernels of $i, j, i + j$, so Φ is exactly the set of characters ramified at all three points and Chevalley–Weil gives $m_\chi = -2 + \sum_P (1 - \delta_{\chi|I_P}) = 1$. σ-orbits on $\widehat{K}$: N fixed points (i, i) and $(N^2 - N)/2$ orbits of size 2; the former extend in two ways ($\sigma = \pm 1$), the latter induce irreducibly (the little group is K). Summing squares: $2N + 4(N^2 - N)/2 = 2N^2$. Inside Φ, a diagonal (i, i) survives iff $2i \not\equiv 0$; for odd N this is all $i \neq 0$, so the diagonal count is $N - 1$ and the off-diagonal character count is $2g - (N - 1) = (N - 1)(N - 3)$, giving $(N - 1)(N - 3)/2$ two-orbits; the printed formula unifies even N. Exact checks at $N = 3, 5, 7, 9, 11, 13, 25, 27$ give $(g, |\Phi|, \mathrm{diag}, \mathrm{orbits}) = (1, 2, 2, 0)$, $(6, 12, 4, 4)$, $(15, 30, 6, 12)$, $(28, 56, 8, 24)$, $(45, 90, 10, 40)$, $(66, 132, 12, 60)$, $(276, 552, 24, 264)$, $(325, 650, 26, 312)$ with $\dim H^1 = 2g$ each time. Each $\rho_{[i,j]}$ is monomial by construction, so its ℓ-adic realisation has CM by a cyclotomic field; $\kappa \leq 1/8$ follows from $\mu \leq g\hat{f}/4$ of Appendix 8.1. □

Theorem 10.16 (Construction B: CM newforms are exponentially scarce in a prime-power level)**.** *For $N = p^n$ ($p \geq 5$), the Atkin–Lehner decomposition $S_2(\Gamma_0(p^n)) = \bigoplus_{j \leq n} \bigoplus_{f \text{ new level } p^j} f^{\oplus(n-j+1)}$ gives, with the genus formula $g(X_0(p^n)) = 1 + \frac{\mu}{12} - \frac{\nu_2}{4} - \frac{\nu_3}{3} - \frac{\nu_\infty}{2}$, $\mu = p^n + p^{n-1}$, $\nu_2 = 2$ if $p \equiv 1\ (4)$ else 0, $\nu_3 = 2$ if $p \equiv 1\ (3)$ else 0, $\nu_\infty = \sum_{k=0}^{n} \varphi(p^{\min(k, n-k)})$,*

$$\dim S_2^{\mathrm{new}}(\Gamma_0(p^n)) = g(p^n) - 2g(p^{n-1}) + g(p^{n-2}) \sim \frac{(p+1)p^{n-1}}{12}.$$

A weight-2 CM newform of level p^n requires a fundamental discriminant D with $|D|f^2 = p^n$. For $p \equiv 1$ (4) the only candidate is $D = -4p$, whose 2^2 cannot divide the odd number p^n: hence there are exactly zero *CM newforms at every level p^n. For $p \equiv 3$ (4), $D = -p$ forces $f = p^{(n-1)/2}$ and n odd; the number of Hecke characters of conductor f is $O(\varphi(f)) = O(p^{n/2})$ and it is exactly* 0 *for n even. Consequently the CM conductor-mass fraction satisfies*

$$w_{\mathrm{CM}}(n) = O(p^{-n/2}) \longrightarrow 0,$$

and takes only countably many values accumulating at 0*; the modular tower $X_0(p^n)$ therefore approaches the single point $\kappa \to 1/3$ and supplies no tunable interior weight.*

Proof. Exact rational computation of $g(p^n)$ and the second difference for $p = 5, 7, 11, 13$, $1 \le n \le 12$: e.g. $g(5^{12}) = 24\,404\,688$, $\dim_{\mathrm{new}}(5^{12}) = 15\,620\,000$ with CM count 0; $g(7^{12}) = 1\,318\,150\,601$, $\dim_{\mathrm{new}} = 968\,443\,350$, CM bound 0 for even n and $3, 13, 85, 589, 4117, 28813$ for $n = 1, 3, 5, 7, 9, 11$, giving ratios 0.541667, 0.072279, 0.010221, 0.001458, 0.000208; for $p = 11$ the ratios are 3.0, 0.210, 0.018264, 0.001654, 0.000150, 0.000014; for $p = 13$ all CM counts are 0. The discriminant criterion is standard CM theory: a newform with CM by K has level $|D_K|\, N_{K/\mathbb{Q}}(\mathfrak{f})$, and a fundamental discriminant divisible only by p is $-p$ ($p \equiv 3 \bmod 4$) or $-4p$ ($p \equiv 1 \bmod 4$), the latter forcing $4 \mid N$. The $O(p^{n/2})$ count is $\#\{\mathfrak{f}\}$ times the class number, both polynomial in $p^{(n-1)/2}$. Since $\dim_{\mathrm{new}} \asymp p^n$, the ratio decays as $p^{-n/2}$; each non-CM newform has conductor p^j with $j \le n$, so the conductor mass is dominated by the top layers. $\kappa \to 1/3$ is the $X_0(N)$ anchor of 00544v2 §2.4. □

Theorem 10.17 (Construction C: multiplicity–conductor cancellation in the fibre product)**.** *Let $C_1 \to Y$ and $C_2 \to Y$ be finite covers of a fixed base curve (take $Y = \mathbb{P}^1$ via the j-line and the Kummer t-line), of degrees d_1, d_2, with $G = \mathrm{Gal}(\widetilde{X}/Y)$ the product of the two Galois closure groups. Put $X = C_1 \times_Y C_2$ (or its Galois closure). Then $J(X) \sim J(C_1)^{d_2} \oplus J(C_2)^{d_1} \oplus P$, where the Prym P carries the $\rho_1 \otimes \rho_2$ tensor summands. Artin conductor additivity gives*

$$\log \mathrm{cond}\, J(X) = d_2\, L_1 + d_1\, L_2 + L_P, \qquad L_i = \log \mathrm{cond}\, J(C_i) = g_i \hat{f}_i, \quad g_i \asymp d_i.$$

Hence the visible CM-to-GL_2 conductor-mass ratio is

$$\boxed{\frac{d_2 L_1}{d_1 L_2} = \frac{\hat{f}_1}{\hat{f}_2} \cdot (1 + o(1))}\,:$$

the Chevalley–Weil multiplicities $d_1 d_2$ cancel on the two sides. In particular the free integer multiplicities a, b of the Reindexing tower (Theorem 10.8) have no curve-side analogue: in a genuine cover they are replaced by degrees, and degrees cancel. Theorem 7.8 of 00544v2 does not *exclude this construction, because $J(X)$ contains P and is not the product ppav.*

Proof. For $G = G_1 \times G_2$, $H^1(X) = \bigoplus_{\rho_1,\rho_2} \rho_1 \otimes \rho_2^{\,m_{\rho_1} m_{\rho_2}}$ by the Künneth/Chevalley–Weil decomposition of a fibre product; the terms with one trivial factor give $J(C_1)^{d_2}$ and $J(C_2)^{d_1}$, the rest give P. Artin conductors are additive over direct sums of ℓ-adic representations, so $\log \mathrm{cond}\, J(X) = \sum_{\rho_1,\rho_2} m_{\rho_1} m_{\rho_2}\, a(\rho_1 \otimes \rho_2) \log S$; the two marginal sums are $d_2 L_1$ and $d_1 L_2$. Riemann–Hurwitz gives $g_i - 1 = d_i(g_Y - 1) + \frac{1}{2}\sum_P (e_P - 1)$, so $g_i = d_i \Theta(1)$ and the ratio is $\hat{f}_1/\hat{f}_2$ up to the $O(1/d_i)$ correction. The theta-divisor argument of Theorem 7.8 applies to $J(C_1) \times J(C_2)$ with the product polarisation; $J(X)$ has P as an extra isogeny factor, so it is not that ppav. The bad set is the union of the two bad sets, fixed when both are fixed. □

Theorem 10.18 (Visible weight is dense)**.** *Take $C_1 = F_{p^n}$ (CM, $\hat{f}_1 = 2n \log p$) and $C_2 = X_0(q^{2m})$ (GL_2, $\hat{f}_2 = 6m \log q$), with $p \ne q$ primes so that $S = \{p, q\}$ is fixed. Then*

$$\frac{\hat{f}_1}{\hat{f}_2} = \frac{n}{m} \cdot \frac{\log p}{3 \log q}, \qquad \log \frac{w}{1 - w} = \log \frac{n}{m} + \log \frac{\log p}{3 \log q},$$

and since $\log p/\log q \notin \mathbb{Q}$, *Kronecker's theorem makes the visible weight* $w = \hat{f}_1/(\hat{f}_1 + \hat{f}_2)$ *dense in* $(0,1)$ *as* (n,m) *ranges over* $\mathbb{Z}^2_{>0}$. *Enumeration for* $p = 3, q = 5$ *gives maximal gaps* 0.092936, 0.070795, 0.055724, 0.045003, 0.031117, 0.019832, 0.011111 *for* $M = 4, 6, 8, 10, 14, 20, 30$.

Proof. $\hat{f}_1 = 2n\log p$ is the Fermat normalised conductor of Appendix 5.8; $\hat{f}_2 = 6m\log q$ is the $X_0(q^{2m})$ value ($\log \mathrm{cond} = 3g\log(q^{2m})$, 00544v2 §2.4). The map $(n,m) \mapsto \log(n/m)$ has dense image modulo $\log(\log p/3\log q)$ because $\log p/\log q$ is irrational (unique factorisation of integers). All seven gap values are exact finite enumerations. □

Theorem 10.19 (Prym dilution caps the fibre product below 1/4)**.** *The tensor summands satisfy, by the Artin-conductor inequality for* $\rho_1 \otimes \rho_2$, $L_P = c\,d_1 d_2(\hat{f}_1 + \hat{f}_2)$ *with* $c = \Theta(1)$, *hence the* GL_2 *share is*

$$w_{\mathrm{GL}_2} = \frac{\hat{f}_2}{(1+c)(\hat{f}_1+\hat{f}_2)} < \frac{1}{1+c} \le \frac{2}{3}.$$

If the tensor Prym is of CM/mixed-CM type, so that $\kappa_P \le 1/8$ *by the CM–Parseval bound (Appendix 8.1), the convex hull gives*

$$\kappa_X \le \frac{1}{8} + \frac{5}{24}\,w_{\mathrm{GL}_2} < \frac{1}{8} + \frac{5}{24(1+c)} \le \frac{11}{48} = 0.2291\overline{6} \quad (c \ge 1),$$

and for $c = 1, 2, \frac{1}{2}$ *the numerical maxima are* 0.227995, 0.193663, 0.262327. *Therefore the fibre product reaches a dense subset of the* lower *interior* $(1/8, \approx 1/4)$ *but cannot reach* $(1/4, 1/3)$*: the three upper anchors* $1/6, 5/24, 1/3$ *are not all reachable this way.*

Proof. $J(X)$ is the isogeny direct sum of the CM, GL_2 and Prym parts, so κ is the conductor-mass weighted average of $\kappa_{\mathrm{CM}} \le 1/8$, $\kappa_{\mathrm{GL}_2} \to 1/3$ and κ_P (Theorem 10.17 and Appendix 3.3). Subtracting the common 1/8 leaves at most $(5/24)w_{\mathrm{GL}_2}$. With $L_P = c\,d_1 d_2(\hat{f}_1 + \hat{f}_2)$ the GL_2 mass fraction is $\hat{f}_2/[(1+c)(\hat{f}_1+\hat{f}_2)] < 1/(1+c)$. For $c \ge 1$ the bound is $\le 11/48 < 1/4$; the three tabulated values are exact maximisations over $1 \le n, m \le 20$. The condition $\kappa_P \le 1/8$ is the remaining analytic hypothesis; mixed CM heights computed in Appendices 7.1–7.5 all satisfy $\kappa < 1/8$, and the Kummer tensor is of exactly that type, but the $\mathrm{GL}_2 \otimes \mathrm{CM}$ tensor has not been computed. □

Corollary 10.20 ((I′b) reduced to one scalar)**.** *The three constructions exhaust the natural mechanisms:* (A) *nonabelian regrouping of a CM curve gives* $w_{\mathrm{GL}_2} = 0$*;* (B) *a pure* GL_2 *tower gives* $w_{\mathrm{CM}} \to 0$*;* (C) *the fibre product gives a dense visible weight but the Prym factor dilutes it to* $\kappa < 11/48$ *conditionally. Hence (I′b) is no longer "find a cover": the cover is exhibited, and the single remaining unknown is*

$$\boxed{\kappa_P = \frac{h_F(P)}{\log \mathrm{cond}(P)} \quad \text{for } P \sim \bigoplus_{\rho_1 \ne 1, \rho_2 \ne 1} A_{\rho_1 \otimes \rho_2}},$$

the average Colmez/Bergman height of the $\rho_1 \otimes \rho_2$ *tensor factor. This is the same analytic input as the mixed-CM height of Appendices 7.1–7.5 and the upper-half scalar of Appendix 7.1. Consequently Problem 10.12 remains open, but its curve-theoretic part is now reduced to a height computation;* (II) *(general* $R \le 3$*) and the convexity warning are unchanged.*

Proof. Immediate from Theorems 10.15–10.19 and Theorem 10.10. No new No-Go is created and12 item, and all external inputs are the standard Chevalley–Weil, Artin-conductor, Atkin–Lehner and CM theories already used in Appendices 9.1–10.2. □

Remark 10.21 (Boundaries). 10.15–10.20 are U relative to 00544v2 Thm 4.7/7.4/7.8/7.9, Chevalley–Weil, Artin conductor additivity, Mackey theory, Atkin–Lehner newform theory, CM discriminant theory and Appendices 2.1–10.2; they are not absolutely unconditional. The fibre-product connectivity, the Galois closure of the non-Galois cover $X_0(q^{2m}) \to X(1)$, the Artin tensor constant $c = \Theta(1)$, and above all κ_P are C; the CM-newform counts for $p \equiv 3$ (4) are generous upper bounds, not LMFDB exact values. This appendix does not verify the proof of Theorem A. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, 47/3, $C \geq 36/7 > 3$, $\kappa \leq 1/3$ on modular towers, Routes 1–3, $Th.A^*$ and *abc* open are unchanged.

10.4 Computing κ_P: corrected multiplicities, Prym dominance, and a conditional solution of (I′b)

Appendix 10.3 exhibited the fibre product $X = F_{p^n} \times_{\mathbb{P}^1} X_0(q^{2m})$ and reduced (I′b) to the single scalar $\kappa_P = h_F(P)/\log \mathrm{cond}(P)$, where P is the tensor isogeny factor. This appendix computes that scalar. The computation forces two corrections to Appendix 10.3 and one refutation.

Theorem 10.22 (Corrected Chevalley–Weil multiplicities)**.** *Let $C_1, C_2 \to \mathbb{P}^1$ be Galois covers with disjoint branch loci and groups G_1, G_2, and let $\tilde{X}$ be the fibre product with group $G = G_1 \times G_2$. For irreducible $\rho_i \in \mathrm{Irr}(G_i)$ write $d_i = \dim \rho_i$ and let $m_{\rho_i}^{(i)}$ be the virtual Chevalley–Weil multiplicity of ρ_i in $H^1(\tilde{C}_i)$ (so $m_{\mathbf{1}}^{(i)} = -2$). Then*

$$m_{\rho_1 \otimes \rho_2} = d_2\, m_{\rho_1}^{(1)} + d_1\, m_{\rho_2}^{(2)} + 2d_1 d_2,$$

and the curve $X = \tilde{X}/(H_1 \times H_2)$ satisfies

$$\boxed{J(X) \sim J(C_1) \oplus J(C_2) \oplus P},$$

each visible Jacobian occurring once*, not d_2 and d_1 times as stated in Theorem 10.17. The p-side characters contribute $N_\chi = |\widehat{G_1} \setminus \mathbf{1}| = 2g_1 + O(\sqrt{g_1})$ reps and the tensor multiplicity for a one-dimensional χ paired with a two-dimensional ρ_f is $m_{\chi \otimes \rho_f} = m_f^{(2)} + 6$.*

Proof. For a Galois cover of $\mathbb{P}^1$ with branch points B the virtual Chevalley–Weil character is $m_\rho = -2\dim\rho + \sum_{P \in B}(\dim\rho - \dim\rho^{I_P})$. On $\tilde{X}$ the branch locus is $B_1 \sqcup B_2$ with $I_P \subseteq G_1 \times \{1\}$ for $P \in B_1$ and $I_P \subseteq \{1\} \times G_2$ for $P \in B_2$, hence $\dim(\rho_1 \otimes \rho_2)^{I_P} = d_2 \dim \rho_1^{I_P}$ on B_1 and $d_1 \dim \rho_2^{I_P}$ on B_2. Since $\sum_{B_i}(d_i - \dim \rho_i^{I_P}) = m_{\rho_i}^{(i)} + 2d_i$, summing gives the displayed formula. The $\rho_2 = \mathbf{1}$ sector has $d_2 = 1$ and $m_{\mathbf{1}}^{(2)} = -2$, hence multiplicity $m_{\rho_1}^{(1)}$: exactly one copy of $H^1(C_1)$; symmetrically for C_2. Exact verification on six tame toy models $G_i = (\mathbb{Z}/a)^2 \times (\mathbb{Z}/b)^2$ with three branch points each: for $(a,b) = (3,5), (3,7), (5,7), (4,5), (3,11), (5,11)$ the total $\sum m_\rho \dim \rho$ equals $2g_{\mathrm{RH}}$ from Riemann–Hurwitz (values $542, 1136, 3642, 1062, 2972, 9462$), and the visible dimensions equal $2g_1 = (a-1)(a-2)$ and $2g_2 = (b-1)(b-2)$ exactly. For the real tower $m_\chi^{(1)} = 1$ (Theorem 10.15), $d_1 = 1$, $d_2 = 2$, so $m = 2 \cdot 1 + m_f^{(2)} + 4 = m_f^{(2)} + 6$. □

Theorem 10.23 (Prym dominance; 10.18 is subleading)**.** *With $\hat{f}_1 = 2n \log p$, $\hat{f}_2 = 6m \log q$, $L_1 = g_1 \hat{f}_1$, $L_2 = g_2 \hat{f}_2$, and Artin's tensor-conductor formula $a(\chi \otimes \rho_f) = 2a(\chi) + a(\rho_f)$ (the bad primes $p \neq q$ have disjoint ramification),*

$$L_P^{(p)} = 2M_2\, L_1, \qquad L_P^{(q)} = N_\chi\, L_2', \qquad M_2 = \sum_f (m_f^{(2)} + 6) = g_2 + O(g_2/m), \quad L_2' = L_2 + O(g_2),$$

$$\boxed{w_P = \frac{L_P}{L_X} = 1 - O(g_1^{-1} + g_2^{-1}) \to 1, \qquad r := \frac{L_P^{(p)}}{L_P^{(q)}} = \frac{\hat{f}_1}{\hat{f}_2} = \frac{2n \log p}{6m \log q}\,(1 + o(1))}.$$

Consequently $\kappa_X = \kappa_P + O(w_{\text{vis}})$ *and the Kronecker-dense visible weight of Theorem 10.18 is a* subleading *correction; the same ratio* r *survives as the leading parameter of the Prym. Numerically* $w_{\text{vis}} = 1.58 \cdot 10^{-2},\ 1.31 \cdot 10^{-4},\ 1.54 \cdot 10^{-6},\ 1.89 \cdot 10^{-8}$ *at* $(n, m) = (2, 2), (4, 4), (6, 6), (8, 8)$ *for* $(p, q) = (3, 5)$*, and* $r - \hat{f}_1/\hat{f}_2 \le 2.3 \cdot 10^{-16}$ *over* 144 *pairs.*

Proof. Artin conductors are additive over isogeny direct sums; for $V = \chi$ (dimension 1, conductor $a(\chi)$) and $W = \rho_f$ (dimension 2, conductor $a(\rho_f)$) with disjoint ramification, $a(V \otimes W) = \dim W\, a(V) + \dim V\, a(W) = 2a(\chi) + a(\rho_f)$. Summing $m_{\chi \otimes \rho_f}$ over all χ, f with $\mu_1 = \sum_\chi m_\chi^{(1)} a(\chi) \log p = L_1$ and $\mu_2' = \sum_f (m_f^{(2)} + 6) a(\rho_f) \log q = L_2 + O(g_2)$ gives $L_P^{(p)} = 2M_2\mu_1$, $L_P^{(q)} = N_\chi \mu_2'$. Atkin–Lehner multiplicities give $\sum_f m_f^{(2)} = g_2$ and the number of newforms is $O(g_2/m)$, so $M_2 = g_2 + O(g_2/m)$; $N_\chi = p^{2n} - 1 = 2g_1 + O(\sqrt{g_1})$. Dividing, $r = 2M_2 g_1 \hat{f}_1/(2g_1 L_2') = \hat{f}_1/\hat{f}_2 + O(1/m)$. Visible mass $L_1 + L_2$ is $O(g_1\hat{f}_1 + g_2\hat{f}_2)$ while $L_P = \Theta(g_1 g_2(\hat{f}_1 + \hat{f}_2))$, giving $w_{\text{vis}} = O(1/g_1 + 1/g_2)$. Faltings height and conductor are additive under isogeny, hence κ_X is the conductor-weighted convex combination (Appendix 3.3). □

Theorem 10.24 (The Prym is dominated by non-CM factors; 10.19 is refuted)**.** *Every isogeny factor of* P *is* $A_{f\otimes\chi}$*, the modular abelian variety attached to the newform* $f \otimes \chi$ *of level* $q^d p^{2c}$ *with nebentypus* $\omega_f \chi^2$*, where* f *runs over newforms of level* $q^d \le q^{2m}$ *and* χ *over characters of* $(\mathbb{Z}/p^n)^2$*. The factor has CM if and only if* f *has CM. By Theorem 10.16 the CM newforms at level* q^d *are exactly* 0 *for* $q \equiv 1$ (4) *and* $O(q^m)$ *for* $q \equiv 3$ (4)*, whereas* $\dim S_2^{\text{new}} \sim (q+1)q^{2m-1}/12$*; hence the CM conductor-mass fraction of* P *is* $O(q^{-m}) \to 0$ *and* P *is dominated by* non-CM GL_2*-type factors. Therefore the CM–Parseval bound* $\kappa \le 1/8$ *of Appendix 8.1 does not apply to* P*: the hypothesis* $\kappa_P \le 1/8$ *of Theorem 10.19 is false, and its conclusion* $\kappa_X < 11/48$ *is withdrawn.*

Proof. Twisting a newform by a Dirichlet character of conductor p^c with $p \nmid q$ produces a newform of level $q^d p^{2c}$ (local conductor $2c$ at p, since the Galois representation is a two-dimensional unramified representation tensored by a ramified character), and CM is preserved exactly when the original form has CM. Isogeny factors of $J(X)$ are classified by the Chevalley–Weil decomposition of Theorem 10.22. The counting is Theorem 10.16. Since κ is a conductor-weighted average over isogeny factors, the non-CM majority determines κ_P; the CM–Parseval inequality $\mu \le g\hat{f}/4$ of Theorem 8.3 requires a CM type, which these factors do not possess. The barrier $\kappa \le 1/3$ on modular towers (00544v2 §4, Theorem 5.2) remains valid and gives $\kappa_P \le 1/3 + o(1)$ unconditionally. □

Theorem 10.25 (Two-channel convex hull: the value of κ_P)**.** *The height and conductor of* $A_{f\otimes\chi}$ *split into the* q*-channel (the base tower* $X_0(q^{2m})$*) and the* p*-channel (the twist tower, whose characters form the full-level group* $(\mathbb{Z}/p^n)^2$*). Denote by* κ_2 *and* κ_χ *the height/conductor ratios of the two channels. Then*

$$\boxed{\kappa_P(r) = \frac{\kappa_2 + r\,\kappa_\chi}{1+r}, \qquad r = \frac{\hat{f}_1}{\hat{f}_2} = \frac{n \log p}{3m \log q}},$$

with $\kappa_2 = 1/3$ *(the* $X_0(N)$ *anchor,* $\bar{\omega}^2 \sim 3g \log N$*, Theorem 5.2) and* $\kappa_\chi \in \{1/4,\ 5/24\}$*: the value* $1/4$ *corresponds to the full-level tower* $X(N)$ *with* $R \to 2$ *(Grados–von Pippich), and* $5/24$ *to* $X(p^n)$ *with* $R \le 3/2$*. For* $\kappa_\chi = 1/4$*,*

$$\kappa_P(r) = \frac{4 + 3r}{12(1+r)}, \qquad R_P = 12\kappa_P - 1 = 2 + \frac{1}{1+r} \in (2, 3].$$

Exact values: $r = 0 \Rightarrow 1/3$*;* $r = 1/2 \Rightarrow 11/36$*;* $r = 1 \Rightarrow 7/24$*;* $r = 3 \Rightarrow 13/48$*;* $r \to \infty \Rightarrow 1/4$*. All seven anchors of 00544v2 §2.4 satisfy* $R = 12\kappa - 1$ *exactly.*

Proof. Faltings height and conductor are additive under isogeny (Appendix 3.3); the twist increment at p is linear in the Artin conductor increment $2a(\chi)\log p$ because the family $\{\chi\}$ is

the full-level Kummer tower, whose height/conductor ratio is the anchor κ_χ. Hence $h_F(P) = \kappa_2 L_P^{(q)} + \kappa_\chi L_P^{(p)}$ and $\log \mathrm{cond}(P) = L_P^{(q)} + L_P^{(p)}$, giving the convex hull with $r = L_P^{(p)}/L_P^{(q)}$ of Theorem 10.23. The anchors: $R = 0, 1/2, 1, 3/2, 2, 3$ give $\kappa = 1/12, 1/8, 1/6, 5/24, 1/4, 1/3$ respectively, verified as exact rational numbers. The R-identity follows by eliminating r. □

Theorem 10.26 (Conditional solution of (I′b) on the upper half)**.** *Since* $\log p/\log q \notin \mathbb{Q}$*, Kronecker's theorem makes* r *dense in* $(0,\infty)$ *as* (n,m) *ranges over* $\mathbb{Z}^2_{>0}$*, hence* κ_P *is countably dense in* $(1/4, 1/3)$ *for* $\kappa_\chi = 1/4$*, and in* $(5/24, 1/3)$ *for* $\kappa_\chi = 5/24$*, with closure* $[1/4, 1/3]$ *(respectively* $[5/24, 1/3]$*). By Theorem 10.23,* $\kappa_X = \kappa_P + o(1)$*, so the fibre product is a* single-curve *tower with fixed bad set* $\{p, q\}$*,* $g_n \to \infty$ *and* $\hat{f}_n \to \infty$*, attaining a dense set of interior values in the upper half. Maximal gaps for* $(p,q) = (3,5)$*:* 0.003750, 0.001653, 0.000591 *at* $M = 10, 20, 40$*; for* $(p,q) = (11,13)$ *(the parameter range of Theorem 5.2):* 0.003537, 0.001425, 0.000563*. This is consistent with Theorem 7.9 of 00544v2 (countable dense attainment, attainable set not closed). The barrier is saturated only at* $r = 0$*; for* $r > 0$ *one has* $R_P < 3$ *strictly.*

Proof. $\log r = \log(n/m) + \log(\log p/3\log q)$, and $\log p/\log q$ is irrational by unique factorisation; Kronecker gives density of $\{\log r\}$. The map $r \mapsto \kappa_P(r)$ is continuous and strictly decreasing with the stated endpoint limits, so density transfers. The tower axioms are those of Definition 1.1: the bad set is the union of two fixed sets, $g \asymp p^{2n}q^{2m} \to \infty$, and $\hat{f} \asymp n\log p + m\log q \to \infty$. All five numerical series are exact finite enumerations. Theorem 7.9 states the attainable set is not closed and exact attainment occurs only on a countable set; the fibre product supplies precisely such a countable dense set. □

Corollary 10.27 (Net verdict; 10.17–10.19 disposition)**.** (i) *Theorem 10.17 is corrected to* $J(X) \sim J(C_1) \oplus J(C_2) \oplus P$*; its multiplicity–conductor cancellation survives in the form* $r = \hat{f}_1/\hat{f}_2$ *inside the Prym, not between visible factors.* (ii) *Theorem 10.18 survives as a statement about a subleading weight; its ratio is the leading parameter of* κ_P*.* (iii) *Theorem 10.19 and its bound* 11/48 *are withdrawn, because* $\kappa_P \le 1/8$ *fails: even with the weakest possible twist anchor* $\kappa_\chi \ge 1/12$ *one has* $\kappa_P \ge (4+r)/(12(1+r)) > 1/8$ *for every* $r < 5$*, a dense set of towers.* (iv) *(I′b) is solved conditionally on the upper-half interior; (II) general* $R \le 3$ *and the convexity warning of §7.6.1 are unchanged. The lower half* $(1/8, 1/4)$ *is not reached by this construction: a CM⊗CM fibre product is of product CM type and converges to the Heisenberg value* $\kappa \to 1/12$ *(Theorem 7.36), while the intermediate anchors* $1/6, 5/24$ *remain modular-curve anchors.*

Proof. Immediate from Theorems 10.22–10.26, Theorem 10.16, Theorem 9.2, and Theorem 6.8. The weakest-anchor computation is exact: $(4+r)/(12(1+r)) > 1/8 \iff 8 + 2r > 3 + 3r \iff r < 5$. No new No-Go is created and12 item. □

Remark 10.28 (Boundaries). 10.22–10.24 and the density statement of 10.26 are U relative to Chevalley–Weil, Artin conductor additivity, Atkin–Lehner theory, the CM discriminant criterion of Theorem 10.16, the convex hull of Appendix 3.3, 00544v2 §2.4/§4/Thm 5.2/Thm 7.9, and Appendices 2.1–10.3; they are not absolutely unconditional. The numerical values of κ_P are C: they require (a) linearity of the twist-height increment in the Artin conductor increment (a standard but here unreproved Chowla–Selberg/Colmez input), (b) the choice between the two full-level anchors 1/4 and 5/24 (00544v2 §4 warns the three modular families have different normalisations), (c) the extension of the modular BMY theorem to the twist tower of level $q^{2m}p^{2n}$ with nebentypus, (d) connectivity of the fibre product and its Galois closure, and (e) the wild Swan constants at p. The barrier $\kappa \le 1/3$ is unconditional on modular towers. No new external citation, no new C ledger item and no No-Go is added; this appendix does not verify the proof of Theorem A. Fourteen No-Go reasons, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, 47/3, $C \ge 36/7 > 3$, Routes 1–3, $Th.A^*$ and *abc* open are unchanged.

11 Routes toward the general bound $R \le 3$

11.1 Verification of $\kappa_\chi = 1/4$ and channel additivity

Throughout, p, q are distinct primes, χ is a primitive odd Dirichlet character of conductor $f_\chi = p^n$ with $p \nmid q$, and f is a weight-two newform of level q^{2m} with associated two-dimensional representation ρ_f. We write $A_{f\otimes\chi}$ for the abelian variety attached to $f \otimes \chi$, h_F for the stable Faltings height normalized by the genus, and $\kappa = h_F/\log \mathrm{cond}$. The objective is to verify the linearity assumption left open in Appendix FFF, Section 6, to fix the value of κ_χ, and to remove the conditional item (iii) there.

Theorem 11.1 (Gauss-sum height increment)**.** *Let $\tau(\chi) = \sum_{a \bmod f_\chi} \chi(a) e^{2\pi i a/f_\chi}$. Then*

$$\Delta h_F := h_F(A_{f\otimes\chi}) - h_F(A_f) = \tfrac{1}{2}\log f_\chi + O(1) = \tfrac{n}{2}\log p + O(1).$$

Consequently, writing $\hat{h}_\chi(n) = \frac{n}{2}\log p + C_{\mathrm{tw}}$ and $\hat{f}_\chi(n) = (2n+1)\log p$, one has

$$\kappa_\chi(n) = \tfrac{1}{4} + \frac{c}{(2n+1)\log p}, \qquad c = C_{\mathrm{tw}} - \tfrac{1}{4}\log p = O(1),$$

so that $\kappa_\chi(n) - \frac{1}{4} = O(1/n)$ with $\lim_n n(\kappa_\chi(n) - \frac{1}{4}) = c/(2\log p)$.

Proof. The period lattice of $A_{f\otimes\chi}$ is obtained from that of A_f by the equivariant scaling $\tau(\chi)$: the ℓ-adic realization is $V_\ell(A_f) \otimes \chi$, and the comparison isomorphism is twisted by the Gauss sum. The Faltings metric transforms by $\log|\det|$, and the standard identity $|\tau(\chi)|^2 = f_\chi$ for a primitive character gives $\log|\tau(\chi)| = \frac{1}{2}\log f_\chi$. Rank-two normalization contributes no further multiplicative factor, and the archimedean Γ-factor contributes an $O(1)$ term independent of n. Hence $\Delta h_F = \frac{1}{2}\log f_\chi + O(1)$. Dividing by $\hat{f}_\chi = (2n+1)\log p$ and expanding the ratio of two affine functions of n yields the displayed finite-layer formula; the limit of $n(\kappa_\chi(n) - \frac{1}{4})$ is $c/(2\log p)$. □

Theorem 11.2 (Artin conductor increment)**.** *With $p \nmid q$ one has $a(\chi \otimes \rho_f) = 2a(\chi) + a(\rho_f)$, hence*

$$\Delta \log \mathrm{cond}(A_{f\otimes\chi}) = 2\log f_\chi = 2n\log p.$$

Proof. Artin conductor is additive under tensor product with a character whose ramification is supported at primes disjoint from those of ρ_f: $a(\chi\otimes\rho_f) = \dim(\rho_f)a(\chi) + a(\rho_f)$, and $\dim\rho_f = 2$, $a(\chi) = n$ for a primitive character modulo p^n. Taking logarithms of the conductor $p^{2n}\cdot q^{a(\rho_f)}$ gives the claim. This is exactly the identity used in Appendix FFF, Section 2. □

Theorem 11.3 (Value of κ_χ)**.** *$\kappa_\chi = \frac{1}{4}$ and $R_\chi = 12\kappa_\chi - 1 = 2$.*

Proof. By Theorems 11.1 and 11.2 the increment ratio is $\frac{1}{2}\log f_\chi/(2\log f_\chi) = \frac{1}{4}$. By the line $R = 12\kappa - 1$ of Appendix BB, $R_\chi = 12\cdot\frac{1}{4} - 1 = 2$. This agrees with the full-level modular tower $X(N)$, whose limit $R \to 2$ is the independent theorem of Grados–von Pippich quoted in the anchor table; the agreement is therefore a cross-check, not a circular derivation. □

Theorem 11.4 (Channel additivity is unconditional relative to citations)**.** *For the fibre product $X = F_{p^n} \times_{\mathbb{P}^1} X_0(q^{2m})$ of Appendix EEE and its Prym factor P,*

$$\kappa_P(r) = \frac{\kappa_2 + r\kappa_\chi}{1+r} = \frac{4+3r}{12(1+r)}, \qquad r = \frac{L_p}{L_q} = \frac{\hat{f}_1}{\hat{f}_2} = \frac{(2n+1)\log p}{3(2m+1)\log q} = \frac{n\log p}{3m\log q}(1+o(1)),$$

with $\kappa_2 = \frac{1}{3}$. Moreover $R_P = 2 + \frac{1}{1+r} \le 3$ for all $r \ge 0$, with equality only at $r = 0$.

Proof. Chevalley–Weil multiplicities (Appendix FFF, Section 1) give $J(X) \sim J(C_1) \oplus J(C_2) \oplus P$ and $m_{\chi\otimes\rho_f} = m_f^{(2)} + 6$. Faltings height and Artin conductor are additive over isogeny direct sums and over the two disjoint ramification channels, so $h_F(P) = L_q\kappa_2 + L_p\kappa_\chi$ and $\log \mathrm{cond}(P) = L_q + L_p$ with $r = L_p/L_q$. Theorems 11.1 and 11.2 supply $\kappa_\chi = \frac{1}{4}$ on the p-channel, and the $X_0(q^{2m})$ anchor supplies $\kappa_2 = \frac{1}{3}$. Substitution and simplification give $(1/3 + r/4)/(1 + r) = (4 + 3r)/(12(1 + r))$; then $12\kappa_P - 1 = 2 + 1/(1 + r) \le 3$. The finite-layer expression for r follows from $\hat{f}_1 = (2n + 1)\log p$ and $\hat{f}_2 = 3(2m + 1)\log q$ (Appendices RR and FFF); the 144 pairs $(n, m \le 12)$ satisfy the identity with residual 0, and for $n = m$ the asymptotic form differs by at most 2.8×10^{-17}. □

Theorem 11.5 (Uniqueness of κ_χ)**.** *For the full-level Kummer twist channel $(\mathbb{Z}/p^n)^2$ of the fibre product, $\kappa_\chi \equiv \frac{1}{4}$; the alternative $\frac{5}{24}$ listed in Appendix FFF, Section 3, is withdrawn. Consequently the attainable interval of κ_P is exactly $(1/4, 1/3]$, and the weight scale of Appendix CCC satisfies*

$$w(r) = \frac{24\kappa_P(r) - 3}{5} = \frac{5 + 3r}{5(1 + r)} \in \left[\tfrac{3}{5}, 1\right].$$

Proof. The three normalizations of the anchor table are those of X_0, X_1 (coefficient 3), $X(N)$ (coefficient 2), and the three-point tower $X(p^n)$ with $R \le 3/2$. The latter corresponds to $\kappa = 5/24$ and is a non-twisted three-point cover. The p-channel of the fibre product is the full-level Kummer twist tower: $a(\chi) = n$, $a(\chi\otimes\rho_f) = 2n + a(\rho_f)$, and $|\tau(\chi)|^2 = p^n$ hold simultaneously, so Theorems 11.1–11.3 force $\kappa_\chi = 1/4$ with no free normalization. Hence the alternative branch and its interval $(5/24, 1/3]$ are withdrawn. Solving $\kappa = \frac{1}{8} + \frac{5}{24}w$ for $\kappa = \kappa_P(r)$ gives $w(r) = (5 + 3r)/(5(1 + r))$, monotone from 1 at $r = 0$ to $3/5$ as $r \to \infty$; exact rational checks at $r = 0, 1/2, 1, 3, 100$ give $1, 13/15, 4/5, 7/10, 61/101$. □

Theorem 11.6 (No independent BMY for the twist tower is needed)**.** *The barrier $\kappa_P \le 1/3$ for the tower of Theorem 11.4 follows from the modular anchor $\kappa_2 = 1/3$ and convexity alone; conditional item (iii) of Appendix FFF, Section 6, is closed. The general problem of $R \le 3$ for arbitrary* GL_2 */ non-*GL_2 *towers remains open.*

Proof. Since $\kappa_\chi = 1/4 < 1/3 = \kappa_2$ and κ_P is a convex combination of κ_2 and κ_χ with nonnegative weights, $\kappa_P \le \max(\kappa_2, \kappa_\chi) = 1/3$ for every $r \ge 0$; equivalently $R_P = 2 + 1/(1 + r) \le 3$. No assertion about $R \le 3$ for a general twisted GL_2 tower is used. Item (iii) asked whether the modular-tower BMY theorem applies at level $q^{2m}p^{2n}$ with nebentypus $\omega_f\chi^2$; the convexity argument renders that question irrelevant for this construction, while leaving the general problem untouched. □

Remark 11.7 (Numerical verification and status). Exact rational checks: $\frac{1}{2}/2 = \frac{1}{4}$, $12 \cdot \frac{1}{4} - 1 = 2$; six anchors $R = 12\kappa - 1$ give $0, 1/2, 1, 3/2, 2, 3$; the table of $\kappa_P(r)$ at $r = 0, 1/2, 1, 2, 3, \infty$ gives $1/3, 11/36, 7/24, 5/18, 13/48, 1/4$. Density: with $r = n\log p/(3m\log q)$, $(p, q) = (3, 5)$, the interior maximal gaps of the κ-set for $M = 10, 20, 40$ are $0.003750, 0.001653, 0.000591$, exactly reproducing Appendix FFF, Section 4; the finite-layer normalization $(2n + 1)/(3(2m + 1))$ gives $0.002978, 0.001432, 0.000571$. Endpoints $1/4$ and $1/3$ are limit points not attained (Theorem 7.9). All statements are U relative to the Gauss-sum identity, Artin conductor additivity, Chevalley–Weil multiplicities, Faltings height additivity, the modular anchors, and Appendices BB–FFF; the $O(1)$ twist constant C_{tw}, wild Swan constants, and fibre-product connectivity remain C.

11.2 Breakthrough routes for the long-term item (II)

We study item (II): the general arithmetic Bogomolov–Miyaoka–Yau inequality $\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g)$ on arbitrary towers (Problem 10.1). Throughout we use the normalized Wilms Noether identity $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}$, the ratio $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$, $\sigma = \Delta_{\mathrm{Del}}/(g\hat{f})$ with $\hat{f} = \log \mathrm{cond}\,/g$, $\varepsilon_{\mathrm{ar}} = \delta_{\mathrm{Fal}}/(g\hat{f})$, and the master identity $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$. We put $W = 2\log(2\pi^4) = 10.5441334479\ldots$.

Theorem 11.8 (Noether reduction of (II))**.** *For every tower,*

$$\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g) \iff h_F \le \tfrac{1}{3}\Delta_{\mathrm{Del}} + \tfrac{1}{12}\delta_{\mathrm{Fal}} + O(g).$$

Consequently, using Wilms' lower bound $\delta_{\mathrm{Fal}} > -Wg$, a sufficient condition is $h_F \le \frac{1}{3}\Delta_{\mathrm{Del}} - \frac{W}{12}g + O(g)$ with $W/12 = 0.8786777873\ldots$; using the archimedean bound $\delta_{\mathrm{Fal}} \le C\Delta_{\mathrm{Del}} + O(g)$ (Theorem 8.7), a necessary condition is $h_F \le (\frac{1}{3} + \frac{C}{12})\Delta_{\mathrm{Del}} + O(g)$.

Proof. Solve the Noether identity for $\bar{\omega}^2 = 12h_F - \Delta_{\mathrm{Del}} - \delta_{\mathrm{Fal}}$. Then $\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g)$ is equivalent to $12h_F \le 4\Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}} + O(g)$; divide by 12. Insert $\delta_{\mathrm{Fal}} > -Wg$ to obtain a sufficient condition (a smaller right hand side suffices), and $\delta_{\mathrm{Fal}} \le C\Delta_{\mathrm{Del}} + O(g)$ to obtain a necessary one. No inequality is used beyond these two archimedean bounds; the equivalence itself is an identity. □

Theorem 11.9 (Which constant must be pressed)**.** *Assume $\bar{\omega}^2 \le C\Delta_{\mathrm{Del}} + O(g)$ along a tower with $\sigma \to 1$ and $\varepsilon_{\mathrm{ar}} = O(1/\hat{f})$. Then*

$$R \le C + O(1/\hat{f}), \qquad \kappa \le \frac{1+C}{12}\sigma + O(1/\hat{f}) \longrightarrow \frac{1+C}{12}.$$

Hence $\limsup \kappa \le 1/3$ requires $C \le 3$, and $C = 4$ already gives $\kappa \le 5/12 > 1/3$. Moreover, with $\varepsilon_{\mathrm{ar}} \le C_{\mathrm{ar}}\sigma/\hat{f}$ and $\sigma = 1 + s$,

$$\kappa \le \frac{1}{3} + \frac{s}{3} + \frac{C_{\mathrm{ar}}(1+s)}{12\hat{f}},$$

so finite layers may exceed $1/3$ by exactly $O(s) + O(1/\hat{f})$.

Proof. Divide $\bar{\omega}^2 \le C\Delta_{\mathrm{Del}} + O(g)$ by $g\hat{f}$ to get $R \le C + O(1/\hat{f})$. Substitute into the master identity: $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}} \le \sigma(1+C) + O(1/\hat{f})$. Take $\sigma \to 1$ and $\varepsilon_{\mathrm{ar}} = O(1/\hat{f})$; the finite-layer form is the same inequality with $\sigma = 1 + s$ and the explicit $\varepsilon_{\mathrm{ar}}$ bound inserted. Exact rational values: $C = 3, 4, 6, 12, 36, 36/7$ give $(1+C)/12 = 1/3, 5/12, 7/12, 13/12, 37/12, 43/84$. □

Theorem 11.10 (The tower statement is strictly weaker than classical BMY)**.** *Problem 10.1 asks for $\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g)$; Parshin's classical inequality is $\omega^2 \le 3\sum_v \delta'_v \epsilon_v + (2g-2)\log|D_K|$. After dividing by $g\hat{f}$ the former is $R \le 3 + O(1/\hat{f})$, i.e. only $\limsup_n R_n \le 3$ is required; the $O(g)$ term is invisible in the tower limit. No finite-layer inequality is needed.*

Proof. $\Delta_{\mathrm{Del}} = \sigma g\hat{f}$ with $\sigma = O(1)$, so $O(g)/(g\hat{f}) = O(1/\hat{f}) \to 0$ along any tower satisfying (T3). The classical form contains the archimedean discriminant contribution and a base-discriminant term; the tower form absorbs both into $O(g)$. Taking lim sup removes the $O(1/\hat{f})$. □

Theorem 11.11 (An unconditional domain for (II))**.** *Let a tower have Jacobian isogenous to a direct sum of factors each belonging to the proved classes: modular towers $X_0(N), X_1(N), X_0(p^n), X(N)$ (with coefficients $3, 3, 3, 2$), Fermat/CM classes (with $R \le 1/2$), and the hyperelliptic/Belyi classes of Wilms and Kühn–Curilla (with $R = O(1)$ on the covered levels). Then $\limsup R \le 3$ by convexity, and the barrier holds without assuming general arithmetic BMY. If a weight w is carried by a coefficient-3 factor and $1-w$ by a coefficient-2 factor, then*

$$R \le 3w + 2(1-w) = 2 + w, \qquad \kappa \le \frac{3+w}{12}.$$

Proof. Faltings height and logarithmic conductor are additive over isogeny direct sums, hence κ is a convex combination of the factor ratios with nonnegative weights w_ρ (Appendix EE); equivalently $R = 12\kappa - 1$ is the same convex combination of the factor R_ρ when $\sigma \to 1$ and $\varepsilon_{\mathrm{ar}} \to 0$. The modular coefficients 3 and 2 are the theorems quoted in the Problem 10.1 table; Fermat/CM satisfy $R \le 1/2$. Thus $R \le \max_\rho R_\rho \le 3$. For two factors, $R \le 3w + 2(1-w) = 2 + w$ and $\kappa = (1+R)/12 \le (3+w)/12$; exact checks give $w = 0 \Rightarrow R = 2, \kappa = 1/4$ and $w = 1 \Rightarrow R = 3, \kappa = 1/3$. □

Theorem 11.12 (Weak layer: refuting $\kappa \to \infty$)**.** *By §9, no $\kappa \to \infty$ is equivalent to $\sigma = O(1)$ and $R = O(1)$; Theorem 9.1 gives $\bar\omega^2 \le 12h_F + Wg$. Yuan's bidirectional bound $c_3 \max\{h_{\mathrm{Fal}}, 1\} \le [\bar\omega_a^2] \le c_4 \max\{h_{\mathrm{Fal}}, 1\}$ (with $[\bar\omega_a^2] = \bar\omega_a^2/[K:\mathbb{Q}]$), together with de Jong's local formula $(\omega,\omega)_a = \frac{2g-2}{2g+1}\sum_v \chi(X_v)\log Nv$, yields the following verifiable criterion: if $c_3(g)^{-1}c_4(g) = O(1)$ as $g \to \infty$ and the normalization $[\cdot]$ matches $\hat f$ when $[K:\mathbb{Q}] = \varphi(p^n)$ grows, then $h_F = O(g\hat f)$ and $R = O(1)$, closing Problem 10.2 without the constant* 3.

Proof. From $h_{\mathrm{Fal}} \le c_3^{-1}[\bar\omega_a^2]$ and $\bar\omega_a^2 = \bar\omega^2 - \sum_v \varepsilon_v \log Nv$, an $O(g\hat f)$ bound on $[\bar\omega_a^2]$ gives $h_F = O(g\hat f)$; substituting into Theorem 9.1 gives $\bar\omega^2 = O(g\hat f)$, hence $R = \bar\omega^2/\Delta = O(1)$ provided $\sigma = O(1)$. The two displayed conditions are exactly those needed to transfer the bound and to keep the normalization consistent; they are not verified here. □

Theorem 11.13 (Circularity boundary and scope of the negative verdict)**.** *On* GL_2*-type towers the barrier is equivalent to [Arith-BMY] (Corollary 5.4), and the barrier implies BMY strictly when $\sigma > 1$ (Theorem 5.3). Hence a proof of the barrier on that class that assumes BMY is circular, unless the tower class or the conclusion is changed. Moreover §10.4 states that the expected return of further search along* the discriminant channel *is negative and leaves both (D) and Problem 10.1 as open problems; the negative verdict does not apply to the curvature channel (II).*

Proof. The two implications are quoted verbatim. Routes P1–P4 change the tower class (explicit covers, isogeny-factor classes) or the conclusion (weak layer $R = O(1)$, leading term of R), so they do not assume BMY on the same class. Scope of §10.4: the sentence preceding the verdict attributes the structural gap to Mestre's covering only the automorphic half of channel (D); the verdict sentence names the discriminant channel only, and §10.5 explicitly records a second rescue mechanism. Hence item (II) is listed as open, not judged negatively. □

Remark 11.14 (Numerical checks and status). Exact and numerical verifications: $W = 2\log(2\pi^4) = 10.544133447915092$; $W/12 = 0.8786777873262577$; the $C \mapsto (1+C)/12$ table of Theorem 11.9; finite-layer excesses $8.33 \cdot 10^{-4}$, $4.18 \cdot 10^{-3}$, $4.25 \cdot 10^{-2}$, $4.17 \cdot 10^{-4}$ for $(\sigma - 1, \hat f) = (0, 100), (1/100, 100), (1/10, 10), (1/1000, 1000)$ with $C_{\mathrm{ar}} = 1$; the six anchors $R = 12\kappa - 1$ are recovered from $R < 12\kappa/\sigma - 1 + W/(\sigma\hat f)$ as $\hat f \to \infty$; and $w = 0, 1/5, 2/5, 3/5, 1$ give $R = 2, 11/5, 12/5, 13/5, 3$. Literature items (Parshin 1988, Moret-Bailly 1989, Wilms 2016, Yuan 2026 Thm 1.4, Yuan–Zhang effective bounds, Kühn/Curilla–Kühn, Faltings 2021, de Jong 2019, de Jong symmetric-roots, Jorgenson–Kramer, Ballaÿ) are checked at text level only. Conditions on c_3, c_4, on the wild-tower applicability of the semistable asymptotics, and on the p^n extension of the Belyi bounds remain C. No new No-Go, no new mathematical assumption, and no claim of an independent verification of Theorem A.

11.3 Pushing P1 and P2: one-sided Yuan bound, degeneration asymptotics relocated, and a single criterion for (II)

We carry out routes P1 (semistable degeneration asymptotics combined with the dual-graph data of Appendices HH, SS, AAA) and P2 (the genus dependence of the constants in Yuan's Theorem 1.4) of Appendix HHH, and sharpen P3. We keep the normalized Wilms Noether

identity $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}$, the ratios $\kappa = h_F/(g\hat{f})$, $\sigma = \Delta_{\mathrm{Del}}/(g\hat{f})$, $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$, $\varepsilon_{\mathrm{ar}} = \delta_{\mathrm{Fal}}/(g\hat{f})$, and $W = 2\log(2\pi^4) = 10.5441334479\ldots$.

Theorem 11.15 (Three-layer separation and the exact form of Problem 10.2)**.** *For every tower,*

$$12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}, \qquad R = \frac{12\kappa}{\sigma} - 1 - \frac{\varepsilon_{\mathrm{ar}}}{\sigma}.$$

In the audit normalization $\varepsilon_{\mathrm{ar}} = O(1/\hat{f})$ (Appendices LL, MM, UU and the $O(g)$ term of Theorem 8.7) one has $\kappa \le (1+R)/12 + O(1/\hat{f})$, so the barrier $\limsup \kappa \le 1/3$ is equivalent to $\limsup R \le 3$. In the coarse reading $\varepsilon_{\mathrm{ar}} \le C_{\mathrm{ar}}\sigma + o(1)$ of Theorem 8.7 one has $\kappa \le (1+R+C_{\mathrm{ar}})/12$, and the barrier needs $R \le 3 - C_{\mathrm{ar}}$. In either reading $\varepsilon_{\mathrm{ar}}$ is bounded above, hence $\kappa \to \infty$ forces $R \to \infty$: Problem 10.2 is exactly the negation of $R = O(1)$, and the archimedean channel is not an independent problem.

Proof. Divide the Noether identity by $g\hat{f} = \log \mathrm{cond}$ and use $\bar{\omega}^2 = R\sigma g\hat{f}$, $h_F = \kappa g\hat{f}$. Solving for R gives the displayed identity. Insert $\varepsilon_{\mathrm{ar}} = O(1/\hat{f})$, respectively $\varepsilon_{\mathrm{ar}} \le C_{\mathrm{ar}}\sigma + o(1)$, and let $\sigma \to 1$. Since both bounds are uniform in R, a divergent κ can only come from a divergent R. The seven anchors $\kappa = 1/12, 1/8, 1/6, 5/24, 1/4, 7/24, 1/3$ give $R = 12\kappa - 1 = 0, 1/2, 1, 3/2, 2, 5/2, 3$ exactly; for $\sigma = 101/100$, $\varepsilon_{\mathrm{ar}} = -1/1000$, $\kappa = 1/8$ one finds $R = 491/1010$. All computations are exact rational ones. □

Remark 11.16 (Normalization used in the barrier proof)*.* Remark 4.6 of the source text derives $\kappa \le (1/12)(1)(1+3) + o(1) = 1/3$ from $d \le C\sigma + o(1)$ and [Arith-BMY]. This implicitly uses that $d = \varepsilon_{\mathrm{ar}}$ has already been absorbed by the product formula (Theorem 8.7 is titled *the archimedean channel is cancelled*), i.e. the first normalization above. Keeping only the coarse bound with $C_{\mathrm{ar}} > 0$ would give $(4 + C_{\mathrm{ar}})/12 > 1/3$ instead. The two readings are compatible once the cancellation is applied; we record the distinction and do not claim an error in the source.

Theorem 11.17 (Yuan's Theorem 1.4 is one-sided; route P2 is closed)**.** *Let $g > 1$. Theorem 1.4 of Yuan (Annals of Mathematics 203 (2026), 15–119; arXiv:2108.05625, Theorems 4.11 and 4.12) asserts verbatim that there is a constant $c_3 > 0$ depending only on g such that for every smooth projective geometrically connected curve C/K of genus g,*

$$[\bar{\omega}^2_{C/K,a}] > c_3 \cdot \max\{h_{\mathrm{Fal}}(C), 1\}.$$

The full text contains no constant c_4, no statement of a two-sided comparison $[\bar{\omega}^2_a] \asymp_g h_{\mathrm{Fal}}$, and no occurrence of the phrases "equivalent invariants" or "conjectural arithmetic BMY". The sentence "the constants (c_1', c_2') depend only on g and are amazingly explicit" refers to the Looper–Silverman–Wilms constants, not to c_3; no explicit bound for $c_3(g)$ is given. Consequently $h_{\mathrm{Fal}} < c_3^{-1}[\bar{\omega}^2_a]$: obtaining the height inequality $h_F \le \frac{1}{3}\Delta_{\mathrm{Del}} + \frac{1}{12}\delta_{\mathrm{Fal}} + O(g)$ of Appendix HHH would require an independent upper *bound on $[\bar{\omega}^2_a]$, which is item (II) itself. Route P2 and the criterion $c_3(g)^{-1}c_4(g) = O(1)$ of Appendix HHH are therefore withdrawn.*

Proof. This is a text-level verification of the published theorem statement and a full-text scan of the article; no proof of Yuan's theorem is reproduced or used. The direction is that of a lower bound on the admissible self-intersection, i.e. an upper bound on h_{Fal} in terms of $[\bar{\omega}^2_a]$. For $c_3 = 1/10, 1, 10$ the resulting inequalities are $h_{\mathrm{Fal}} < 10[\bar{\omega}^2_a]$, $< [\bar{\omega}^2_a]$, $< \frac{1}{10}[\bar{\omega}^2_a]$, respectively; even $c_3 \ge 3$ leaves the required upper bound on $[\bar{\omega}^2_a]$ unavailable. Hence no choice of $c_3(g)$ closes Problem 10.2 by this route. □

Theorem 11.18 (Degeneration asymptotics belong to the archimedean channel, not to Remark 4.6)**.** *Faltings (J. reine angew. Math. 771 (2021), 65–84) gives asymptotics of the Arakelov metric, the Green functions and the delta-invariant for semistable degenerations over a base of arbitrary*

dimension, involving the local discriminants of the singularities and the dual graph; de Jong (J. Differential Geom. 111 (2019), 241–301) gives the one-dimensional version

$$\delta(\mathcal{X}_t) \sim -(\delta(\Gamma_\pi) + \varepsilon(\Gamma_\pi)) \log|t| - 6g_0(\Gamma_\pi) \log(-\log|t|), \qquad g_0(\Gamma) = b_1(\Gamma),$$

with a combinatorial interpretation via Zhang's admissible Green functions on polarized metrized graphs. The proof of Theorem 8.7 of the source text already quotes the identical Wilms singularity order $\delta_{\mathrm{Fal}} \sim -(\delta + \varepsilon) \log|t| - 6 \log\det \operatorname{Im} \Omega(t)$ *and states that the logarithmic coefficient is tied to the vanishing-cycle length. Hence these asymptotics are the mechanism of the* archimedean channel *(closed in Appendix ZZ), not the missing step of Remark 4.6, which is an upper bound on* $\bar{\omega}^2$ *itself. Through the Noether identity they only give* $\bar{\omega}^2 = 12h_F - \Delta_{\mathrm{Del}} - \delta_{\mathrm{Fal}}$*, which still requires an upper bound on* h_F*.*

Proof. The two asymptotic formulae are quoted at text level. The source text's Theorem 8.7 proof is quoted verbatim from §8.3. Noether is an identity, so a bound on δ_{Fal} alone cannot bound $\bar{\omega}^2$ without a bound on h_F; the Bergman-kernel estimate of Theorem 4.3 concerns a different quantity. Hence the claimed relocation of the tool in Appendix HHH is withdrawn. □

Theorem 11.19 (Dual-graph size of the delta-invariant on mixed Kummer towers)**.** *For the mixed Kummer tower* C_{p^n,q^m} *with* $\gcd(p,q) = 1$*, the stable dual graphs* Γ_p, Γ_q *are trees (Appendix SS), so* $b_1(\Gamma) = 0$ *and the* $-6g_0(\Gamma)\log(-\log|t|)$ *term of de Jong's asymptotic vanishes identically. With* $|V_p| = (p^n + p - 2)/(p-1)$*,* $|E_p| = |V_p| - 1$*, bounded thickness* $N_i \le N_0$ *(Appendix AAA), and tree-graph invariants* $\delta(\Gamma) + \varepsilon(\Gamma) = O(|E|)$*, one gets*

$$\delta_{\mathrm{Fal}} = O\big(|E_p| \log p + |E_q| \log q\big), \qquad \varepsilon_{\mathrm{ar}} \lesssim \lambda\,(\sigma - 1),$$

where λ *is a bounded edge-length convention constant, and* $\sigma - 1 = O(p^{-n}/(q^m \hat{f}) + q^{-m}/(p^n \hat{f}))$*. Thus* $\varepsilon_{\mathrm{ar}}$ *is exponentially small on these towers, and the archimedean upper bound acquires a second derivation independent of Theorem 8.7.*

Proof. Combine the tree property and the vertex count of Appendix SS with the bounded-thickness bound of Appendix AAA: every edge parameter $-\log|t_e| \asymp \log p$ or $\asymp \log q$ is uniformly bounded, and the tree invariant is linear in $|E|$. Dividing by $g\hat{f}$ with $g = \frac{1}{2}(p^n - 1)(q^m - 1)$ and $\hat{f} = \hat{f}_p + \hat{f}_q$ (Appendix 6.24) gives the displayed bound; the same graph mass divided by $g\hat{f}$ is exactly $\sigma - 1$ by Appendix 5.10, so the ratio is the convention constant λ. Numerical verification for $p = 3, q = 5$ and $(n,m) = (4,4), (8,6), (12,12), (20,15), (50,40)$ gives $\sigma - 1 = 5.94 \cdot 10^{-4}, 5.52 \cdot 10^{-6}, 2.41 \cdot 10^{-8}, 2.95 \cdot 10^{-12}, 4.74 \cdot 10^{-27}$, with $\varepsilon_{\mathrm{ar}}$ of the same order for $\lambda = 1$ and $\lambda = \log p$. □

Theorem 11.20 (P1 does not change the leading term of R; (II) is trivial on CM/Kummer towers)**.** *On any tower with* $\sigma \to 1$ *and* $\varepsilon_{\mathrm{ar}} = o(1)$*,* $\limsup R \le 3$ *is equivalent to* $\limsup \kappa \le 1/3$*, independently of the precision of the delta-invariant asymptotics. On CM towers and on mixed Kummer towers Appendices YY and ZZ give* $\kappa \le 1/8$*, hence* $R \le 1/2 + o(1) < 3$ *without any Bogomolov–Miyaoka–Yau input; on modular towers* $R \to 3$ *is a theorem (Dolce–Mercuri, Mayer). Therefore the ranking of P1 in Appendix HHH as the route that closes (II) on Kummer, mixed and CM towers is withdrawn: those cases were already closed by the height bound. P1 contributes a second archimedean derivation and an explicit graph-level remainder; it is needed, if at all, only for towers with* $b_1(\Gamma) > 0$ *or unbounded thickness, where its applicability remains C.*

Proof. The equivalence is Theorem 11.15. Insert the established bounds $\kappa \le 1/8$ (CM–Parseval, Appendices YY and ZZ) and $R \to 3$ (modular towers). Since $\varepsilon_{\mathrm{ar}}/\sigma = O(1/\hat{f})$ is subleading, no refinement of δ_{Fal} can alter $\limsup R$. □

Theorem 11.21 (Exact finite-layer margin of the P3 convex hull)**.** *Assume all isogeny factors of* $\mathrm{Jac}(X_n)$ *lie in the proved classes of Appendix HHH. By the convex hull of Appendix 3.17 and the per-factor bound of Appendix 11.9,*

$$\kappa_X \le \frac{\sigma}{3} + \frac{C_{\mathrm{ar}}\sigma}{12\hat{f}} = \frac{1}{3} + \frac{\sigma-1}{3} + \frac{C_{\mathrm{ar}}\sigma}{12\hat{f}}, \qquad R_X \le 3 + \frac{C_{\mathrm{ar}}+W}{\hat{f}}.$$

The finite-layer excess above 3 *is* $O(1/\hat{f})$ *and is dominated by the archimedean constant* W*, not by* $\sigma-1$*; the factor* σ *cancels in* $12\kappa_X/\sigma$*. Hence the exponentially small* $\sigma-1$ *of Appendix AAA cannot make* $R \le 3$ *hold layer by layer, which is consistent with the limsup formulation of Problem 10.1.*

Proof. Apply the Wilms lower bound $\varepsilon_{\mathrm{ar}} \ge -W/\hat{f}$ to $R = 12\kappa/\sigma - 1 - \varepsilon_{\mathrm{ar}}/\sigma$. With $C_{\mathrm{ar}} = 1$ and $p = 3, q = 5$ the values $(n,m) = (4,4), (8,6), (12,12), (20,15), (50,40)$ give $R \le 3.5800$, 3.3299, 3.1830, 3.1278, 3.0488, matching $3 + (1+W)/\hat{f}$ to 10^{-4}. All other ingredients are exact. □

Theorem 11.22 (Final criterion and surviving routes)**.** *The following are equivalent on towers with* $\sigma \to 1$*:*

$$\limsup \kappa \le \tfrac{1}{3}, \quad \limsup R \le 3, \quad h_F \le \tfrac{1}{3}\Delta_{\mathrm{Del}} + \tfrac{1}{12}\delta_{\mathrm{Fal}} + O(g), \quad \bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g).$$

The constant 3 *is the only target:* $C = 4$ *gives* $\kappa \le 5/12$ *and* $C = 36/7$ *gives* $\kappa \le 43/84$*. Routes P1 and P2 are closed or downgraded, P3 enlarges the proved domain to towers whose isogeny factors are all in proved classes, and P4 serves only the weak layer* $R = O(1)$*. The only surviving positive routes are P5 (the displayed height inequality, i.e. general arithmetic BMY, requiring new analytic input) and P6 (an arithmetic analogue of* Ω^1_X*). The circularity boundary of Theorems 5.3 and 5.4 and the restriction of the negative verdict of §10.4 to the discriminant channel are unchanged.*

Proof. The first three equivalences follow from Theorems 11.15 and Appendix 11.8; the fourth is the definition of (II). The constant table $(1+C)/12$ for $C = 3, 4, 6, 12, 36, 36/7$ is $1/3, 5/12, 7/12, 13/12, 37/12, 43/84$ by exact rational arithmetic. The route assessment collects Theorems 11.17–11.21. No new No-Go, no new mathematical assumption, and no independent verification of Theorem A is claimed. □

Remark 11.23 (Numerical checks and status). Exact rational checks: the seven anchors of Theorem 11.15; the perturbed sample $R = 491/1010$; $W = 10.544133447915092$, $W/12 = 0.8786777873262577$; the $C \mapsto (1+C)/12$ table; the $C_{\mathrm{ar}} = 0, 1, 2$ bounds $(1 + R + C_{\mathrm{ar}})/12$; the $\lambda = 1, \log p$ delta bounds and the P3 margins of Theorems 11.19 and 11.21. Text-level additions: Yuan 2026 Theorem 1.4 (full-text scan: no c_4, no two-sided comparison), Faltings 2021 (Crelle 771, 65–84), de Jong 2019 (JDG 111, 241–301), and §8.3 of the source text. Remaining C items: wild-tower applicability of the semistable asymptotics, the edge constant λ, an explicit $c_3(g)$, the p^n extension of the Belyi bounds, and the normalization note of Remark 11.16. All pointwise conclusions, the fourteen No-Go reasons, and the status of *abc* are unchanged.

11.4 Pushing P3, P5 and P6: three routes converge on the constant four

We carry out the three surviving routes of Appendix HHH and Appendix III: P3 (enlarging the proved domain by isogeny convexity and Belyi-type bounds), P5 (the height inequality equivalent to general arithmetic BMY), and P6 (an arithmetic analogue of Ω^1_X). We keep the normalized Wilms Noether identity $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}$, the ratios $\kappa = h_F/(g\hat{f})$, $\sigma = \Delta_{\mathrm{Del}}/(g\hat{f})$, $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$ and $\varepsilon_{\mathrm{ar}} = \delta_{\mathrm{Fal}}/(g\hat{f})$. Throughout, C denotes a constant in an upper bound $\bar{\omega}^2 \le C\Delta_{\mathrm{Del}} + O(g)$; by Appendix HHH the corresponding barrier is $\kappa \le (1+C)/12$.

Theorem 11.24 (Route P6: the target is three and not four)**.** *In Soulé's IHÉS paper the geometric input is the rank-two Bogomolov inequality $c_1^2(\Omega^1_X) \le 4c_2(\Omega^1_X)$, not $c_1^2 \le 3c_2$; the arithmetic Bogomolov inequality available for the rank-two hermitian extension $\bar{E}$ (Miyaoka, Moriwaki, Soulé) is likewise $\hat{c}_1(\bar{E})^2 \le 4\hat{c}_2(\bar{E})$. The constant 3 belongs to the surface inequality of Yau and Miyaoka (1977), whose equality case forces a ball quotient, and is not part of Soulé's transplantation chain. Consequently route P6 is not the construction of an arithmetic analogue of Ω^1_X alone, but the construction of one that yields the coefficient 3 rather than 4. The two obstructions stated in the source — no arithmetic analogue of Ω^1_X, and the archimedean summand of $\hat{c}_2(\bar{E})$ not being easy to bound from above — remain, and the byproduct $\mu_k \ge \hat{c}_1(\bar{N})^2/(2nd) - A$, $A = (nd)^{-1} + \log(m(n+g-1))$, is a lower bound on successive minima, not an upper bound on $\bar{\omega}^2$.*

Proof. The four statements are quoted verbatim from the IHÉS text: the displayed Bogomolov inequality for Ω^1_X, the sentence "we do not know any arithmetic analog for the sheaf of absolute differentials Ω^1_X", the description of the archimedean summand of $\hat{c}_2(\bar{E})$ as an integral over the complex points of X, and "we are unable to prove Parshin's conjecture by Miyaoka's argument". The arithmetic Bogomolov inequality used there is the rank-two one "à la Bogomolov" for semistable hermitian bundles, with constant 4. The surface inequality $c_1^2 \le 3c_2$ is a separate theorem of Yau and Miyaoka; equality implies that the universal cover is the complex ball. Since the transplantation chain only carries the rank-two constant, any arithmetic analogue obtained inside it can at best reproduce 4. Inserting $C = 4$ into $(1+C)/12$ gives $5/12$, while $C = 3$ gives $1/3$; the difference is $5/12 - 1/3 = 1/12$, the Heisenberg anchor of Appendix BB. All of this is exact rational arithmetic. The direction of the successive-minima bound is opposite to that required by Appendix HHH, so it cannot substitute for an upper bound. □

Theorem 11.25 (Exact normalization of the cyclotomic discriminant)**.** *Let $K = \mathbb{Q}(\zeta_{p^n})$, $d = [K:\mathbb{Q}] = p^{n-1}(p-1)$. Then $\log|D_{K/\mathbb{Q}}| = p^{n-1}\big((p-1)n-1\big)\log p$ and hence*

$$\frac{\log|D_{K/\mathbb{Q}}|}{d} = \Big(n - \frac{1}{p-1}\Big)\log p$$

exactly. In the audit normalization of Appendix NN, $\hat{f}_n = (2n+1)\log p$ and $g_n = \frac{1}{2}(p^n-1)(p^n-2)$.

Proof. For a cyclotomic field of conductor $N = p^n$ one has

$$\log|D_{K/\mathbb{Q}}| = \varphi(N)\log N - \varphi(N)\sum_{\ell \mid N}\frac{\log \ell}{\ell - 1}.$$

Here $\varphi(N) = p^{n-1}(p-1)$, $\log N = n\log p$ and the sum equals $\log p/(p-1)$, giving

$$\log|D| = p^{n-1}(p-1)n\log p - p^{n-1}\log p = p^{n-1}((p-1)n-1)\log p.$$

Division by d gives the display. The formula was checked in exact rational arithmetic for the 25 pairs $p \in \{3,5,7,11,13\}$, $n \in \{1,2,3,12,20\}$; every difference is 0. The last two normalizations are those of Appendices NN and GG respectively. □

Theorem 11.26 (Route P5 reduces to one normalization question)**.** *Parshin's original inequality (9) reads*

$$\omega_{V/B}\cdot\omega_{V/B} \le 3\sum_{v\in B\cup B_\infty}\delta'_v\,\epsilon_v + (2g-2)\log|D_{K/\mathbb{Q}}|.$$

Assume that δ'_v and ϵ_v are aligned with Δ_{Del} after division by $d = [K:\mathbb{Q}]$, and set

$$B_n(p) = \frac{(2g_n-2)\big(n - \frac{1}{p-1}\big)}{g_n(2n+1)}.$$

Then $0 \le B_n(p) < 1$ *for all* $p \ge 3$, $n \ge 1$, $B_n(p) \to 1^-$ *as* $n \to \infty$, *and*

$$1 - B_n(p) = \frac{1 + \frac{2}{p-1}}{2n+1} + O(g_n^{-1}).$$

If the term $(2g-2)\log|D_{K/\mathbb{Q}}|/d$ *is* not *already absorbed in* Δ_{Del}, *Parshin's inequality produces* $C_n = 3 + B_n(p) \to 4$ *and* $\kappa \le (4 + B_n(p))/12 \to 5/12$ *instead of* $1/3$. *Hence route P5 is exactly the question whether* Δ_{Del} *absorbs that term; the alignment of* δ'_v *and the archimedean weights* ϵ_v *is a cited normalization, not a derivation.*

Proof. Divide (9) by d and use Theorem 11.25: the discriminant term becomes $(2g_n - 2)(n - \frac{1}{p-1})\log p$. By Appendix NN, $\Delta_{\mathrm{Del}} = g_n \hat{f}_n = g_n(2n+1)\log p$, so the ratio of the two terms is exactly $B_n(p)$. Write $B_n(p) = (2 - 2/g_n) \cdot \frac{2n - \frac{2}{p-1}}{2n+1}$; the second factor is < 1 because $\frac{2}{p-1} > 0$, and the first is < 1, so $B_n(p) < 1$. Both factors tend to 1, hence $B_n \to 1^-$. Subtracting from 1 and dropping the $O(g_n^{-1})$ correction gives $1 - B_n = \frac{1+\frac{2}{p-1}}{2n+1} + O(g_n^{-1})$. Adding the leading coefficient 3 gives $C_n = 3 + B_n(p)$, and $(1 + C_n)/12 = (4 + B_n(p))/12$, whose limit is $5/12$. Exact rational checks for 32 pairs $p \in \{3,5,7,11\}$, $n \in \{1,2,3,5,8,12,20,50\}$ give, for instance, $B_{12}(3) = 0.92$ with $1 - B_{12}(3) = 2/25$ exactly, and $B_{50}(11) = 0.9881188119$; the corresponding κ bounds are 0.41 and 0.415676568, both below $5/12 = 0.41\overline{6}$. If instead the discriminant term is absorbed in Δ_{Del}, it must not be counted again and the coefficient is 3. Nothing else in Appendix HHH changes. □

Theorem 11.27 (Route P3: the executable alignment and the Belyi layer)**.** *Kühn's Theorem I bounds* ω^2_{Ar} *by three contributions,*

$$(2g-2)\log|\Delta_K|, \qquad [K:\mathbb{Q}]\big(\kappa_1 \log b_{\max} + \kappa_2\big), \qquad \sum_{p \text{ bad}} a_p \log \mathrm{Nm}(\mathfrak{p}),$$

with $a_p \in \mathbb{Q}$ *determined by local intersection numbers on the bad fibres. For the towers of this audit:*

1. *the proved Belyi/Fermat classes (Curilla–Kühn, Javanpeykar, Kühn) are families with prime or squarefree exponent* $p \to \infty$, *whereas the audit towers fix* p *and let* $n \to \infty$; *extending the bounds to* p^n *is a change of tower direction, not a constant improvement;*
2. *the Belyi degree satisfies* $\log b \asymp n \log p$, *hence* $\rho_n = \log b / \hat{f}_n = \frac{n}{2n+1} \to \frac{1}{2}$ *exactly, so this class can reach only the weak layer* $R = O(1)$ *of Problem 10.2, never* $C = 3$;
3. *the local coefficient* a_p *is no longer a black box: Appendix HH gives* $\Delta_{\mathrm{Del,p}} - f_p = \operatorname{rank} L(\Gamma_p)$, *Appendix SS gives* $\operatorname{rank} L = 0$ *on the mixed Kummer cluster trees, and Appendices QQ and RR give* $\mathrm{Sw}_p = f_p$ *exactly.*

Therefore the unique executable step of route P3 is to align a_p *with* $\operatorname{rank} L + \mathrm{Sw}_p$; *this enlarges the proved domain to all wild Kummer towers at the level* $R = O(1)$ *and does not close item (II).*

Proof. Item (1) is quoted from the sources: Curilla–Kühn treat Fermat curves of prime exponent and state that the squarefree-exponent case is deferred, while Javanpeykar bounds ω^2_{Ar} by the minimal Belyi degree. The audit towers of Appendices GG–YY fix p and send $n \to \infty$, so the two families are transverse. For item (2), the standard Belyi map of F_{p^n} has $\log b = n \log p + O(1)$; combining with Theorem 11.25 gives $\rho_n = n/(2n+1)$, checked exactly for $n = 1, 2, 3, 5, 8, 12, 20, 50, 200$ with limit $1/2$. Since $\kappa_1 \log b + \kappa_2$ is independent of g, its contribution per $g\hat{f}$ is $o(1)$; the remaining term is the bad-fibre sum, whose scale is $O(1)$ but whose coefficient is not 3 without further input. This is precisely the separation of Appendix HHH, Theorem 5 (weak layer) and Appendix III, Theorem 1. For item (3), the three cited appendices are exact identities of this volume, so a_p is computable from the dual graph and the Swan conductor. As the output is $R = O(1)$ rather than $R \le 3$, (II) is not closed. Paragraph 9.5 of the source text states verbatim that arithmetic Bogomolov is strictly weaker than BMY and that refuting $\kappa \to \infty$ requires only $R = O(1)$. □

Theorem 11.28 (Why the constant three is rigid)**.** *The constant* 3 *is not a formal consequence of rank-two tools. Yau's equality case in* $c_1^2 \le 3c_2$ *forces the universal cover to be the complex ball, and Jang (Mich. Math. J. 59 (2010)) constructs, for every* $M > 0$*, a smooth proper surface of general type over a field of positive characteristic with smooth Picard scheme and* $c_1(X) > Mc_2(X)$*. Hence transplanting the surface inequality to arithmetic surfaces requires a Kähler–Einstein or ball-quotient type rigidity that has no arithmetic counterpart in the rank-two chain. This is a bibliographic warning only: Jang's varieties are surfaces over fields of characteristic* p*, not arithmetic surfaces, and no inference about the truth of item (II) over number fields is drawn from it.*

Proof. Both assertions are quoted verbatim from the cited sources. The first is the equality statement of the Yau–Miyaoka theorem; the second is the main theorem of Jang's paper, proved via generically ordinary fibrations. The domain of validity of the counterexample is different from that of Parshin's conjecture, so it is recorded as a structural indication that 3 is not formal, not as a counterexample to (II). This is consistent with Theorem 11.24, where the rank-two chain produces 4. The negative verdict of Section 10.4 of the source text remains restricted to the discriminant channel, and the Bergman-kernel gap of Remark 4.6 is unchanged. □

Theorem 11.29 (Final assessment of the three routes)**.** *The three surviving routes of Appendix III reduce as follows.*

Route	*Status*	*Remaining criterion*
P3	*executable, computation only*	*align* a_p *with* $\mathrm{rk}\, L + \mathrm{Sw}_p$*; gives* $R = O(1)$
P5	*normalization criterion*	*is* $(2g-2)\log\lvert D\rvert/d$ *absorbed in* Δ_{Del}*?*
P6	*constant criterion*	*produce* 3 *instead of* 4*; both Soulé obstructions remain*
proved classes	*the only ones reaching* 3	*Dolce–Mercuri, Mayer, Michel–Ullmo (coeff.* 2 *or* 3*)*

Two independent appearances of the constant 4 *occur: the rank-two Bogomolov inequalities (geometric and arithmetic) of Theorem 11.24, and the normalized Parshin coefficient* $C_n = 3 + B_n(p) \to 4$ *of Theorem 11.26. They have different dimensions and the first is only a formal analogy, but both identify* $C = 4$*, equivalently* $\kappa \le 5/12$*, as the natural ceiling of rank-two and relative-discriminant tools. Crossing from* 4 *to* 3 *requires either the Yau-type rigidity of Theorem 11.28 or the already proved modular classes. Item (II) remains open*

Proof. The table collects Theorems 11.24, 11.26, 11.27 and 11.28 together with Appendix III, Theorem 7, and the class table of Section 10.1 of the source text, which states that in every known class $R \to c$ with $c \in \{2, 3\}$ and that BMY is attained asymptotically only in the coefficient-3 families. The two appearances of 4 are computed independently: one is a quoted constant in a Chern-class inequality, the other is the limit of the exact rational sequence $3 + B_n(p)$ of Theorem 11.26. The dictionary $C \mapsto (1 + C)/12$ is exact rational arithmetic of Appendix HHH. Since none of the three routes supplies the coefficient 3, the equivalence of Appendix III, Theorem 7 is untouched. □

Remark 11.30 (Numerical checks and status). Exact rational checks: the cyclotomic identity of Theorem 11.25 (25 pairs, all differences zero); $B_n(p)$ for 32 pairs, all < 1 and monotone, with $1 - B_{12}(3) = 2/25$ and $1 - B_{50}(3) = 2/1001$; $C \mapsto (1 + C)/12$ for $C = 3, 4, 6, 12, 36, 36/7$ giving $1/3, 5/12, 7/12, 13/12, 37/12, 43/84$; the gap $5/12 - 1/3 = 1/12$; $\rho_n = n/(2n + 1)$ for $n = 1, 2, 3, 5, 8, 12, 20, 50, 200$ with limit $1/2$; and the seven anchors $R = 12\kappa - 1$ of Appendix BB. Text-level additions: Soulé (IHÉS, verbatim $4c_2$ and the successive-minima bound), Parshin 1988 (inequality (9)), Yau–Miyaoka 1977 and the equality case (Bogomolov–Miyaoka–Yau survey), Kühn arXiv:0906.2056 (Theorem I and (0.2.2)), Curilla–Kühn arXiv:0906.3891 (prime exponent), Javanpeykar (minimal Belyi degree), and Jang 2010. None is reproved. Three cited C items are added: the alignment of δ'_v with Δ_{Del}, the archimedean weights ϵ_v, and the explicit coefficients κ_1, a_p. All fourteen No-Go reasons, the pointwise conclusions, the status of Theorem A as an external input, and *abc* as open are unchanged.

11.5 Completing the P3 alignment: the weak layer on all wild Kummer towers

We carry out the only executable computational step left by Appendix JJJ, namely the alignment of the bad-fibre coefficient a_p of Kühn's inequality (0.2.2) with the dual-graph data $\mathrm{rk}_{\mathrm{Lap}}\, L + \mathrm{Sw}_p$ computed in Appendices HH, SS, AAA, QQ and RR. We keep the normalized Wilms Noether identity $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}$, the ratios $\kappa = h_F/(g\hat{f})$, $\sigma = \Delta_{\mathrm{Del}}/(g\hat{f})$, $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$, $\varepsilon_{\mathrm{ar}} = \delta_{\mathrm{Fal}}/(g\hat{f})$, and the notation $C_{a,b} : u^{p^n} + v^{q^m} = 1$ with $a = p^n$, $b = q^m$, $p \neq q$, $g = (a-1)(b-1)/2$, $S = \{p, q\}$.

Theorem 11.31 (The alignment identity and a correction to 11.26(3))**.** *For every mixed Kummer tower, the local Deligne discriminant at p splits as*

$$\Delta_{\mathrm{Del,p}}/\log p = f_p + \big(|V(\Gamma_p)| - 1\big) = \mathrm{Sw}_p + \mathrm{rk}_{\mathrm{Lap}}\, L(\Gamma_p),$$

with the exact finite-layer values

$$\mathrm{Sw}_p = (a-1)(b-1)\Big(n + \frac{n}{a-1} - \frac{1}{p-1}\Big), \qquad \mathrm{rk}_{\mathrm{Lap}}\, L(\Gamma_p) = \frac{a-1}{p-1}.$$

Consequently the geometric share of the bad-fibre intersection number is

$$\theta_p := \frac{\mathrm{rk}_{\mathrm{Lap}}\, L}{\mathrm{rk}_{\mathrm{Lap}}\, L + \mathrm{Sw}_p} = \frac{1}{1 + (p-1)(b-1)\big(n + \frac{n}{a-1} - \frac{1}{p-1}\big)} = O(q^{-m}/n).$$

In particular the quantity called "$\mathrm{rank}\, L = 0$" in 11.26(3) is the first Betti number $b_1(\Gamma_p) = 0$; the Laplacian rank is nonzero, and the object to be aligned with a_p is $\mathrm{Sw}_p + \mathrm{rk}_{\mathrm{Lap}}\, L$, not zero. With this correction $\sigma \neq 1$, in agreement with 8.8.

Proof. By (L1) of Appendix AAA, valid without semistability, $\Delta_{\mathrm{Del,p}} - f_p \log p = (m_p - 1)\log p = (|V(\Gamma_p)| - 1)\log p$; by Appendix SS the special fibre is a cluster dendrogram with $|V| = 1 + (a-1)/(p-1)$ and $b_1(\Gamma_p) = 0$, so the Laplacian rank equals $|V| - 1 = (a-1)/(p-1)$. The semistable conductor decomposition $f_p = b_1 + \mathrm{Sw}_p$ of HH therefore gives $\mathrm{Sw}_p = f_p$. Appendix RR gives the exact normalized conductor $\hat{f}_p/\log p = 2(n + n/(a-1) - 1/(p-1))$, and $f_p = g\hat{f}_p/\log p$; substituting $g = (a-1)(b-1)/2$ yields the displayed Sw_p. Dividing the two exact expressions gives $\mathrm{rk}_{\mathrm{Lap}}\, L/\mathrm{Sw}_p = 1/((p-1)(b-1)U_p)$ with $U_p = n + n/(a-1) - 1/(p-1)$, and $\theta_p = 1/(1 + (p-1)(b-1)U_p)$. The identity was checked in exact rational arithmetic for the eight parameter sets $(3,5,1,1)$, $(3,5,4,4)$, $(3,5,8,8)$, $(3,5,12,12)$, $(3,5,20,20)$, $(5,7,6,6)$, $(7,11,10,10)$, $(3,11,50,40)$; every difference is 0. Since $U_p \sim n$ and $b = q^m$, the share is $O(q^{-m}/n)$: wild Swan dominates the local intersection number, while the graph thickness is subleading. This is the quantitative version of the orthogonality of tower directions noted in 11.26. □

Theorem 11.32 (Exact finite-layer $\sigma - 1$ and a two-channel cross-check)**.** *With $U_p = n + n/(a-1) - 1/(p-1)$ one has the exact cancellation*

$$\frac{\theta_p}{1 - \theta_p}\,\hat{f}_p = \frac{2\log p}{(p-1)(b-1)},$$

independent of n, and therefore

$$\sigma_{n,m} - 1 = \frac{\dfrac{2\log p}{(p-1)(b-1)} + \dfrac{2\log q}{(q-1)(a-1)}}{\hat{f}_{\mathrm{tot}}}.$$

This coincides with the independent 8.8 expression $\big[(|V_p| - 1)\log p + (|V_q| - 1)\log q\big]/(g\hat{f}_{\mathrm{tot}})$.

Proof. By 11.31, $\theta_p/(1-\theta_p) = \mathrm{rk}_{\mathrm{Lap}}\, L/\,\mathrm{Sw}_p = 1/((p-1)(b-1)U_p)$, while $\hat{f}_p = 2U_p \log p$; the product is the display. Since $\sigma - 1 = \sum_p (\Delta_{\mathrm{Del,p}} - f_p \log p)/(g\hat{f}_{\mathrm{tot}})$ and $\mathrm{rk}_{\mathrm{Lap}}\, L_p/g = 2/((p-1)(b-1))$, summing the p and q channels gives the second display. The 8.8 route gives $(a-1)\log p/((p-1)g\hat{f})$ for the p channel, which equals $2\log p/((p-1)(b-1)\hat{f})$ because $g = (a-1)(b-1)/2$. Numerically, for $(p,q,n,m) = (3,5,4,4), (3,5,8,8), (3,5,12,12), (3,5,20,20), (3,5,50,40)$ the two routes agree to relative error $\le 2.8 \times 10^{-16}$, with $\sigma - 1 = 5.942008 \times 10^{-4}$, 3.028927×10^{-6}, 2.407236×10^{-8}, 2.168817×10^{-12} and 4.735941×10^{-27}, reproducing the two values quoted in 8.8. □

Theorem 11.33 (The compositum discriminant and $B_{n,m} < 1$)**.** *For $K = \mathbb{Q}(\zeta_{p^n}, \zeta_{q^m})$ with coprime conductors, $d = [K:\mathbb{Q}] = \varphi(p^n)\varphi(q^m)$ and*

$$\log |D_{K/\mathbb{Q}}|/d = \Big(n - \frac{1}{p-1}\Big)\log p + \Big(m - \frac{1}{q-1}\Big)\log q.$$

Define $B_{n,m} = \frac{2g-2}{g} \cdot X/\hat{f}_{\mathrm{tot}}$ with $X = (n - \frac{1}{p-1})\log p + (m - \frac{1}{q-1})\log q$. Then, with $D = \frac{n}{a-1}\log p + \frac{m}{b-1}\log q > 0$,

$$B_{n,m} = \Big(1 - \frac{1}{g}\Big)\frac{X}{X+D} < 1, \qquad B_{n,m} \to 1^-.$$

For a single-prime Fermat tower this reduces to $B_n(p)$ of 11.25.

Proof. For coprime conductors $D_K = D_1^{d_2} D_2^{d_1}$, so $\log|D_K|/d = \log|D_1|/d_1 + \log|D_2|/d_2$; by the cyclotomic formula of 11.25 each summand is $(n - 1/(p-1))\log p$ respectively $(m - 1/(q-1))\log q$. This was verified exactly for five compositum pairs. By 6.24, $\hat{f}_{\mathrm{tot}} = 2(X+D)$: the conductor carries the extra $n/(a-1)$ and $m/(b-1)$ terms. Hence $B_{n,m} = (2 - 2/g)X/(2(X+D)) = (1-1/g)X/(X+D)$, which is strictly below 1 because $X, D, g - 1 > 0$. Five exact-rational surrogate checks (with $\log p, \log q$ replaced by fixed positive rationals) all give $B < 1$; the single-prime specialization $q^m \to 1$ recovers 11.25's $B_n(p)$, whose 24 values are exact Fractions with $B_n < 1$ and $B_n \to 1^-$. □

Theorem 11.34 (Finite-layer Kühn bound; Problem 10.2 closed on the whole class)**.** *After relativizing $\bar{\omega}^2 = \omega_{\mathrm{Ar}}^2/d$ and aligning a_p with $\Delta_{\mathrm{Del,p}}$ via 11.31, Kühn's inequality (0.2.2) reads, for every finite layer,*

$$\frac{\bar{\omega}^2}{g\hat{f}} \le B_{n,m} + \sigma_{n,m} + C_B \rho_{n,m} + O(1/\hat{f}), \qquad \rho_{n,m} = \frac{n\log p + m\log q}{\hat{f}_{\mathrm{tot}}} \to \frac{1}{2},$$

where C_B denotes the Belyi coefficient (Kühn's κ_1). Hence

$$\sigma R \le B + \sigma + C_B\rho, \qquad R \le 1 + \frac{B}{\sigma} + \frac{C_B\rho}{\sigma} \longrightarrow 2 + \frac{C_B}{2} = O(1),$$

and, taking the conservative $\varepsilon_{\mathrm{ar}} \le 0$,

$$\kappa \le \frac{3 + C_B/2}{12} + O(1/\hat{f}).$$

For every finite C_B the ratio κ is bounded: Problem 10.2 ($\kappa \to \infty$) is closed on all wild Kummer and Belyi-type towers with fixed S and $\hat{f} \to \infty$, independently of the numerical value of C_B.

Proof. Divide (0.2.2) by $d\, g\hat{f}$. The number-field discriminant term gives $B_{n,m}$ by 11.33; the Belyi term $\kappa_1 \log b_{\max}$ gives $C_B\rho_{n,m}$ because $\log b_{\max} \asymp n\log p + m\log q$ and $\rho_{n,m} \to 1/2$ (11.26); the bad-fibre term $\sum_p a_p \log \mathrm{Nm}(\mathfrak{p})$, once a_p is identified with $\Delta_{\mathrm{Del,p}}/\log p$ by 11.31, gives $\sigma_{n,m}$. If a normalization makes two of these terms overlap, the displayed sum is a coarsening and the bound only weakens. By Appendix BB, $\bar{\omega}^2/(g\hat{f}) = \sigma R$, which gives the bound on R. Wilms' identity $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$ then gives $12\kappa \le \sigma + B + \sigma + C_B\rho + \varepsilon_{\mathrm{ar}}$; with $\varepsilon_{\mathrm{ar}} \le 0$, $\sigma \to 1$, $B \to 1$, $\rho \to 1/2$ this tends to $3 + C_B/2$. Since C_B is a fixed finite constant of the cited inequality, κ cannot diverge. The values $C_B = 0, 1, 2, 4, 8, 36$ give $\kappa \le 1/4$, $7/24$, $1/3$, $5/12$, $7/12$, $7/4$ respectively. □

Theorem 11.35 (Exact barrier criterion: $C_B \le 2$)**.** *The Kühn channel yields $\kappa \le 1/3$ (equivalently $R \le 3$) if and only if*

$$C_B \le C_B^{\max} = \frac{4 - 2\sigma - B}{\rho} \longrightarrow \frac{4 - 2 - 1}{1/2} = 2.$$

On a single-prime Fermat tower every ingredient is an exact rational:

$$\sigma_n - 1 = \frac{2}{(p-1)(p^n-2)(2n+1)}, \qquad \rho_n = \frac{n}{2n+1}, \qquad B_n(p) = \frac{(2g_n - 2)(n - \frac{1}{p-1})}{g_n(2n+1)}.$$

For all $n \ge 3$ and $p \in \{3, 5, 7, 11\}$ one has $C_B^{\max}(p, n) < 4$; the layer $n = 1$ is a degenerate boundary with $C_B^{\max}(3, 1) = 4$ exactly.

Proof. $\kappa \le 1/3$ is $12\kappa \le 4$; by 11.34, $12\kappa \le 2\sigma + B + C_B\rho$, so the condition is $C_B \le (4 - 2\sigma - B)/\rho$. Taking $\sigma \to 1$, $B \to 1$, $\rho \to 1/2$ gives the limit 2. For the Fermat tower 5.60 gives $\hat{f}_n = (2n+1)\log p$, $g_n = (p^n - 1)(p^n - 2)/2$, and 11.32 with $q^m \to 1$ gives $\sigma_n - 1 = 2/((p-1)(p^n-2)(2n+1))$, an exact rational because $\log p$ cancels. The 24 exact values $(p, n) \in \{3, 5, 7, 11\} \times \{1, 3, 8, 12, 20, 50\}$ give, for instance, $C_B^{\max}(3, 12) = 2.249999686$, $C_B^{\max}(3, 50) = 2.060000000$, $C_B^{\max}(11, 50) = 2.044000000$, each monotone decreasing to 2; for $n = 1$, $B_1(3) = 0$ and $\sigma_1 = 4/3$ give $(4 - 8/3 - 0)/(1/3) = 4$. The mixed-tower values $C_B^{\max} = 1.8463$, 1.9414, 1.9895 for $(3, 5, 4, 4)$, $(3, 5, 12, 12)$, $(3, 11, 50, 40)$ approach 2 from below. □

Theorem 11.36 (Third appearance of the constant four; (II) remains open)**.** *Under the normalization in which the Belyi coefficient equals the constant four already met in P6 (rank-two Bogomolov $4\hat{c}_2$) and P5 ($3 + B_n \to 4$), one has $C_B = 4 > 2 = C_B^{\max}(\infty)$, hence the Kühn channel gives*

$$R \le 4, \qquad \kappa \le \frac{5}{12} = \frac{1}{3} + \frac{1}{12},$$

the gap $1/12$ being the Heisenberg anchor of Appendix BB. This is the third independent convergence on the constant four and does not close (II): the only classes reaching $R \le 3$ remain the proved modular towers of 00544, Section 10.1. The computational part of P3 is therefore complete.

Proof. Insert $C_B = 4$ into 11.34: $R \le 2 + 4/2 = 4$ and $\kappa \le (3 + 2)/12 = 5/12$; the difference to $1/3$ is $5/12 - 1/3 = 1/12$. By 11.35, $R \le 3$ would require $C_B \le 2$, which fails for $C_B = 4$ on every layer with $n \ge 3$. The three occurrences of four are logically independent: 11.24 gives $4\hat{c}_2$ and $4c_2(\Omega^1_X)$ in the rank-two transplantation chain, 11.25 gives $3 + B_n(p) \to 4$ from the cyclotomic discriminant term of Parshin's inequality (9), and 11.34–11.35 give the Belyi coefficient $C_B = 4$ of the Kühn bound. None of them supplies the coefficient 3; the modular classes with coefficient 2 or 3 remain the only proved inputs, so (II) stays open while P3 has no further purely computational step. □

Remark 11.37 (Status of the inputs). Theorems 11.31–11.33 and the exact criteria of 11.35 are unconditional relative to the audit's standard citations (Appendices HH, SS, AAA, QQ, RR, NN, JJJ and the cyclotomic discriminant formula). 11.34 and 11.36 are conditional on the form of Kühn's inequality (0.2.2), on the relativization $\bar{\omega}^2 = \omega^2_{\mathrm{Ar}}/d$, on the identification $C_B = \kappa_1$, and on the normalization $C_B = 4$; these are textual-level citations, not reproved here. The closure of Problem 10.2 is unconditional in C_B: it needs only finiteness of the Belyi coefficient. One textual conditional is removed (the a_p alignment of 11.26); no new No-Go, no new mathematical hypothesis and no new global conditional are introduced.

11.6 P5: the normalization criterion is settled by the source, and (II) is not a normalization issue

We explore route P5 of Appendix HHH, namely Parshin's inequality (9) and the question left open by Appendix JJJ: whether Δ_{Del} absorbs the number-field discriminant term $(2g-2)\log|D_K|$. We keep the Wilms identity $12h_F = \bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}$, the ratios $\kappa = h_F/(g\hat{f})$, $\sigma = \Delta_{\mathrm{Del}}/(g\hat{f})$, $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$, $\varepsilon_{\mathrm{ar}} = \delta_{\mathrm{Fal}}/(g\hat{f})$, and write $\delta_{\mathrm{base}} := (2g-2)\log|D_K|/d$, $B := \delta_{\mathrm{base}}/\Delta_{\mathrm{Del}}$. Throughout $K = \mathbb{Q}(\zeta_{p^n})$ unless stated otherwise, and $c := \lim \hat{f}_n/(n\log p)$ is the conductor density of Theorem 3.4 of the source.

Theorem 11.38 (The P5 criterion is settled verbatim by the source)**.** *Section 2.1 of 00544v2 defines* Δ_{Del} *as the logarithm of the sum of the* local *discriminant exponents over the finite places, and Remark 2.5 gives* $\Delta_{\mathrm{Del,p}} = f_p + m_p - 1$ *place by place, with no semistability, genus or residue-characteristic hypothesis. Section 5.1 of the same source states Parshin's inequality (9) verbatim as*

$$\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + 3\delta_{\mathrm{arch}} + (2g-2)\log|D_K|.$$

Hence $(2g-2)\log|D_K|$ *is a third, independent term and is* not *absorbed into* Δ_{Del}*. Consequently the "absorbed" branch of 11.25 is void and the "not absorbed" branch is the actual content of (9): with* $B_n(p) = \frac{2g_n-2}{g_n}\frac{n-1/(p-1)}{2n+1}$*,*

$$C_n = 3 + B_n(p) \longrightarrow 4, \qquad \kappa \le \frac{4+B_n(p)}{12} \longrightarrow \frac{5}{12}.$$

Moreover Section 5.1 itself performs the alignment of Parshin's original $\sum \delta'_v \epsilon_v$ *with* $3\Delta_{\mathrm{Del}} + 3\delta_{\mathrm{arch}}$*, so the two textual conditionals recorded in 11.25 (*$\delta'_v \leftrightarrow \Delta_{\mathrm{Del}}$ *and the archimedean weight of* ϵ_v*) are removed.*

Proof. The definition of Δ_{Del} in Section 2.1 is a sum over finite places of local exponents δ_p; Remark 2.5 identifies $\delta_p = f_p + m_p - 1$ by Saito's theorem and Liu's identity $-\mathrm{Art} = m - 1 + \varphi$. The number-field discriminant is instead $\log|D_K| = \sum_p v_p(D_K)\log p$, a global invariant of the base field; the two are distinct quantities, and Section 5.1 displays them as separate summands. For the cyclotomic field, $d = p^{n-1}(p-1)$ and $v_p(D_K) = p^{n-1}(n(p-1)-1)$, so $\log|D_K|/d = (n - \frac{1}{p-1})\log p$ exactly (25 exact rational checks, $p \in \{3,5,7,11,13\}$, $n \in \{1,2,3,12,20\}$). With $\hat{f}_n = (2n+1)\log p$ and $g_n = (p^n-1)(p^n-2)/2$ one obtains the stated $B_n(p)$; 35 exact values are strictly increasing and < 1, with $1 - B_n = (1 + \frac{2}{p-1})/(2n+1) + O(1/g_n)$, for instance $B_{12}(3) = 0.92 = 1 - 2/25$ and $B_{50}(11) = 0.9881188119$. Dividing (9) by $\Delta_{\mathrm{Del}} = g\hat{f}\sigma$ gives $R \le 3 + B_n$; the identity $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$ with $\sigma \to 1$, $\varepsilon_{\mathrm{ar}} \le 0$ gives $\kappa \le (4+B_n)/12$. The two textual conditionals are withdrawn because the rewritten form of (9) is quoted explicitly in Section 5.1 of the source, with $3\delta_{\mathrm{arch}}$ carrying the archimedean weight. □

Theorem 11.39 (A unified conductor-density dictionary: $B_\infty = 2/c$)**.** *Let a tower be defined over* $K = \mathbb{Q}(\zeta_{p^n})$ *with* $\hat{f}_n = c\,n\log p + O(\log p)$*,* $c > 0$ *fixed or tending to infinity, and put* $B = \delta_{\mathrm{base}}/\Delta_{\mathrm{Del}}$*. Then*

$$B_\infty = \frac{2}{c}, \qquad C_{\mathrm{Parshin}} = 3 + \frac{2}{c}, \qquad \kappa \le \frac{1}{3} + \frac{1}{6c}, \qquad \kappa = \frac{c+1}{12c},$$

the last equality being Theorem 7.7 of the source for the Heisenberg family. The two anchors of the source are recovered exactly: $c = 2$ *(Theorem 3.4, Fermat/CM conductor density) gives* $\kappa = 1/8$*,* $B = 1$*,* $C_{\mathrm{Parshin}} = 4$*,* $\kappa \le 5/12$*; and the three rows* $c = 6,547, 1.79 \times 10^6$ *of Theorem 7.7 give* $\kappa = 7/72 = 0.097222$*,* 0.083486*,* 0.083333 *with ratios* $1/c$ *equal to* 1.67×10^{-1}*,* 1.83×10^{-3}*,* 5.59×10^{-7}*.*

Proof. By the exact cyclotomic identity above, $\delta_{\mathrm{base}} = (2g-2)(n - \frac{1}{p-1})\log p/g$, while $\Delta_{\mathrm{Del}} = g\hat{f}\sigma$ and $\sigma \to 1$ by Appendix 6.1; hence $B = \frac{2g-2}{g}\frac{(n-\frac{1}{p-1})\log p}{c\,n\log p}(1+o(1)) \to 2/c$. Parshin's inequality gives $R \le 3 + B$, and $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$ gives $12\kappa \le 4 + B = 4 + 2/c$, i.e. $\kappa \le 1/3 + 1/(6c)$. Theorem 7.7 of the source states $\kappa = \frac{c+1}{12c}$ for the Heisenberg tower and records that the discriminant term is of lower order, with main term $\sim c\,n\log p$ and discriminant term $\sim n\log p$; its ratio column is therefore $O(1/c)$, consistent with $B = 2/c$ up to the factor $\frac{2g-2}{g} \to 2$. Exact rational evaluation gives $(c+1)/(12c) = 1/8$ at $c = 2$, $7/72$ at $c = 6$, and $548/6564 = 0.083485679\ldots$ at $c = 547$; the mixed-tower values $B_{n,m}$ of Appendix 11.33 (0.993399878, 0.999999211, 1.000000000) converge to the same limit $2/c = 1$ with $c = 2$. □

Theorem 11.40 (Absorption is not a change of normalization)**.** *Set* $\Delta_{\mathrm{P}} := \Delta_{\mathrm{Del}} + \delta_{\mathrm{base}} = (1+B)\Delta_{\mathrm{Del}}$. *Parshin's inequality (9) is equivalent to*

$$\bar{\omega}^2 \le (3+B)\Delta_{\mathrm{Del}} = \frac{3+B}{1+B}\,\Delta_{\mathrm{P}}, \qquad \frac{3+B}{1+B} \longrightarrow 2 \quad (B \to 1),$$

whereas "coefficient 3 relative to Δ_{P}*" would read* $3\Delta_{\mathrm{P}} = 3(1+B)\Delta_{\mathrm{Del}} \to 6\Delta_{\mathrm{Del}}$*, an inequality strictly stronger than (9) by* $2\delta_{\mathrm{base}}$*. For* $B = 0, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, 1$ *the absorbed coefficient equals* 3, 2.5, 7/3, 2.2, 2*. In particular no choice of normalization produces the coefficient* 3 *when* $B > 0$*; the branch "absorbed* $\Rightarrow 3$*" of 11.25 is withdrawn on two independent grounds.*

Proof. $\delta_{\mathrm{base}} = B\Delta_{\mathrm{Del}}$ by definition, so $3\Delta_{\mathrm{Del}} + \delta_{\mathrm{base}} = (3+B)\Delta_{\mathrm{Del}}$ and $\Delta_{\mathrm{P}} = (1+B)\Delta_{\mathrm{Del}}$; dividing gives the first identity. The coefficient function $(3+B)/(1+B) = 1 + 2/(1+B)$ is strictly decreasing in $B \ge 0$, equals 3 only at $B = 0$, and tends to 2 as $B \to 1$; the five stated values are exact rationals. Replacing the right-hand side by $3\Delta_{\mathrm{P}}$ adds $3\delta_{\mathrm{base}} - \delta_{\mathrm{base}} = 2\delta_{\mathrm{base}}$, so it is not a rewriting of (9) but a different, stronger assertion. Hence the "absorbed" branch of 11.25 fails both factually (11.38) and structurally (here). □

Theorem 11.41 (Disjointness: coefficient three only at the Heisenberg endpoint)**.** *On towers over* $K = \mathbb{Q}(\zeta_{p^n})$ *one has the equivalences*

$$C_{\mathrm{Parshin}} \to 3 \iff B \to 0 \iff c \to \infty \iff \kappa \to \frac{1}{12},$$

and the Parshin upper bound exceeds the actual height–conductor ratio by

$$\kappa_{\mathrm{bound}} - \kappa = \left(\frac{1}{3} + \frac{1}{6c}\right) - \frac{c+1}{12c} = \frac{3c+1}{12c} = \frac{1}{4} + \frac{1}{12c} > \frac{1}{4}.$$

Consequently the discriminant term of (9) has no bearing on the barrier on this whole class of towers: the barrier $\kappa \le 1/3$ *is already unconditional there by the CM–Parseval bound* $\kappa \le 1/8$ *of Appendix ZZ and by the joint limit of Appendix YY.*

Proof. $C_{\mathrm{Parshin}} = 3 + 2/c$ by 11.39, so $C_{\mathrm{Parshin}} \to 3$ iff $c \to \infty$, which by $B = 2/c$ is equivalent to $B \to 0$; Theorem 7.7 gives $\kappa = (c+1)/(12c) \to 1/12$. The displayed difference is an exact rational identity, checked at $c = 2, 6, 547, 1.79 \times 10^6$; at $c = 2$ it is $7/24 = 5/12 - 1/8$. Since $\kappa \le 1/8 < 1/3$ holds on this class by ZZ (the CM–Parseval inequality $\mu \le g\hat{f}/4$), the Parshin bound, whether it carries coefficient 3 or 4, is never the binding constraint. The claim that (9) could supply the missing coefficient 3 therefore concerns only a region where the barrier is already proved by other means. □

Theorem 11.42 (Fixed base field: the modular towers have $B = 0$)**.** *For the modular towers* $X_0(N), X_1(N)$ *over* $K = \mathbb{Q}$ *one has* $|D_K| = 1$*, hence* $\delta_{\mathrm{base}} = 0$ *and* $B = 0$ *exactly; the right-hand side of (9) is* $3\Delta_{\mathrm{Del}} + 3\delta_{\mathrm{arch}}$ *with coefficient* 3 *and no relaxation. Correspondingly Theorem 5.2 of the source gives* $R(X_0(N)) \to 3$ *(finite level* $\approx 3g/(g+1)$*), proved by the automorphic results of Dolce–Mercuri, Michel–Ullmo and Mayer, not by Parshin's inequality. Thus the class on which the barrier is sharp and non-trivial is precisely the class on which the number-field discriminant term is absent.*

Proof. $D_{\mathbb{Q}} = 1$ gives $\log|D_K| = 0$ and therefore $\delta_{\text{base}} = 0$, $B = 0$, $C_{\text{Parshin}} = 3$ and $\kappa \le 1/3$ from 11.39; all these are exact, not asymptotic. Theorem 5.2 of the source computes $\Delta_{\text{Del}} = (g+1)\log p$ and $\omega^2 \sim 3g\log p$ on $X_0(p)$, gives $\sigma = 1 + 1/g$, $\sigma(1+R) = 4 + 1/g$ and $\kappa = 1/3 + 1/(12g) + d/12$ with $d = \delta_{\text{Fal}}/(g\log p) \to 0$ by Jorgenson–Kramer; the convergence $R \to 3$ is one-sided and the constant 3 is attained only asymptotically. Section 5.1 states explicitly that the upper bound in (9) remains a conjecture, so this computation is not a consequence of Parshin's inequality. □

Theorem 11.43 (Route (II) is the general upper bound, not a normalization issue)**.** *The three statements*

$$\bar{\omega}^2 \le 3\Delta_{\text{Del}} + O(g), \qquad h_F \le \frac{1}{3}\Delta_{\text{Del}} + \frac{1}{12}\delta_{\text{Fal}} + O(g), \qquad \limsup R \le 3$$

are equivalent by the Wilms identity and the definitions of $\kappa, \sigma, R, \varepsilon_{\text{ar}}$ together with $\sigma \to 1$; their validity on general towers is general arithmetic BMY, i.e. Problem 10.1 of the source, and remains open. Route P5 is therefore downgraded from a normalization criterion to a dictionary plus the disjointness theorem 11.41: what is missing is an upper bound on $\bar{\omega}^2$, not a convention for $\log|D_K|$. The constant four has its first structural explanation, $4 = 3 + 2/c$ with $c = 2$ the conductor density of Theorem 3.4; the constants $4\hat{c}_2$ of 11.24 and $C_B = 4$ of 11.36 are normalization constants of different dimension and their coincidence in value is not a mathematical convergence.

Proof. The first two statements are related by $12h_F = \bar{\omega}^2 + \Delta_{\text{Del}} + \delta_{\text{Fal}}$ after division by Δ_{Del}; dividing further by $g\hat{f}$ and using $\sigma = \Delta_{\text{Del}}/(g\hat{f}) \to 1$ and $\varepsilon_{\text{ar}} = O(1/\hat{f})$ gives the equivalence with $\limsup R \le 3$, as in 11.8 and 11.21. Section 5.1 of the source states verbatim "the upper bound remains a conjecture" and "one may not cite arithmetic BMY as a theorem", and Problem 10.1 records the general case as open; 11.38–11.42 show that the number-field discriminant term is absent on the sharp modular class and irrelevant on the cyclotomic class. Hence no normalization choice affects (II). Finally $3 + 2/c = 4$ at $c = 2$ by 11.39, whereas $4\hat{c}_2$ is the rank-two Bogomolov constant and C_B is the Belyi coefficient; the three constants have different dimensions, so the earlier phrasing "three independent convergences on four" is tightened to a numerical coincidence with one structural source. No new No-Go and □

Remark 11.44 (Status of the inputs). Theorems 11.38–11.43 are unconditional relative to the audit's standard citations: 00544v2 Sections 2.1, 2.5, 3.4, 5.1, 5.2, 7.7 and Problem 10.1, Parshin's inequality (9) as quoted there, the Wilms identity, and Appendices BB, HHH, III, JJJ, KKK, YY, ZZ, OO of this volume. The external literature is verified at the textual level only; in particular the source itself records that the implication "BMY implies Szpiro" in Parshin's original paper has not been independently verified, and that the upper bound in (9) is a conjecture. The finite-layer constants in $\hat{f}_n = c\,n\log p + O(\log p)$, the values of c in Theorem 7.7, $\varepsilon_{\text{ar}} \le 0$ and $(2g-2)/g \to 2$ are cited, not reproved. Two textual conditionals are removed (δ'_v alignment and the archimedean weight); the remaining ones are the general validity of (9), Problem 10.11, connectivity of the fibre products and the absolute normalization of the mixed CM height. The fourteen No-Gos and all constants of the pointwise side are unchanged, and this appendix does not constitute an independent verification of Theorem A.

11.7 General-tower upper bounds: the ladder, the optimal closure, and the rigidity gap

We place routes P3, P5 and P6 of Appendix HHH on a single ledger. Throughout we use the Wilms–Noether identity $12h_F = \bar{\omega}^2 + \Delta_{\text{Del}} + \delta_{\text{Fal}}$ and divide by $g\hat{f}$:

$$12\kappa = \sigma(1+R) + \varepsilon_{\text{ar}}, \qquad R = \bar{\omega}^2/\Delta_{\text{Del}}, \quad \sigma = \Delta_{\text{Del}}/(g\hat{f}), \quad \varepsilon_{\text{ar}} = \delta_{\text{Fal}}/(g\hat{f}). \tag{MMM.0}$$

If $\bar{\omega}^2 \le C\Delta_{\text{Del}} + O(g)$, $\sigma \to 1$ and $\varepsilon_{\text{ar}} \to 0$, then $\kappa \le (C+1)/12$ and $R \le C$. We write $c = \lim \hat{f}_n/(n\log p) \ge 2$ for the conductor density of Theorem 3.4 of the source, and $B = \delta_{\text{base}}/\Delta_{\text{Del}} = 2/c$ for the number-field discriminant ratio of Appendix LLL.

Theorem 11.45 (The three-step ladder of general-tower upper bounds)**.** *For every $c \geq 2$,*

$$\underbrace{3}_{\text{target (II)}} \quad \leq \quad \underbrace{3+\frac{2}{c}}_{\text{Parshin-type, conditional}} \quad \leq \quad \underbrace{4}_{\text{rank-two Bogomolov, unconditional}}, \tag{11.45}$$

and correspondingly $\frac{1}{3} \leq \frac{1}{3}+\frac{1}{6c} \leq \frac{5}{12}$ for κ and $3 \leq 3+\frac{2}{c} \leq 4$ for R. The middle bound is conditional: Section 5.1 of the source states Parshin's inequality (9) with the explicit warning that "the upper bound remains a conjecture". The unconditional closure is the constant 4 described in Theorem 11.46 and does not depend on c.

Proof. Appendix LLL gives $C_{\mathrm{Parshin}} = 3 + 2/c$ with $B_\infty = 2/c$; since $c \geq 2$, $2/c \leq 1$ and the two inequalities of (11.45) are exact rational inequalities. Dividing by twelve, $(C+1)/12 = (4+2/c)/12 = \frac{1}{3}+\frac{1}{6c}$, and $R = 12\kappa - 1 = 3 + \frac{2}{c}$. At $c = 2$ (Theorem 3.4, Fermat/CM conductor density) the middle and right constants both equal 4; as $c \to \infty$ the middle constant tends to 3, but then $\kappa = (c+1)/(12c) \to 1/12$ by Theorem 7.7 of the source, which is the Heisenberg endpoint where the barrier is already unconditional by the CM–Parseval bound of Appendix ZZ. Ten exact rational checks, $c \in \{2, 3, 4, 6, 9, 12, 20, 100, 547, 1790000\}$, verify $3 \leq C \leq 4$, $\kappa_{\mathrm{b}} = (C+1)/12$, $R_{\mathrm{b}} = 12\kappa_{\mathrm{b}} - 1$ and $\kappa_{\mathrm{b}} - \kappa = (3c+1)/(12c)$; for $c = 2$ the gap is $7/24$, for $c = 6$ it is $19/72$ with $\kappa = 7/72$, and for $c = 547$ one obtains $\kappa = 0.083485679$, matching the table of Theorem 7.7. Conditional status follows from the verbatim literature status of Section 5.1. □

Theorem 11.46 (The proved inputs have optimal closure 5/12, and do not imply (II))**.** *Let $\mathcal{I}$ consist of (I1) the Wilms–Noether identity, (I2) the archimedean upper bound of Theorem 8.7 in the form $\varepsilon_{\mathrm{ar}} \leq C_{\mathrm{ar}}\sigma/\hat{f} \to 0$, (I3) the rank-two arithmetic Bogomolov–Gieseker bound $R \leq 4+o(1)$, (I4) the isogeny convex hull $\kappa = \sum_\rho w_\rho \kappa_\rho$, (I5) the proved subclasses $R \leq 1/2$ for CM-type factors and $R \to 3$ for modular towers, and (I6) $\sigma \to 1$. Then all linear/identity consequences of $\mathcal{I}$ give*

$$\sup_{\mathcal{I}} \kappa = \frac{5}{12} + o(1), \qquad \sup_{\mathcal{I}} R = 4,$$

and the supremum is attained in the limit. In particular $\mathcal{I}$ does not imply $\limsup \kappa \leq 1/3$.

Proof. Fix $\sigma = 1$; this is permitted by (I6) and by the finite-layer excess control of Appendix III. By (MMM.0), $12\kappa = 1 + R + \varepsilon_{\mathrm{ar}} \leq 1 + 4 + o(1)$ using (I2) and (I3), hence $\kappa \leq 5/12 + o(1)$ and $R \leq 4 + o(1)$. The bound is sharp: consider a single non-CM-type isogeny factor (so (I4) is trivial and the CM subclass of (I5) does not apply), with $c = 2$ and

$$\sigma = 1, \quad \bar{\omega}^2 = 4\Delta_{\mathrm{Del}}, \quad \delta_{\mathrm{Fal}} = (-L_2 + \hat{f}^{-1})g, \quad L_2 = 2\log(2\pi^4), \quad h_F = \tfrac{1}{12}(\bar{\omega}^2 + \Delta_{\mathrm{Del}} + \delta_{\mathrm{Fal}}).$$

Then (I1) holds by definition with residual 0; (I3) holds with equality; $\varepsilon_{\mathrm{ar}} = (-L_2 + \hat{f}^{-1})/\hat{f} < 0$ satisfies (I2), while $-L_2 + \hat{f}^{-1} > -L_2$ keeps the Wilms lower bound strict; (I6) holds exactly; and the modular subclass of (I5) is not assumed. Eight values $\hat{f} \in \{10, 20, 50, 100, 200, 500, 1000, 10000\}$ give $12\kappa - (1 + R + \varepsilon_{\mathrm{ar}}) = 0$ exactly, with $\kappa \to 5/12$ and $R \equiv 4$. For small $\hat{f}$ the negative archimedean term puts κ below $1/3$ (for example $\kappa = 0.329632221$ at $\hat{f} = 10$); the violation of (II) occurs only in the limit, which is consistent with the finite-layer excess of Appendix III. Since the six inputs are identities or linear inequalities in $(\kappa, R, \varepsilon_{\mathrm{ar}})$ at $\sigma = 1$, a linear-programming duality argument shows $5/12$ is their exact optimal closure: no nonlinear or rigidity inequality belongs to $\mathcal{I}$. □

Theorem 11.47 (A sufficient condition: CM or GL_2-automorphic factors)**.** *Suppose every simple isogeny factor of the Jacobian of a curve tower is either of CM type or* GL_2 *automorphic (modular-tower class). Then, relative to the standard citations,*

$$R_X \leq (1-w)\cdot\tfrac{1}{2} + w \cdot 3 \leq 3, \qquad \kappa_X \leq \tfrac{1}{3},$$

where w is the conductor mass carried by the GL_2 *factors. More generally, if the remaining factors enjoy only the Bogomolov bound $R_{\mathrm{o}} \le 4$ and carry conductor weight w_{o}, then*

$$R_X \le 3 + w_{\mathrm{o}}, \qquad \kappa_X \le \tfrac{1}{3} + \frac{w_{\mathrm{o}}}{12}.$$

Consequently (II) is equivalent to $R_\rho \le 3$ for non-GL_2 *simple factors, or to $w_{\mathrm{o}} \to 0$.*

Proof. Faltings height, Deligne discriminant and $\bar{\omega}^2$ are additive under isogeny decomposition, so R_X and κ_X are conductor-mass weighted means of the factor values (Appendix EE). Appendix ZZ gives $R \le 1/2$ for CM-type factors by the CM–Parseval bound $\mu \le g\hat{f}/4$, and Theorem 5.2 of the source gives $R \to 3$ on modular towers by the Dolce–Mercuri / Mayer theorem. The first display is the resulting convex combination; six exact weights $w \in \{0, 1/10, 1/4, 1/2, 3/4, 1\}$ give $\kappa_X \in \{1/8, 7/48, 17/96, 11/48, 9/32, 1/3\}$, all $\le 1/3$, with equality only at $w = 1$. If a third class with $R_{\mathrm{o}} \le 4$ is present, the same convexity gives $R_X \le 3 + w_{\mathrm{o}}$ and $\kappa_X \le (4 + w_{\mathrm{o}})/12$; five exact weights $w_{\mathrm{o}} \in \{0, 1/10, 1/4, 1/2, 1\}$ give $1/3$, $41/120$, $17/48$, $3/8$, $5/12$. The endpoint $w_{\mathrm{o}} = 1$ reproduces Theorem 11.46. Existence of non-CM simple factors outside the automorphic class follows from Tsimerman (Annals 176, 2012), but their occurrence inside the Torelli locus is the Coleman–Oort problem and remains conditional; the coarse dimension count of Appendix PP cannot force $w_{\mathrm{o}} \to 0$. □

Theorem 11.48 (P6: the constants 4 and 3 belong to different rigidity layers)**.** *The constant* 4 *is the optimal universal constant of rank-two semistability: Bogomolov's theorem gives $c_1^2 \le 4c_2$ for rank-two semistable bundles, with violation equivalent to an unstable extension, and Moriwaki's arithmetic Bogomolov–Gieseker inequality $\hat{c}_2 - \frac{r-1}{2r}\hat{c}_1^2 \ge 0$ gives the same constant for $r = 2$, with equality if and only if the complex bundle is projectively flat and the metric weakly Einstein. The constant* 3 *is instead the Miyaoka–Yau constant $c_1^2 \le 3c_2$, whose equality case is a ball quotient $\mathbb{B}_2/\Gamma$ with $\Gamma \subset PU(2,1)$ a cocompact lattice. In positive characteristic Jang constructs, for every $M > 0$, surfaces with smooth Picard scheme and $c_1^2 > Mc_2$. Hence* 3 *cannot be obtained by strengthening the stability argument that yields* 4*: it requires an arithmetic Kaehler–Einstein / ball-quotient rigidity theorem.*

Proof. The three statements are verbatim literature checks. Bogomolov's rank-two theorem is exposed by Reid (1977), where $c_1^2 > 4c_2$ is equivalent to a destabilising extension $0 \to L \to E \to I_Z M \to 0$. Moriwaki (Duke Math. J. 74 (1994), 713–761) proves $\hat{c}_2 - \frac{r-1}{2r}\hat{c}_1^2 \ge 0$ for semistable hermitian bundles on arithmetic surfaces and (Amer. J. Math. 117 (1995), 1325–1347) extends it to all dimensions with the projective-flatness / weakly-Einstein equality case; Miyaoka's stable arithmetic version appeared in Soulé (Invent. Math. 116 (1994), 577–599). Miyaoka–Yau gives $c_1^2 \le 3c_2$ with equality exactly for ball quotients, and the orbifold version states the equality case as $\mathcal{X} = \mathbb{B}_2/\Gamma$. Jang (Mich. Math. J. 59 (2010), 169–178) constructs, for any $M > 0$, a surface of general type over a finite field with smooth Picard scheme and $c_1^2 > Mc_2$; Urzúa makes c_1^2/c_2 dense in $[2, \infty)$. Projective flatness is strictly weaker than being a ball quotient, so the equality cases are separated; and the characteristic-p examples show that neither 4 nor 3 is a formal consequence of the Noether identity alone. Consequently the obstacle list of Section 1.3 of the parent paper (no arithmetic analogue of Ω^1_X, and difficulty of controlling the archimedean integral of $\hat{c}_2$) is untouched: P6 asks for a new rigidity theorem, not for a refinement of Bogomolov. The three appearances of 4 in the audit — $3+2/c$ at $c = 2$ (Appendix LLL), $4\hat{c}_2$ (P6) and $C_B = 4$ (Appendix KKK) — therefore share the rank-two threshold but have different sources and dimensions; they are neither a pure numerical coincidence nor a single quantity. □

Theorem 11.49 (A three-scale dictionary for the gap)**.** *With the exact line $R = \frac{3}{2}\mu_{\mathrm{ratio}} - 1$ and $\kappa = \mu_{\mathrm{ratio}}/8$ of Appendix BBB, the four landmarks of the audit are*

$$(\mu_{\mathrm{ratio}}, R, \kappa) \in \{(2/3, 0, 1/12),\ (1, 1/2, 1/8),\ (8/3, 3, 1/3),\ (10/3, 4, 5/12)\},$$

corresponding respectively to the Heisenberg endpoint, the CM–Parseval bound, the barrier (II), and the Bogomolov / optimal closure of Theorem 11.46. The gap is exactly $\Delta\mu = 2/3 = 8\Delta\kappa$, $\Delta R = 1$, $\Delta\kappa = 1/12$.

Proof. Substitute the four rational values into $R = \frac{3}{2}\mu - 1$ and $\kappa = \mu/8$: $\mu = 2/3$ gives $R = 0$, $\kappa = 1/12$; $\mu = 1$ gives $R = 1/2$, $\kappa = 1/8$; $\mu = 8/3$ gives $R = 3$, $\kappa = 1/3$; $\mu = 10/3$ gives $R = 4$, $\kappa = 5/12$. All four identities are exact in $\mathbb{Q}$, and $12\kappa - 1 = R$ holds in each case. The differences are $10/3 - 8/3 = 2/3$, $4 - 3 = 1$ and $5/12 - 1/3 = 1/12$, with $8 \cdot (1/12) = 2/3$. The value $1/12$ is precisely the Heisenberg anchor, so the gap is not an estimation remainder: it is the distance between the semistability threshold and the rigidity threshold of Theorem 11.48. Route (II) is therefore equivalent to lowering the universal bound on μ_{ratio} from $10/3$ to $8/3$, or the constant C from 4 to 3. □

Theorem 11.50 (Terminal state of the breakthrough matrix)**.** *Routes P1 and P2 are closed (Appendix III); P3 gives the weak layer* $R = O(1)$ *and closes Problem 10.2 on the whole wild Kummer / Belyi class (Appendix KKK) but not the constant* 3*; P5 is a normalization dictionary and is irrelevant to (II) (Appendix LLL), with its upper-bound function conditional on Parshin's conjectural inequality; P6 is the rigidity gap of Theorem 11.48. Together with Theorem 11.47, the only remaining inputs are (a) an arithmetic ball-quotient rigidity inequality and (b)* $R_\rho \le 3$ *for non-*GL_2 *simple isogeny factors inside the Torelli locus. All finite computation, combinatorial reasoning and textual alignment on the tower side are exhausted.*

Proof. Each clause is a restatement of the cited appendix together with the present one: P1 was relocated to the already-closed archimedean channel, P2 was withdrawn after full-text verification that Yuan's Theorem 1.4 is one-sided, P3 produced $R \le 1 + B/\sigma + C_B\rho/\sigma$ with $C_B = 4$, P5 proved the irrelevance of the discriminant normalization, and P6 is settled by Theorem 11.48. Theorem 11.46 shows no further linear consequence of the proved inputs can reach 3, and Theorem 11.47 localises the missing inequality to non-GL_2 factors. Both surviving inputs require new analytic or rigidity information rather than further finite computation. □

Remark 11.51 (Corrections and honest boundary). Three phrasings are sharpened. (i) The bound $\kappa \le \frac{1}{3} + \frac{1}{6c}$ of Appendix LLL is conditional on Parshin's conjectural inequality (9); the unconditional closure is $5/12$, independent of c. (ii) The three occurrences of the constant 4 share the rank-two semistability threshold but have distinct sources; the wording "numerical coincidence" in Appendix LLL is replaced by "same threshold, different sources". (iii) The finite-layer excess $R \le 3 + (C_{\mathrm{ar}} + W)/\hat{f}$ of Appendices HHH and III is conditional on the proved subclasses; on a general tower the unconditional excess is $R \le 4 + \varepsilon$, as in Theorem 11.46. All statements are U relative to the Wilms–Noether identity, the source text, Parshin (9), Bogomolov, Moriwaki, Miyaoka–Yau, Jang, Urzúa, Tsimerman and Appendices BB–LLL; the external literature is checked verbatim only and no proof is independently re-verified. The non-derivability of Theorem 11.46 is a statement about the assignment model and the linear inequality closure of $\mathcal{I}$, not a model-theoretic claim about every theorem of arithmetic geometry. The Coleman–Oort conditional is merged into the existing Problem 10.12 item All fourteen No-Go reasons and every pointwise conclusion of the parent paper are unchanged, *abc* remains O, and Theorem A is not independently verified.

11.8 Torelli non-GL2 factors: three existence layers and a literature ledger

Section 11.7 localises route (II) to simple isogeny factors that are neither of CM type nor GL_2 automorphic. This appendix audits that class against eleven external sources, four of them checked verbatim by full-text retrieval, and splits it into three existence problems of different strength. Throughout we keep the Wilms–Noether normalisation of Section 11.7 and write w_{o} for the conductor mass carried by non-GL_2 factors:

$$\kappa \le \frac{1}{3} + \frac{w_{\mathrm{o}}}{12}, \qquad R \le 3 + w_{\mathrm{o}}, \qquad R = 12\kappa - 1. \tag{NNN.0}$$

Theorem 11.52 (Three existence layers)**.** *The phrase "non-CM, non-GL_2 simple isogeny factor of a Jacobian" conflates three problems. (L1) Existence of an abelian variety isogenous to no Jacobian is a theorem for $g > 3$ (Tsimerman), non-constructive and outside the Torelli locus. (L2) Prescribing a given abelian variety as an isogeny factor of some Jacobian is Matsusaka's theorem; a* 2026 *preprint refines it by making the complementary factor absolutely simple of arbitrarily large dimension, but is not peer reviewed. (L3) Exhibiting one explicit curve with a proved non-CM, non-GL_2 factor, together with bad primes, semistable model, conductor and Faltings height under one normalisation, is open. Even if (L2) is granted,* (NNN.0) *bounds κ by $\frac{1}{3} + \frac{w_o}{12}$, so (II) is not settled by existence alone.*

Proof. The three statements are disjoint by quantifier strength: (L1) asserts $\exists A\ \forall C$, (L2) asserts $\forall A\ \exists C$ up to isogeny, (L3) asks in addition for explicit equations and arithmetic invariants. Six exact rational weights $w_o \in \{0, 1/10, 1/4, 1/2, 3/4, 1\}$ give $\kappa \in \{1/3, 41/120, 17/48, 3/8, 19/48, 5/12\}$ and $R = 3 + w_o \in \{3, 31/10, 13/4, 7/2, 15/4, 4\}$, with $12\kappa - 1 = R$ identically; the endpoint $w_o = 1$ reproduces the saturating model of Theorem 11.46. The verbatim literature status of (L1) and (L2) is recorded in Theorems 11.53 and 11.55; no construction in the literature supplies (L3). □

Theorem 11.53 (The Tsimerman label is corrected)**.** *The Annals paper is entitled "The existence of an abelian variety over $\overline{\mathbf{Q}}$ isogenous to no Jacobian", and its abstract reads "We prove the existence of an abelian variety A of dimension g over $\mathbf{Q}$ that is not isogenous to any Jacobian, subject to the necessary condition $g > 3$." Consequently it produces objects outside the Torelli locus, not samples of the dangerous factor inside it, and it does not assert that those objects are non-automorphic. The file's phrasing "non-automorphic simple abelian varieties exist, the obstruction being their Torelli realisation" is sharpened to "abelian varieties isogenous to no Jacobian exist; their automorphy is not decided by this theorem; the obstruction remains Torelli realisation".*

Proof. The title and abstract are verbatim retrieval results; the weaker reading "not isomorphic to any Jacobian" in the delegated survey is retracted. The genus threshold is exact: for $g = 1, 2, 3$ the Torelli locus is dense in $\mathcal{A}_g$ and the theorem does not apply, while for every $g \geq 4$ it supplies a non-constructive example; the nine-value threshold table $g = 1, \ldots, 9$ is checked in the computation file. The logical separation is strict: "isogenous to no Jacobian" neither implies nor is implied by "no automorphic parameterisation", since the latter concerns L-functions and the former the Torelli morphism. Hence the corrected label removes one over-claim without weakening the conclusion of Section 9.1 that the Torelli realisation is the bottleneck. □

Theorem 11.54 (Coleman–Oort status and a downgrade of Appendix BBB)**.** *The Coleman–Oort conjecture is false for $g \leq 7$ and is restated with the hypothesis $g \geq 8$; the Moonen–Oort reduction says that an infinite family of CM Jacobians for $g \geq 8$ forces a positive-dimensional special subvariety $S \subset \mathcal{T}_g$ meeting the open Torelli locus with $\dim S = 1$, or $\dim S = 2$ and S complete, or S a complete ball quotient. Known $g \leq 7$ counterexamples come from abelian or Galois covers of $\mathbb{P}^1$ with fixed group and local monodromy, but there is no complete table of $\overline{\mathbf{Q}}$-simple factors and End^0 for all such subvarieties. The claim of Section 9.1 that those examples are all of CM or GL_2 type is downgraded to an empirical induction from the known lists.*

Proof. The verbatim source states "The conjecture is known to be false for $g \leq 7$, so a corrected version of the conjecture should have as an assumption that $g \geq 8$"; the three cases are the Moonen–Oort main theorem resting on Hain and de Jong–Zhang. A 2022 computation finds no new families for $g \leq 100$, but this is a large case exclusion, not a theorem for arbitrary g. "All known constructions are Galois covers" does not entail "all factors are CM or GL_2": a group-algebra component can have a large endomorphism algebra while its $\overline{\mathbf{Q}}$-simple constituents have unclassified Albert type. Consistent with this, generic cyclic-cover Prym components satisfy $\mathrm{End}_{\mu_d}(P_d) \simeq \mathbf{Z}[\zeta_d]$, so Galois symmetry tends to produce many endomorphisms rather than non-GL_2 factors; the Kummer families of Appendices EEE–FFF exhibit exactly that behaviour. The

CM–Parseval bound of Section 8.1 therefore applies only to factors already verified to be of CM type. □

Theorem 11.55 (Non-GL2 is not non-automorphic: the genuinely GL4 test)**.** *Ribet's abstract states, conditionally on Serre's conjecture, that an abelian variety over* $\mathbf{Q}$ *is of* $\mathrm{GL}(2)$*-type if and only if it is a simple factor over* $\mathbf{Q}$ *of* $J_1(N)$*; the conjecture is now a theorem of Khare–Wintenberger, so the characterisation is unconditional. Fite–Florit–Guitart define* A *to be of* GL_n*-type if* $\mathrm{End}^0(A)$ *contains a number field of degree* $2\dim(A)/n$*, and genuinely so if* $A_{\overline{\mathbf{Q}}}$ *has no* GL_m*-type factor with* $m < n$*; they construct a family of genus-two curves over a quadratic field whose restriction of scalars is a genuinely* GL_4*-type fourfold. Hence "non-*GL_2*" is nonempty in the abelian-variety layer, does not exclude higher-rank automorphy, and yields neither* $R \le 3$ *nor unbounded height growth.*

Proof. The dimension–degree arithmetic is exact. For the fourfold, $\dim A = 4$ and $[\mathrm{End}^0(A) : \mathbf{Q}] = 2$, so $2\dim(A)/n = 8/n = 2$ gives $n = 4$, while the $\mathrm{GL}(2)$ condition requires a field of degree $2\dim(A)/2 = 4 \neq 2$; thus the example is not of $\mathrm{GL}(2)$-type. At $n = 2$ the Fite–Florit–Guitart definition reduces to $2\dim(A)/2 = \dim(A)$, which is Ribet's condition that $\mathrm{End}^0(A)$ contain a field of degree $\dim(A)$; the two definitions agree verbatim. Genuine GL_4-type excludes all $m < 4$. Since the example is a restriction of scalars and no single curve is produced whose $\overline{\mathbf{Q}}$-Jacobian contains it as a simple factor, it belongs to layer (L1)–(L2), not (L3) of Theorem 11.52. The automorphy statement is one-directional: failing the $\mathrm{GL}(2)$-type criterion removes the weight-two newform parametrisation, not the possibility of a higher-rank automorphic representation. □

Theorem 11.56 (Height–conductor literature boundary)**.** *Lu's Theorem 1.1 bounds the local base-change conductor* $c(A, \mathfrak{p})$ *uniformly and, with* $h_{\mathrm{Fal}}(A_L) - h_{\mathrm{Fal}}(A) = -[k : \mathbf{Q}]^{-1} \sum_{\mathfrak{p} \in P} c(A, \mathfrak{p}) \log N\mathfrak{p}$*, reduces the generalised Szpiro conjecture to semistable reduction. It does not supply a linear upper bound for* $h_F / \log \mathrm{cond}$*. No verbatim source found in this audit proves* $R = O(1)$ *or* $R \le 3$ *for a single non-CM, non-*GL_2 *factor. Consequently the saturating model of Theorem 11.46 is consistent with the literature: a single such factor with* $R \equiv 4$ *and* $\kappa \to 5/12$ *is not excluded by any proved input.*

Proof. The eight-point saturation table is reproduced exactly: with $L_2 = 2\log(2\pi^4) = 10.544133448$, $\varepsilon_{\mathrm{ar}} = (-L_2 + \hat{f}^{-1})/\hat{f}$ and $\kappa = (5 + \varepsilon_{\mathrm{ar}})/12$, the values $\hat{f} \in \{10, 20, 50, 100, 200, 500, 1000, 10000\}$ give $\kappa = 0.329632221$, 0.372941111, 0.399126444, 0.407888222, 0.412275361, 0.414909644, 0.415788072, 0.416578800, converging to $5/12 = 0.416666667$, with $12\kappa - (1 + R + \varepsilon_{\mathrm{ar}}) = 0$ to machine precision. This is the assignment model of Section 11.7, now checked against the literature status: Theorem 11.46 is a statement about the linear closure of proved inputs, and the absence of a proved factor-level height bound is an independent, literature-level reason why no linear consequence can reach 3. The closure $\kappa \le 5/12$ is internal to this file (relative to standard citations); the external literature audit did not independently verify its constant history, so it must not be quoted as an externally proved theorem. □

Theorem 11.57 (Net verdict and a three-tier grading)**.** *Route (II) remains open. The dangerous class is resolved into: five items decidable at literature level (Tsimerman's theorem and* $g > 3$*; the* $g \ge 8$ *Coleman–Oort restatement; the Ribet plus Khare–Wintenberger sufficiency boundary; Lu's semistable reduction; the genuinely* GL_n *definition); three items pushable by computation but not across a theorem threshold (a single-object database of* $\overline{\mathbf{Q}}$*-decomposition,* End^0*, bad primes, semistable model and conductor under one normalisation; an explicit descent of the genuinely* GL_4 *example; an explicit computation of* w_{o} *feeding (NNN.0)); and three items requiring new analytic input (a universal height–conductor bound for non-CM, non-*GL_2 *factors; a proof or counterexample for* $\kappa \le 1/3$ *on general towers; upgrading "non-*GL_2*" to "non-automorphic of every rank").*

Proof. The grading is a partition by what is needed to close each item: the first tier needs only verbatim source checking, the second needs finite computation with existing formulas, the third needs a theorem not present in the audited corpus. Counts $5, 3, 3$ and the eleven-source ledger (4 full-text retrievals, 7 verbatim-quoted secondary items) are checked in the computation file. Theorem 11.52 gives the three layers, Theorem 11.53 fixes the Tsimerman label, Theorem 11.54 downgrades one file claim, Theorem 11.55 separates "non-GL_2" from "non-automorphic", and Theorem 11.56 shows the remaining factor-level inequality is not in the literature. Together they leave the sufficient condition of Theorem 11.47 unchanged and sharpen only the description of the residual gap. □

Remark 11.58 (Corrections and honest boundary). Four phrasings are tightened. (i) The delegated survey's weak reading "not isomorphic to any Jacobian" is retracted in favour of the verbatim "not isogenous to any Jacobian" with the condition $g > 3$. (ii) Section 9.1's claim that all $g \leq 7$ counterexamples are of CM or GL_2 type is downgraded to an empirical induction. (iii) The 2026 preprint on Jacobians with prescribed factors is recorded at preprint status only: its abstract confirms the Matsusaka refinement and an absolutely simple complementary factor of arbitrarily large dimension, but the full "prescribed A" statement, the author list and the reported AI-assistance disclaimer were not verified verbatim on the abstract page. (iv) The bound $\kappa \leq 5/12$ is an internal linear closure of this file, not an externally verified theorem. All eleven external sources are checked verbatim or at text level only; no proof is independently re-verified. Three new text-level or preprint items are added to the conditional ledger All fourteen No-Go reasons and every pointwise conclusion are unchanged, *abc* remains O, and Theorem A is not independently verified.

11.9 Verbatim verification of the prescribed-factor preprint and its effect on route (II)

Section 11.8 recorded the 2026 preprint on Jacobians with prescribed factors at abstract level only, leaving the author, the exact "prescribed A" statement and the AI-assistance disclaimer unverified. The full text is now available. This appendix audits all three items verbatim, promotes layer (L2) to theorem status relative to that preprint, and shows precisely which parts of route (II) remain untouched. We keep the normalisation of Section 11.7 and write w_{o} for the conductor mass carried by *all* factors that are not proved to lie in the CM or GL_2 class:

$$\kappa \leq \frac{1}{3} + \frac{w_{\mathrm{o}}}{12}, \qquad R \leq 3 + w_{\mathrm{o}}, \qquad R = 12\kappa - 1. \tag{OOO.0}$$

Theorem 11.59 (Verbatim verification and closure of ledger item F3)**.** *The full text of arXiv:2606.04344v1 [math.AG], dated 3 June 2026, is by Congling Qiu and is entitled* A Note on Jacobians with Prescribed Factors*. Its Theorem 1 reads, verbatim: for k an infinite field and A/k an abelian variety of dimension $g \geq 2$, for every integer $N > 0$ there exists a smooth projective curve C/k contained in A such that A is an isogeny factor of J_C and the complementary isogeny factor is absolutely simple of dimension $> N$. The Acknowledgment states, verbatim, that ChatGPT and Gemini were used as writing assistants, that the author checked and revised all such output, and that the author takes full responsibility for the mathematical content. Consequently the three text-level reservations of Section 11.8 (missing author, unverified prescribed-A statement, unverified AI disclaimer) are closed: ledger item 8 becomes verbatim full-text verified, with the proof itself not independently re-proved.*

Proof. Direct full-text retrieval. The string "CONGLING QIU" occurs five times, "arXiv:2606.04344v1 [math.AG] 3 Jun 2026" once, the phrase "Theorem 1. Let k be an infinite field" once, and "used ChatGPT and Gemini" once. The Corollary 1 statement is also verbatim: it answers a question of Poonen learnt from W. Zhang, asserting that if A and B have no common simple isogeny factor over k then some C/k has A as an isogeny factor of J_C while

J_C shares no simple isogeny factor with B; the case $\dim A = 1$ is handled by applying Theorem 1 to A^2. Verbatim verification establishes the text, not the correctness of the monodromy, torsion-freeness, equidistribution or irreducibility steps used in its proof; that proof remains cited at level C. □

Theorem 11.60 (Layer (L2) upgraded; number-field specialisation)**.** *Let k be an infinite field and A/k an abelian variety with $\dim A \geq 2$. For every $N > 0$ there is a smooth projective curve C/k with $J_C \sim A \times A_{\mathrm{van}}$ up to isogeny, where A_{van} is absolutely simple and $\dim A_{\mathrm{van}} > N$. If k is a number field then infinitely many such curves are defined over k; the case $\dim A = 1$ follows from Corollary 1 applied to A^2. Hence layer (L2) of Section 11.8 is a theorem relative to the cited preprint: an arbitrary prescribed non-GL_2 simple factor can be placed inside the Jacobian of a curve over a number field, and Tsimerman's "isogenous to no Jacobian" varieties do not obstruct that placement.*

Proof. The preprint takes $X = A$, chooses a projective embedding and a smooth complete-intersection flag $A \supset S_1 \supset \cdots \supset S_{g-2} = S$ by Bertini, and for the generic curve of a Lefschetz pencil obtains $J_{C_\eta} \sim A_{k(\eta)} \times A_{\mathrm{van}}$ with $H^1(A_{\mathrm{van}}, \mathbf{Q}_\ell) \simeq H^1_{\mathrm{van}}(C_{\bar{\eta}}, \mathbf{Q}_\ell)$. Choosing the hypersurface degree large forces $\dim A_{\mathrm{van}} > N$, a condition preserved under specialisation. A number field is not algebraic over a finite field, hence is Hilbertian after reduction to a finitely generated subfield; Proposition 2 with Serre's Hilbert irreducibility theorem gives infinitely many $t \in U(k)$ for which $A_{\mathrm{van},t}$ stays absolutely simple. The finite-field-algebraic case is covered by Proposition 3 (Chavdarov plus Lang–Weil). Tsimerman's theorem concerns abelian varieties with $\forall C$ no isogeny to J_C, a different quantifier, so it cannot block an $\exists C$ statement for a prescribed A. □

Theorem 11.61 (The complementary factor is itself a third-class factor)**.** *Let A_{van} be as in Theorem 11.60. Then $\mathrm{End}^0(A_{\mathrm{van},\bar{\eta}}) = \mathbf{Q}$, and for every $d = \dim A_{\mathrm{van}} \geq 2$ this excludes both CM type and GL_2 type. Consequently both isogeny summands of J_C may belong to the unverified class, and the sufficient condition of Theorem 11.47 ("all remaining factors are CM or GL_2") cannot be applied to these curves by default.*

Proof. Proposition 1 of the preprint shows that $D = \mathrm{End}^0(A_{\mathrm{van},\bar{\eta}})$ has $D \otimes \mathbf{Q}_\ell$ inside the centraliser of the symplectic group $\mathrm{Sp}(H^1_{\mathrm{van}})$; by Schur's lemma that centraliser is $\mathbf{Q}_\ell$, so $D = \mathbf{Q}$. For an absolutely simple CM abelian variety of dimension d, the endomorphism algebra contains a commutative semisimple algebra of degree $2d$, hence $[\mathrm{End}^0 : \mathbf{Q}] \geq 2d \geq 4$. For a $\mathbf{Q}$-simple factor of GL_2 type, the Ribet characterisation and the FFG definition give a number field of degree $[\mathrm{End}^0 : \mathbf{Q}] = d \geq 2$. Degree 1 violates both bounds for every $d \geq 2$, verified exactly for $d = 1, \ldots, 8$. Since $\dim A_{\mathrm{van}} > N$ with N arbitrary, such factors exist in every sufficiently large dimension. Choosing A itself in the unverified class makes both summands unverified. □

Theorem 11.62 (Convex-hull bound unchanged; a geometric candidate for the saturating model)**.** *With w_{o} redefined as the total conductor-mass weight of all unverified factors, (OOO.0) still holds, and $w_{\mathrm{o}} = 1$ gives exactly $\kappa = 5/12$, $R = 4$, the saturating endpoint of Theorem 11.46. Corollary 1 of the preprint allows one to avoid a prescribed B, hence to push w_{o} towards 1. Therefore the saturating assignment model of Section 11.7 acquires a geometric existence candidate, while the numerical bound and the openness of route (II) are unchanged.*

Proof. Isogeny additivity of Faltings height and conductor gives $\kappa(J_C)$ as a convex combination of the factor ratios with nonnegative weights summing to one; the proved factor ratios are at most $1/3$, so the excess over $1/3$ is at most $w_{\mathrm{o}}/12$ and $R \leq 3 + w_{\mathrm{o}}$. Six exact rational weights $w_{\mathrm{o}} = 0, 1/10, 1/4, 1/2, 3/4, 1$ give $\kappa = 1/3, 41/120, 17/48, 3/8, 19/48, 5/12$ and $R = 3, 31/10, 13/4, 7/2, 15/4, 4$, with $12\kappa - 1 = R$ in each case. Corollary 1 supplies curves whose Jacobians avoid any given simple isogeny class B, so the proved-class mass can be removed. No arithmetic invariant of the surviving factor is supplied, so the bound is not improved. □

Theorem 11.63 (No arithmetic invariants; layer (L3) and (I' b) remain open; no tower)**.** *The full text of the preprint contains zero occurrences of "Faltings", "conductor", "height", "Δ", "Swan", "bad reduction" and "BMY". Its method is monodromy of complete intersections, Hilbert irreducibility and Chavdarov equidistribution, all existential. Hence none of* κ_A, R_A, w_{o}, Δ_{Del} *or* δ_{Fal} *can be computed from it; layer (L3) stays open and the weight-tunable single-curve realisation (I' b) is not achieved. Moreover the infinitely many specialisations* $t \in U(k)$ *carry no control of* $g_n \to \infty$*, a fixed bad set* S *or* $\hat{f}_n \to \infty$*, so they do not form a tower in the sense of axioms (T1)–(T3).*

Proof. A case-insensitive census of the 401-line full text gives counts $0, 0, 0, 0, 0, 0, 0$ for the seven invariant keywords above; the single occurrence of "arithmetic" is the adjective in the finite-field discussion of the isogeny category. The proofs of Lemma 1, Lemma 2, Propositions 1–3 and Theorem 1 use Deligne's big monodromy, Gabber's torsion-freeness, Katz–Deligne–Hall mod-ℓ monodromy, Chavdarov's irreducibility theorem, Lang–Weil and Serre's Hilbert irreducibility, none of which produces heights or conductors. Hilbertian specialisation yields a set of k-points of a fixed pencil, not a sequence of curves with controlled genus growth and controlled bad primes; hence the tower axioms fail and the tower-side theorems of Sections 6.1, ZZ and AAA cannot be invoked. □

Theorem 11.64 (Three computable interfaces to Sections 3–4 and Remark 1)**.** *(i) Section 3 defines* $\mathcal{Z}_A$ *as the union of Hecke translates of* $A_0 \times \mathcal{A}_{g-a}$ *and states verbatim that it is special when* A *is CM; hence for non-CM* A *the* A*-special locus is weakly special but not special. (ii) Lemma 3 (from Corollary E of Lam–Moretti– Passeri) gives* $\dim V \le 2g - 4a + 2$ *whenever* $a > 1$ *and* $g > 2a - 2$*, while Lemma 4 gives* $\dim V = 2g - 3$ *for* $a = 1$*; Proposition 5 follows for* $g \le 14$ *because* $b_1^2 - 7b_1 + 2 < 0$ *holds exactly for* $1 \le b_1 \le 6$ *and* $b_1 \le \lfloor (g-1)/2 \rfloor \le 6$*. (iii) Remark 1 links the* $\dim A = 1$ *case to ramified double covers and Prym varieties, with Biswas–Paranjape proving* $\mathrm{NS} = \mathbf{Z}$ *for a general Prym; Section 4 defines* $G_{\mathrm{Jac}}(k)$ *and proves that along an* A*-special family all complementary factors have the same class there.*

Proof. (i) is a direct quote: "Each such Hecke translate is a weakly special subvariety of $\mathcal{A}_g$. It is special when A is CM." For (ii), $2g - 4a + 2 < 2g - 3$ reduces to $4a > 5$, true for every integer $a \ge 2$; twenty exact pairs (a, g) with $2 \le a \le 6$ and $g \in \{2a - 1, 2a, 2a + 2, 14\}$ verify the strict inequality. The polynomial $b_1^2 - 7b_1 + 2$ takes values $-4, -8, -10, -10, -8, -4$ at $b_1 = 1, \ldots, 6$ and $+2$ at $b_1 = 7$; $g \le 14$ gives $b_1 \le 6$, so every splitting locus has dimension strictly below $\dim V$. For (iii), the three statements are verbatim in Remark 1, Proposition 6 and Lemma 5. The Prym route connects to Sections 4.1 and DDD, and the preprint itself records that the required big monodromy on anti-invariant cohomology is not in the literature; the Grothendieck-group statements carry no height or conductor data and therefore do not feed (OOO.0). □

Remark 11.65 (Corrections and honest boundary). Section 11.8's ledger item 8 and its phrasing "preprint candidate, author unverified, AI disclaimer unverified" are withdrawn in favour of Theorem 11.59; layer (L2) of Theorem 11.52 is upgraded by Theorem 11.60. One new tightening is recorded: by Theorem 11.61 the complementary factor of the Qiu curve is itself non-CM and non-GL_2, so the sufficient condition of Theorem 11.47 must not be applied to these curves with the remaining factors assumed proved; using w_{o} for all unverified factors restores the same bound. Theorem 11.52's layer (L3), the upper-half problem (I' b), the fourteen No-Go reasons and every pointwise conclusion are unchanged; route (II) remains open. The verification is verbatim text retrieval plus exact rational arithmetic; the preprint is dated 3 June 2026, carries an AI writing-assistance disclaimer with the author taking full responsibility, and its proof is not independently re-verified. All external sources are cited at text level only.

11.10 Final reduction of the remaining open problem: dilution, limited decoupling, and the positive-weight gap

Section 11.9 promoted the prescribed-factor theorem of Qiu (arXiv:2606.04344v1) to verbatim status and suggested that the saturation model of Section 11.7 had acquired a geometric existence candidate. Here we put that candidate back into the isogeny convex hull and into the tower axioms. We use the normalisation of Section 11.7,

$$\kappa = \frac{h_F}{\log \mathrm{cond}}, \qquad R = \frac{\bar{\omega}^2}{\Delta_{\mathrm{Del}}}, \qquad 12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}, \qquad R = 12\kappa - 1, \tag{PPP.0}$$

and write $L_\rho = \log \mathrm{cond}(\rho) = g_\rho \hat{f}_\rho$ for the conductor mass of a simple isogeny factor. Route (II) is the assertion $\limsup \kappa \le 1/3$, equivalently $\limsup R \le 3$, for every curve tower.

Theorem 11.66 (Exact weight law of the insertion operator)**.** *Let $J_C \sim A \times B$ up to isogeny, with A the prescribed factor and B the complementary factor. Put $\lambda = L_A/L_B = g_A\hat{f}_A/(g_B\hat{f}_B)$. Then*

$$\kappa(J_C) = \frac{\lambda\,\kappa_A + \kappa_B}{1+\lambda}.$$

In particular λ depends on the two conductor masses only; the genus g_B enters only through L_B. The law is consistent with the three-dial dictionary $\mu_{\mathrm{ratio}} = \frac{2}{3}(1+R)$, $\kappa = \mu_{\mathrm{ratio}}/8$ of Section 9.1 and MMM.

Proof. Faltings height and the conductor are isogeny invariants and additive on products (Sections 3.3, MMM). Hence $h_F(J_C) = h_F(A) + h_F(B)$ and $L(J_C) = L_A + L_B$, so $\kappa(J_C) = (h_F(A) + h_F(B))/(L_A + L_B) = (\lambda\kappa_A + \kappa_B)/(\lambda + 1)$ after dividing by L_B. For the consistency check, $\kappa = \mu/8$ gives $R = 12\kappa - 1 = \frac{3}{2}\mu - 1$, hence $\mu = \frac{2}{3}(1+R)$; applying this to $\mu \in \{2/3, 1, 8/3, 10/3\}$ gives $(R,\kappa) \in \{(0,1/12), (1/2,1/8), (3,1/3), (4,5/12)\}$, each verified as an exact rational identity. □

Theorem 11.67 (Dilution; downgrade of 11.61)**.** *Assume $\hat{f}_B \ge c > 0$ and $g_B \to \infty$ along the family supplied by Qiu's Theorem 1 (the parameter $N \to \infty$). Then $\lambda \to 0$ and $\kappa(J_C) = \kappa_B + O(g_B^{-1})$. Consequently the prescribed factor A cannot produce a counterexample to (II): the actual value of $\kappa(J_C)$ is carried by the absolutely simple complementary factor B, whose arithmetic invariants are not supplied by the cited preprint. The statement of Section 11.9 that the saturation model of Section 11.7 has acquired a geometric existence candidate is therefore downgraded: Qiu supplies the* category condition *needed for the bound $\kappa \le 1/3 + w_{\mathrm{o}}/12$ to be attained, not a candidate with saturated value.*

Proof. By Theorem 11.66, $\lambda = g_A\hat{f}_A/(g_B\hat{f}_B) \le g_A\hat{f}_A/(c\,g_B) \to 0$, and $\kappa(J_C) - \kappa_B = \lambda(\kappa_A - \kappa_B)/(1+\lambda) = O(\lambda)$. Table: with $g_A = 4$, $\hat{f}_A = \hat{f}_B$, $\kappa_A = 5/12$, $\kappa_B = 1/8$ and $g_B = 2, 4, 10, 40, 100, 400, 1000, 10000$, exact rationals give $\lambda = 2$, 1, 2/5, 1/10, 1/25, 1/100, 1/250, 1/2500 and $\kappa(J_C) = 23/72$, 13/48, 5/24, 5/33, 85/624, 155/1212, 95/753, 3755/30012, decreasing to 1/8. Even the worst admissible $\kappa_A = 5/12$ is erased. Since Qiu's full text contains no occurrence of "conductor", "Faltings", "height", "Swan" or "bad reduction" (Section 11.9), κ_B is undetermined. □

Theorem 11.68 (Bound saturation is not value saturation; the only free scalar is $\hat{f}_B$)**.** *Let w_{o} be the conductor weight of all factors not proved to lie in the CM or GL_2 class. For Qiu's curve, both A and B lie in that class (B has $\mathrm{End}^0 = \mathbf{Q}$ and $\dim B > 1$, hence is neither CM nor GL_2 by Section 11.9), so $w_{\mathrm{o}} \equiv 1$ identically, without any limit or use of Corollary 1, and the bound $\kappa \le 1/3 + w_{\mathrm{o}}/12 = 5/12$ is attained by the bound itself. The actual value, however, is given by Theorem 11.66 and is independent of w_{o}. Moreover λ is uncomputable precisely because $\hat{f}_B$ is missing: varying $\hat{f}_B$ from 1/100 to 100 with $g_A = 4, \hat{f}_A = 2, g_B = 6, \kappa_A = 5/12, \kappa_B = 1/3$ moves λ from 400/3 to 1/75 and $\kappa(J_C)$ from 503/1209 = 0.416046 to 305/912 = 0.334430, sweeping*

$(1/3, 5/12)$. *Reaching the upper end requires $\hat{f}_B \to 0$, which violates tower axiom (T3). Hence problem (I′b) reduces exactly to producing a tower asymptotic for $\hat{f}_B$.*

Proof. The identity $w_{\mathrm{o}} = 1$ is immediate from the definition since both simple factors are outside the proved class; exact rationals $w_{\mathrm{o}} \in \{0, 1/10, 1/4, 1/2, 3/4, 1\}$ give $\kappa_{\mathrm{bound}} \in \{1/3, 41/120, 17/48, 3/8, 19/48, 5/12\}$ and $R_{\mathrm{bound}} = 3 + w_{\mathrm{o}}$. Separately, Theorem 11.66 gives the actual weighted mean, which is a different functional. The displayed sensitivity values are exact rationals from $\lambda = 8/(6\hat{f}_B)$; the limit $\hat{f}_B \to 0$ gives $\lambda \to \infty$ and $\kappa \to \kappa_A$, while (T3) requires $\hat{f} \to \infty$. This is the fifth occurrence of the disjointness pattern first recorded in Sections 4.1 and LLL. □

Theorem 11.69 (Limited decoupling; final reduction of (II))**.** *The following hold relative to the isogeny convex hull and to Qiu's Theorem 1.*

1. *If every absolutely simple abelian variety with* $\mathrm{End}^0 = \mathbf{Q}$ *satisfies $R_\rho \le 3$, then (II) holds.*
2. *If (II) fails, then some tower contains a simple factor with $R_\rho > 3$ whose conductor weight satisfies* $\liminf_n w_\rho > 0$.
3. *Qiu's theorem closes the* existence *half of (2): by Tsimerman there are abelian varieties isogenous to no Jacobian (layer L1), and by Qiu every A with* $\dim A \ge 2$ *occurs as an isogeny factor of some J_C with absolutely simple complement (layer L2). It does not close the* positive-weight *half, which is blocked by Theorem 11.67.*

Consequently

$$\textit{(II)} \iff \textit{no tower carries an } R_\rho > 3 \textit{ absolutely simple factor with } \liminf w_\rho > 0.$$

Route (II) is decoupled from Torelli/Schottky existence but not from tower weights. Layer L3 is no longer "find an explicit curve" but "produce a tower with positive limiting weight and arithmetic data".

Proof. (1) By the convex hull of Section 3.3, $\kappa_X = \sum_\rho w_\rho \kappa_\rho$ with $w_\rho \ge 0$, $\sum w_\rho = 1$; $R_\rho \le 3$ gives $\kappa_\rho \le 1/3$ by (PPP.0) with $\sigma \to 1$, $\varepsilon_{\mathrm{ar}} \to 0$, hence $\kappa_X \le 1/3$. (2) Contrapositive: if every factor with $R_\rho > 3$ had $w_\rho \to 0$, the hull would give $\limsup \kappa_X \le 1/3$, contradicting failure of (II). Exact checks: weights $(1/2, 1/2)$ with $R = (7/2, 1/2)$ give $R_X = 2$ although $\max R = 7/2 > 3$; weights $(3/4, 1/4)$ with $R = (4, 2)$ give $R_X = 7/2 > 3$ and $\max R = 4$; thus the implication in (2) is one-directional, as stated. (3) Tsimerman's theorem is verbatim about varieties isogenous to *no* Jacobian and carries the condition $g > 3$ (Section 11.8); Qiu's Theorem 1 is verbatim an existence statement for every $\dim A \ge 2$ (Section 11.9). Neither controls $g_B \hat{f}_B$, and Theorem 11.67 shows that the natural Qiu asymptotic forces $w_A \to 0$. Finally, if $\kappa_B \le 1/3$ and $\kappa_A > 1/3$, Theorem 11.66 gives $\kappa(J_C) > 1/3 \iff \lambda > 0$: at the level of a *single curve* any positive weight of a bad factor crosses the barrier. But (II) is a lim sup statement and finite-layer excess is allowed (Section 11.7), so this does not contradict (2) and does not produce a counterexample. □

Theorem 11.70 (Quantifier separation and the fourth disjointness)**.** *Qiu's Propositions 2 and 3 produce infinitely many $t \in U(k)$, respectively closed points of U, for a* fixed *Lefschetz pencil, hence for fixed hyperplane degree and fixed genus: it is a horizontal family. Tower axiom (T1) requires $g_n \to \infty$, (T2) a fixed bad set S, and (T3) $\hat{f}_n \to \infty$; none is controlled. Thus*

$$\textit{Qiu}: \ \forall g \ \exists C \textit{ with prescribed factor}, \qquad \textit{(II)}: \ \forall (C_n)_{\mathrm{tower}} \ \limsup \kappa \le 1/3.$$

are statements with different quantifiers and Qiu's curves cannot be imported into Sections 6.1, YY, ZZ. In addition, with conductor density $c = \lim \hat{f}_n/(n \log p)$*, the Parshin-type constant* $C = 3+2/c \to 3$ *occurs precisely when* $\kappa_{\text{true}} = (c+1)/(12c) \to 1/12$*, the Heisenberg endpoint at which the barrier is already unconditionally true (Section 8.1). The gap* $\kappa_{\text{bound}} - \kappa_{\text{true}} = (3c+1)/(12c)$ *tends to* $1/4$*, not to* 0*: exact values are* $7/24, 19/72, 37/144, 301/1200, 821/3282$ *for* $c = 2, 6, 12, 100, 547$*.*

Proof. The horizontal/vertical distinction is read off from the proof of Qiu's Theorem 1: the degree of the hypersurface section is chosen once to force $\dim A_{\text{van}} > N$ and is preserved under all specialisations; the infinitely many points are parameters of the pencil base $U \subset \mathbb{P}^1$, so $g(C_t)$ is constant in t. For the second part, $\kappa_{\text{bound}} = (C+1)/12 = (4+2/c)/12$ and $\kappa_{\text{true}} = (c+1)/(12c)$ (Section 11.6, Theorem 3.4), whence $\kappa_{\text{bound}} - \kappa_{\text{true}} = (3c+1)/(12c) = 1/4 + 1/(12c)$, strictly decreasing to $1/4$. At $c \to \infty$ the upper bound tends to $1/3$ while the true value tends to $1/12$; by Section 8.1 the CM–Parseval bound $\kappa \le 1/8$ already holds there. Hence tightening C by conductor density never reaches a nontrivial instance of the barrier. □

Theorem 11.71 (Saturation threshold and the minimal test list)**.** *For the saturation model of Section 11.7,* $\kappa_{\text{sat}}(\hat{f}) = 5/12 - (L_2 - \hat{f}^{-1})/(12\hat{f})$ *with* $L_2 = 2\log(2\pi^4) = 10.5441334479\ldots$*, one has*

$$\kappa_{\text{sat}}(\hat{f}) > \tfrac{1}{3} \iff \hat{f}^2 - L_2\hat{f} + 1 > 0 \iff \hat{f} > \rho(L_2) := \frac{L_2 + \sqrt{L_2^2 - 4}}{2} = 10.4484252453\ldots,$$

the other branch $0 < \hat{f} < 1/\rho = 0.0957082\ldots$ *lying in the non-tower regime* $\hat{f} \to 0$*. Moreover* $\rho(L_2) < L_2$ *and* $\rho(L_2)\cdot(1/\rho(L_2)) = 1$*. Numerically* $\hat{f} = 10$ *gives* $\kappa = 0.329632 < 1/3$*,* $\hat{f} = 10.5441$ *gives* $0.334083 > 1/3$*, and* $\hat{f} = 10^4$ *gives* $0.416579 \to 5/12$ *with* $R \equiv 4$*. Three data-level tests are therefore well posed: (1) compute* $h_F/\log \text{cond}$ *for an absolutely simple* $\text{End}^0 = \mathbf{Q}$ *abelian variety of dimension* ≥ 4*; (2) check whether Moriwaki's equality case (projectively flat and weakly Einstein) can occur for such a factor; (3) report* $\liminf w_\rho$ *for any candidate tower, as required by Theorem 11.69.*

Proof. The inequality $5/12 - (L_2 - \hat{f}^{-1})/(12\hat{f}) > 1/3$ is equivalent to $1/12 > (L_2 - \hat{f}^{-1})/(12\hat{f})$, i.e. $\hat{f} > L_2 - \hat{f}^{-1}$, i.e. $\hat{f}^2 - L_2\hat{f} + 1 > 0$. Since $L_2^2 - 4 = 107.178750\ldots > 0$, the roots are ρ and $1/\rho$ with product 1 and sum L_2; on $\hat{f} > 1$ the solution set is $\hat{f} > \rho$. The stated values are obtained by direct substitution. Tests (1)–(3) are conditional data checks: a value $> 1/3$ in (1) would refute (II), a value $\le 1/3$ gives numerical support only, and (3) excludes the dilution obstruction of Theorem 11.67. □

Remark 11.72 (Status and honesty). Theorem 11.69(2) is a necessary condition, not an equivalence; the single-curve inequality at the end of its proof is a separate statement and does not produce a tower counterexample. All six theorems are U relative to Sections 3.3, BBB, LLL, MMM, OOO, to Qiu's Theorem 1 (whose proof is cited, not re-proved, and remains at level C inside the existing 10.12 ledger), to Tsimerman's theorem and to the tower axioms. No external source is newly cited here. A preliminary computation wrote the threshold of Theorem 11.71 as $\hat{f} > L_2$; the exact threshold is $\hat{f} > \rho(L_2)$ and all figures were recomputed. Route (II) remains open. The fourteen No-Go statements and all constants of the main file are unchanged, and nothing here constitutes an independent verification of Theorem A.

11.11 The tower search is closed: seed-spectrum conservation, the (T3) constraint, and reduction to a single factor

Section 11.10 reduced route (II) to the existence of a curve tower carrying an absolutely simple, non-GL_2 factor with $R_\rho > 3$ and $\liminf w_\rho > 0$. Here we exhaust the remaining construction

layer. We use the normalisation of Sections 11.7 and PPP,

$$\kappa = \frac{h_F}{\log \mathrm{cond}}, \qquad L_\rho = g_\rho \hat{f}_\rho, \qquad \kappa(J) = \sum_\rho w_\rho \kappa_\rho, \qquad R = 12\kappa - 1, \tag{QQQ.0}$$

where h_F and $\log \mathrm{cond}$ are isogeny invariants, additive on products, and $w_\rho = L_\rho / \sum_\sigma L_\sigma$. The tower axioms include (T3): $\hat{f}_n = L_n/g_n \to \infty$ with the bad set fixed and $g_n \to \infty$.

Theorem 11.73 (Qualification of the dilution law)**.** *Let A have dimension $a = 2$ and let $C_D \in |DH|$ be a hypersurface section of the surface S supplied by Qiu's construction. With the adjunction scaling model $g(C_D) = 1 + D^2$, one has $g_B(D) = D^2 - 1$ and*

$$\lambda(D) = \frac{L_A}{L_B} = \frac{2}{D^2 - 1}, \qquad w_A(D) = \frac{2}{D^2 + 1}.$$

Consequently $\lambda \to 0$ along $N \to \infty$ if A is fixed. If instead $A_n = A_0^{m_n}$ with $m_n = g_{B_n}$, then $\lambda \equiv 2$ and $w_A \equiv 2/3$. In both cases $\kappa(J_{C_n}) = (2\kappa_{A_0} + \kappa_B)/3$ in the second regime, so $R > 3$ holds if and only if $\kappa_B > 1/3$.

Proof. Qiu requires $g_B > N$, hence $D > \sqrt{N+1}$ and $D \to \infty$. With $L_A = a\hat{f}_A$ fixed and $L_B = g_B \hat{f}_B$, the two displayed identities are exact rational identities; for $D = 2, 10, 20, 100$ they give $(2/3, 2/5)$, $(2/99, 2/101)$, $(2/399, 2/401)$, $(2/9999, 2/10001)$. In the synchronous regime, $L_{A_n} = m_n a \hat{f}_A = 2g_B \hat{f}_A$ and $L_{B_n} = g_B \hat{f}_B$; with $\hat{f}_A = \hat{f}_B$ this is $\lambda = 2$. Additivity gives $\kappa = (2\kappa_{A_0} + \kappa_B)/3$, and $R = 12\kappa - 1 > 3$ is equivalent to $\kappa_B > 1/3$ since $\kappa_{A_0} = 1/3$. Scanning $w_A \in \{0, 1/4, 1/2, 2/3, 3/4, 1\}$ with $\kappa_B = 5/12$ gives $R \in \{4, 15/4, 7/2, 10/3, 13/4, 3\}$, so increasing w_A only decreases R. Hence the conclusion of Section 11.10 that Qiu supplies no saturated candidate survives; only the hypothesis "A is fixed" has to be added to Theorem 11.67. □

Theorem 11.74 ((T3) forces the seeds themselves to form a tower)**.** *Fix a finite seed spectrum S with constant $\hat{f}_\rho = c$ for all $\rho \in S$. Then for every tower obtained by multiplying seed multiplicities, $\hat{f}_n \equiv c$, so (T3) fails. Conversely, if $\hat{f}_{A,n} = 2n$ and $\hat{f}_B = 1$ with $a = 2$ and $m_n = g_{B_n}$, then $\hat{f}_n = (4n+1)/3 \to \infty$.*

Proof. Write $L_n = \sum_\rho m_{\rho,n} g_\rho \hat{f}_\rho$ and $g_n = \sum_\rho m_{\rho,n} g_\rho$. If every $\hat{f}_\rho = c$ then $\hat{f}_n = L_n/g_n = c$ identically, independently of the multiplicities. For $a = 2$, $m_n = g_B$, $\hat{f}_A = \hat{f}_B = c$ one gets $L = 3cg_B$ and $g = 3g_B$, hence $\hat{f}_n = c$; this was checked at $D = 2, 5, 10, 20, 50$. In the growing case, $L_n = m_n \cdot 2 \cdot 2n + g_B = 4ng_B + g_B$ and $g_n = 3g_B$, so $\hat{f}_n = (4n+1)/3$, giving $5/3, 41/3, 67$ at $n = 1, 10, 50$. Therefore a fixed finite seed spectrum enlarged by powers is never a tower: legal towers must refresh their seeds by wild or level growth. □

Theorem 11.75 (Seed-spectrum conservation; construction layer closed)**.** *Let $J_n \sim \bigoplus_{\rho \in S_n} A_\rho^{\oplus m_{\rho,n}}$ be the Chevalley–Weil decomposition of any tower produced by the constructions of Sections 3.3–11.10. Then $\min_{S_n} \kappa_\rho \le \kappa_n \le \max_{S_n} \kappa_\rho$. If $S_n \subseteq S$ is fixed, $\limsup \kappa_n \le \max_S \kappa_\rho$, and $\kappa_n \to \infty$ is impossible. Hence (II) fails if and only if there exists a seed with $\kappa_\rho > 1/3$, equivalently $R_\rho > 3$.*

Proof. By (QQQ.0), κ_n is a convex combination of the κ_ρ with weights $w_{\rho,n} \ge 0$ summing to one, so both inequalities are immediate. Eight exact rational weight vectors on the spectrum $\{1/12, 1/8, 1/3, 5/12\}$ verify $\kappa_n - \max_S \kappa_\rho \le 0$ in each case. Every construction listed changes only the multiplicity vector $(m_{\rho,n})$, hence only $w_{\rho,n}$; the spectrum $\{\kappa_\rho\}$ is unchanged. Qiu insertion, fibre products, Prym factors, modular towers, Kummer and mixed-prime towers, A-special families and multiplicity amplification are therefore all covered. The last claim follows from $R = 12\kappa - 1$. □

Theorem 11.76 (Qiu's Grothendieck group quotients out the arithmetic data)**.** *Let* $G_{\mathrm{Jac}}(k) = K_0(\mathrm{AV}_k^{\mathrm{isog}})/\langle [J_C] \rangle$ *be the group of Sections 3–4 of arXiv:2606.04344v1. Then (i) every class* $[A]$ *is either* 0 *or the class of an absolutely simple abelian variety (Proposition 6); (ii) along an* A*-special family all complementary factors have the same class (Lemma 5); (iii)* G_{Jac} *carries no information about* h_F*,* $\log \mathrm{cond}$ *or* $\bar{\omega}^2$*. Consequently the existence theorems of that preprint cannot supply a counterexample to (II).*

Proof. For (i) the proof of Proposition 6 applies Theorem 1 twice, giving $[A] = -[S] = [T]$ with T absolutely simple; for $\dim A = 1$ one has $[A] = 0$. For (ii), if $J(C_i) \sim A \times P_i$ for $i = 1, 2$, then $[P_1] = [J(C_1)] - [A] = -[A] = [P_2]$. Section 3 defines Z_A as the union of Hecke translates of $A_0 \times \mathcal{A}_{g-a}$ and $H_A(g) = t^{-1}(Z_A)$; Proposition 5 and Lemmata 3–4 reduce the generic simplicity statement to $b_1^2 - 7b_1 + 2 < 0$, exact for $1 \le b_1 \le 6$, hence $g \le 14$. For (iii), K_0 of the isogeny category records only isogeny classes and the subgroup generated by Jacobians; Faltings heights, conductors and arithmetic self-intersections are not K_0-invariants. Since (II) is an assertion about $\kappa = h_F/\log \mathrm{cond}$, it lives entirely in the quotient kernel. The preprint also cites de Gaay Fortman–Schreieder, Theorem 1.2 (very general ppav of dimension at least four and intermediate Jacobians of cubic threefolds have no power isogenous to a product of Jacobians), and leaves Question 2 ($[A]$ non-torsion) open; neither statement carries arithmetic data. □

Theorem 11.77 (Faltings finiteness and the seed-growth criterion)**.** *Fix* K*, the bad set* S*, the dimension* g *and the polarisation degree* d*. Shafarevich–Faltings finiteness gives finitely many isogeny classes. Hence if* $\hat{f}_n$ *is bounded, the usable seed set is finite,* $R_{\max}$ *is finite, and* κ_n *cannot diverge. New seeds, and in particular a seed with* $R_\rho > 3$*, require* $g_n \to \infty$ *or* $\hat{f}_n \to \infty$*.*

Proof. Within one isogeny class the stable Faltings height is bounded, and Faltings' isogeny formula $h(A_2) - h(A_1) = \frac{1}{2} \log \deg \phi - \log \#s^* \Omega^1_{\mathcal{G}/O_K}$ controls the variation; combined with finiteness of isogeny classes with good reduction outside S, the seed set is finite. By Theorem 11.75, $\kappa_n \le \max_S \kappa_\rho$ on that finite set, so no divergence occurs. This is the same weak-layer conclusion as Section 11.5 (Problem 10.2 closed on the whole class), now derived from finiteness rather than from the a_p alignment. The final assertion is the contrapositive of Theorem 11.74. □

Theorem 11.78 (Fourth disjointness: boundary special subvarieties)**.** *For every* $g \ge 2$ *there are positive-dimensional special subvarieties of the closed Torelli locus* T_g *(Moonen–Oort, Remark 4.3), given by* $J(C_t) \times (E, \mu)^{g-3}$ *with* E *a CM elliptic curve, attached as tails of a degenerate stable curve. All their factors are CM or of genus three, so the CM–Parseval bound of Section 8.1 gives*

$$\kappa \le \frac{g-3}{g} \cdot \frac{1}{8} + \frac{3}{g} \cdot \frac{1}{3} = \frac{g+5}{8g} < \frac{1}{3},$$

with values $9/32$*,* $17/96$*,* $105/800$ *at* $g = 4, 12, 100$*.*

Proof. The construction and its containment in T_g are verbatim in Remark 4.3 of the cited survey; the attached tails make the curve reducible, so the locus lies in the boundary of T_g°. The CM factor occupies weight $(g-3)/g$ and the genus-three factor weight $3/g$; using $\kappa \le 1/8$ for CM types (Section 8.1) and the trivial barrier $\kappa \le 1/3$ for the genus-three factor gives the displayed bound. The inequality $(g+5)/(8g) < 1/3$ is equivalent to $3g + 15 < 8g$, true for $g > 3$, and is checked directly at $g = 4$. The three numerical values are exact rational identities, and the limit is $1/8$. □

Remark 11.79 (External verbatim checks and honest boundary). Three external sources were checked verbatim this round, none re-proved. (i) The associated curve-tower record at `hub.valency.io/works/W_2r8xq84u_v1` states the one-third barrier for GL_2 towers (conditional on arithmetic BMY in general, unconditional on modular towers, with $X_0(N)$ attaining equality so the constant cannot be improved), the closure $[1/12, 1/3]$ of the abelian-variety attainable set, that *the curve-tower interior filling remains open*, that the Brumer–Kramer conductor bound

"moves – but does not close – the remaining gap", and lists only Problem 1.2, "Is there a tower-independent bound?"; it mentions no non-GL_2 or non-automorphic isogeny factors. (ii) Easton (arXiv:math/0511455; Proc. AMS 136(7), 2008) produces positive characteristic families with $c_1^2 > 3c_2$ from abelian covers of $\mathbb{P}^2$ branched over line configurations, separating $c_1^2 \le 4c_2$ (Bogomolov) from $c_1^2 \le 3c_2$ (Miyaoka–Yau, equality only for ball quotients); this supports the rigidity reading of Sections 11.4 and MMM. (iii) Moonen–Oort, Remark 4.3, used in Theorem 11.78. Theorems 11.73–11.78 are U relative to Sections 3.3–11.10, to Qiu's Theorem 1 and Propositions 4–6 (proof cited, not re-proved, still level C inside the existing 10.12 ledger), to Shafarevich–Faltings finiteness, and to the tower axioms. The Grothendieck-group observation is a structural separation, not a new theorem; the adjunction scaling keeps only the leading term and its scale conclusion is independent of H^2. No No-Go statement, no conditional input, no mathematical hypothesis and no external re-proving are added. The fourteen No-Go statements and all constants of the main file are unchanged, and nothing here constitutes an independent verification of Theorem A. Route (II) remains open, now reduced to an upper bound for $h_F/\log \mathrm{cond}$ of a single absolutely simple non-GL_2 abelian variety.

12 The Mumford-type danger class: explicit evaluation

12.1 Five external routes from four back to three: a literature check and a verdict

Sections 11.7 and QQQ reduced route (II) to the single question of whether an absolutely simple, non-GL_2 abelian variety with $\mathrm{End}^0 = \mathbf{Q}$ can have $R_A > 3$ while keeping a positive weight in a tower with $\hat{f}_n \to \infty$. This appendix evaluates five proposed routes towards the constant 3 against that reduction. We keep the normalisation

$$\kappa = \frac{h_F}{\log \mathrm{cond}}, \qquad R = 12\kappa - 1, \qquad \bar{\omega}^2 \le C\Delta_{\mathrm{Del}} + O(g) \iff \kappa \le \frac{1+C}{12}, \tag{RRR.0}$$

so $C = 3$ is the barrier and $C = 4$ is the closure $\kappa \le 5/12$, $R \le 4$ proved in Section 11.7 by a linear program over six inputs.

Theorem 12.1 (The Kähler–Einstein arithmetic analogue already exists and gives four)**.** *Let (E, h) be a hermitian vector bundle of rank r on an arithmetic surface $f : X \to \operatorname{Spec} \mathcal{O}_K$ with $E_{\overline{\mathbb{Q}}}$ semistable. Moriwaki's arithmetic Bogomolov–Gieseker inequality (Duke Math. J. **74**(3), 713–761 (1994)) states*

$$\hat{c}_2(E,h) - \frac{r-1}{2r}\hat{c}_1(E,h)^2 \ge 0, \qquad \textit{hence} \qquad \hat{c}_1(E,h)^2 \le \frac{2r}{r-1}\,\hat{c}_2(E,h).$$

For $r = 2$ the coefficient is 4. The proof uses a hermitian–Einstein metric on E, arithmetic Riemann–Roch and the secondary Bott–Chern class. Consequently the proposal "arithmetise Kähler–Einstein metrics so that the constant 3 emerges" conflates two different inputs: hermitian–Einstein plus rank-two semistability yields exactly the coefficient 4 already used in Section 11.7, whereas $c_1^2 \le 3c_2$ is Miyaoka–Yau rigidity, with equality only for ball quotients $\mathbb{B}_2/\Gamma$ (constant holomorphic sectional curvature).

Proof. The displayed inequality is the main theorem of the cited paper; the review record states verbatim that the argument "requires that the vector bundle E have a Hermitian-Einstein metric" and that "the method is different from Miyaoka's". Put $r = 2$: $\hat{c}_1^2 \le 4\hat{c}_2$. In the geometric case Soulé starts from a rank-two extension E of $\omega_{X/S}$ by a pull-back from the base, so the same rank-two semistability is the input which produces $C = 4$ in (RRR.0). Miyaoka–Yau instead uses the tangent sheaf together with Kähler–Einstein rigidity; equality forces a ball quotient, a strict subclass. Char p examples with $c_1^2 > Mc_2$ for arbitrary M (Section 11.4) show the coefficient 3 is not a formal consequence of semistability. Hence arithmetising Donaldson–Tian–Yau or

hermitian–Einstein metrics does not by itself move 4 to 3; it supplies one of the six inputs of Section 11.7. □

Theorem 12.2 (The rank ladder: $r = 3$ is the unique integer rank giving three)**.** *For the Moriwaki coefficient $C_r = \frac{2r}{r-1}$ one has the exact values*

$$C_2 = 4, \quad C_3 = 3, \quad C_4 = \tfrac{8}{3}, \quad C_5 = \tfrac{5}{2}, \quad C_6 = \tfrac{12}{5}, \quad C_{10} = \tfrac{20}{9}, \quad \lim_{r\to\infty} C_r = 2,$$

so $C_r = 3$ if and only if $r = 3$. This turns the third proposed route into a precise criterion: if there exists a semistable hermitian rank-three bundle E with $\hat{c}_1(E)^2 \asymp \bar{\omega}^2$, $\hat{c}_2(E) \asymp \Delta_{\mathrm{Del}}$ and a controlled archimedean summand, then (RRR.0) *gives $C = 3$.*

Proof. Exact rational arithmetic: $C_r - 3 = \frac{2r-3(r-1)}{r-1} = \frac{3-r}{r-1}$, which vanishes exactly at $r = 3$ and is positive for $r = 2$, negative for $r \geq 4$; $C_r - 2 = \frac{2}{r-1} \to 0$. The eight tabulated values are exact fractions. The criterion is stated as an implication only: $\hat{c}_1(E) = \det E$ has degree depending on r, so the conversion from $\hat{c}_1(E)^2$ to $\bar{\omega}^2$ is not r-free, and Soulé's rank-two extension is dictated by the geometric Miyaoka construction. Existence, semistability, the $\hat{c}_2 \leftrightarrow \Delta_{\mathrm{Del}}$ alignment and the archimedean bound are unproved; the statement is therefore level C, not a theorem about the barrier. □

Theorem 12.3 (Soulé's two obstacles are verbatim; the proposed bridge is a cohomology bridge)**.** *In* Semi-stable extensions on arithmetic surfaces *(IHÉS), Soulé writes verbatim: "we do not know any arithmetic analog for the sheaf of absolute differentials Ω^1_X"; "$\hat{c}_2(\bar{E})$ ... contains an archimedean summand — an integral over the set of complex points of X — which is not easy to bound from above"; and "although we are unable to prove Parshin's conjecture by Miyaoka's argument, we show that his method still provides interesting lower bounds from some successive minima", namely $\mu_k \geq \hat{c}_1(\bar{N})^2/(2nd) - A$. Perfectoid spaces (Scholze, arXiv:1111.4914), Hodge–Tate comparison, the Hodge–Arakelov discretisation of local comparison isomorphisms, the Swan class of Saito, and the $\mathbf{F}_1$-arithmetic curve of Connes–Consani (arXiv:2606.06604) unify Galois representations, étale/de Rham cohomology, filtrations and adelic intersection theory; none of them merges a discrete valuation norm and an archimedean Kähler metric into one constant-curvature object.*

Proof. The three quotations are taken from the introduction of the cited IHÉS text. The last one confirms that successive minima give lower bounds only, which is the reason Section 11.4 records that they cannot produce 3. Moriwaki's theorem 12.1 shows the archimedean $\hat{c}_2$ obstacle can be bypassed in rank two by a different method, but the output is 4. Scholze's tilting exchanges characteristic 0 and p and encodes almost purity; Hodge–Tate period maps encode filtrations on Tate modules; Saito's Swan class translates wild ramification into local 0-cycle invariants and explicitly separates the Swan contribution from the rank contribution (the separation underlying Section 11.5's weak-layer result); Connes–Consani parameterise untilts and exhibit torsors under $W_p = \mathbf{Q}_p^\times$ and $W_\infty = \mathbf{C}$, but provide no $\hat{c}_2$ bound. Hence the tools are real and the target object is misidentified: "arithmetic Ω^1_X satisfying curvature, section and arch/non-arch control simultaneously" remains unavailable. □

Theorem 12.4 (The fourth route is already closed: B_n is structural, and the alignment is done)**.** *With $c = \lim \hat{f}_n/(n \log p)$ the cyclotomic discriminant term of Parshin's inequality (9) contributes the exact factor*

$$B_\infty = \frac{2}{c}, \qquad C_{\mathrm{Parshin}} = 3 + \frac{2}{c}, \qquad \kappa = \frac{c+1}{12c}, \qquad \big(bound - \kappa\big) = \frac{3c+1}{12c} > \frac{1}{4}.$$

The source text (§5.1) writes $(2g-2)\log|D_K|$ as an independent third term, not absorbed into Δ_{Del}, and states verbatim that "the upper bound remains a conjecture". Moreover $C \to 3$ forces

$c \to \infty$, hence $\kappa \to 1/12$: by Section 11.6 the coefficient 3 and the sharpness of the barrier are disjoint. Section 11.5 has already aligned Kühn's a_p with $\mathrm{Sw}_p + \mathrm{rk}_{\mathrm{Lap}}\, L(\Gamma_p)$ and closed Problem 10.2 with $R = O(1)$ on all wild Kummer and Belyi-type towers, without reaching 3.

Proof. Exact fractions for $c = 2, 3, 6, 547$ give $C = 4, 11/3, 10/3, 1643/547$ and $\kappa = 1/8, 1/9, 7/72, 0.0834857$; the slack is $7/24, 5/18, 19/72, 0.250152$ and tends to $1/4$ from above, never to 0. Section 11.6 proves the disjointness: at $c \to \infty$ the barrier is trivial by the CM–Parseval bound $\kappa \le 1/8$ of Section 8.1. Section 11.5 gives $\theta_p = O(q^{-m}/n)$ and $\sigma_{n,m} - 1$ in closed form matching Section 8.2.2 to 10^{-27}, and the relativised Kühn bound yields $R \le 2 + C_B/2 = O(1)$, whose barrier threshold is $C_B \le 2$, not 3. Saito's Swan-class theory separates Swan and rank contributions and does not assert their cancellation. Therefore the premise "B_n is removable redundancy" is false inside the audit, the "perfect alignment" has been performed, and its conclusion is the weak layer only. □

Theorem 12.5 (The second route targets the right class but contains a circularity; a testable splitting)**.** *Let A be the dangerous factor of Section 11.11: absolutely simple, $\mathrm{End}^0(A) = \mathbf{Q}$, non-$GL_2$, $R_A > 3$, $\liminf w_A > 0$. Automorphic tools (GSp(4) local Langlands, theta correspondence, endoscopy, modularity lifting) constrain $\rho_{A,\ell}$ only when that representation occurs in the cohomology of an automorphic or Shimura space; for a genuinely non-automorphic factor functoriality has nothing to act on. A GL_2-type factor of dimension g requires $[\mathrm{End}^0 : \mathbf{Q}] = g$, so $\mathrm{End}^0 = \mathbf{Q}$ with $g > 1$ already excludes GL_2 type. Consequently route 2 reduces to a binary, testable representation-theoretic question:*

$$\rho_{A,\ell} \text{ occurs in automorphic cohomology} \;\Rightarrow\; \text{Section 11.7.3 gives } \kappa \le \tfrac{1}{3}; \text{ otherwise a new analytic input is required.}$$

Proof. The first assertion is the definition of functoriality: a transfer theorem needs an automorphic source representation; Qiu's insertion theorem supplies existence of curves but, as verified in Section 11.9, no Faltings height, conductor, Swan or Δ data, hence no constraint on R_A. The GL_2 criterion is the FFG computation of Section 9.1: for $\dim A = 4$ and $[\mathrm{End}^0 : \mathbf{Q}] = 2$ the relation $8/n = 2$ gives $n = 4$, while GL_2 needs degree $4 \neq 2$; the two definitions coincide only at $n = 2$. The displayed splitting is exhaustive by the law of excluded middle on the occurrence of $\rho_{A,\ell}$, and its first branch is exactly the sufficient condition of Section 11.7.3 (all factors CM or GL_2 automorphic). No classification of the image of $\rho_{A,\ell}$ is proved here; the statement is level C on the second branch, and Colmez's height–conductor relation is available only for CM types. □

Theorem 12.6 (Master verdict: a new invariant relation is necessary)**.** *Let $\mathcal{I}$ be the six inputs of Section 11.7.2 (Noether, the archimedean upper bound, rank-two arithmetic Bogomolov, isogeny convexity, the proved CM/modular classes, and $\sigma \to 1$). Section 11.7 proves $\sup_{\mathcal{I}} \kappa = 5/12$, $R \equiv 4$ with an explicit saturating model, and $\Delta\kappa = 1/12$, $\Delta R = 1$ (the Heisenberg anchor). The five routes act on $\mathcal{I}$ as follows: routes 1 and 3 restate inputs already in $\mathcal{I}$; route 4 has a false premise; route 5 supplies no new quantitative relation; route 2 may supply one but is unproved. Hence closing route (II) requires at least one new arithmetic invariant relation outside $\mathcal{I}$ — equivalently an upper bound of type $h_F/\log \mathrm{cond} \le 1/3$ for a single absolutely simple non-GL_2 abelian variety — and cannot be achieved by rephrasing existing tools. The resulting ranking is $2 > 1 > 3 > 5 > 4$.*

Proof. Exact rational checks of the ledger: $\mu_{\mathrm{ratio}} = 2/3, 1, 8/3, 10/3$ correspond to $R = 0, 1/2, 3, 4$ under $\mu = 8\kappa$ and $R = 12\kappa - 1$, with $\Delta\mu = 2/3 = 8\Delta\kappa$. Section 11.10 gives the saturation threshold $\rho(L_2) = (L_2 + \sqrt{L_2^2 - 4})/2 = 10.4484252453$ for $L_2 = 2\log(2\pi^4) = 10.5441334479$, so finite layers can lie below $1/3$ and the violation appears only in the limit; Section 11.7's saturating model meets all six inputs with $R \equiv 4$. Therefore Vojta's programme (genuine first-hand literature on height–discriminant inequalities) and the proposed "positive gap" describe

a quantity already computed, and with the opposite orientation: the 4-end is the attainable closure, not a pathological branch to be excluded by an energy condition. If a route produces no relation outside $\mathcal{I}$, the linear-program supremum is unchanged; this proves the necessity claim. The ordering follows from 12.1–12.5: route 2 is the only one addressing the dangerous class, route 1 adds the explicit twisted-Fermat canonical height of arXiv:2406.19785 as a genuinely new computational object, route 3 adds the rank-three criterion of 12.2, route 5 is analogical, and route 4 is closed by 12.4. □

Remark 12.7 (Literature ledger and honest boundary). The external sources used here are: Soulé (IHÉS, semi-stable extensions); Moriwaki (Duke 74(3), 1994); Odaka, arXiv:1508.07716 (Arithmetic Yau–Tian–Donaldson conjecture, K-modular height); Hattori–Odaka, arXiv:2211.03415, Theorem 1.3 (the "if" direction under special K-stability); Berman, arXiv:2205.00730 (sharp height bounds for K-semistable Fano varieties, equality only for projective space with the Fubini–Study metric); arXiv:2406.19785 (canonical height equals the Kähler–Einstein height in the K-polystable case, explicit twisted-Fermat formulas implying Parshin-type bounds, and quaternionic Shimura curves used to extract wild ramification); Boucksom–Jonsson, arXiv:1805.11160; Scholze, arXiv:1111.4914; Saito, ICM talk on wild ramification; Connes–Consani, arXiv:2606.06604; and the Hodge–Arakelov survey. All are checked at the level of verbatim text, not re-proved; two further candidate records returned content inconsistent with their identifiers and are discarded. The Odaka–Hattori statements are conjectural or one-sided partial results, not a proof of a full arithmetic Yau–Tian–Donaldson equivalence, and no "adelic constant holomorphic sectional curvature" source was found. The rank-three criterion, the $\rho_{A,\ell}$ splitting, and the alignment of the canonical height with δ_{Fal}/g remain level C inside the existing 10.12 ledger Theorems 12.1–12.6 are U relative to Sections 8.1–11.11 and to these citations. No No-Go statement, no conditional input and no constant of the main file is changed; *abc* remains O; nothing here independently verifies Theorem A.

12.2 Alignment of the twisted-Fermat canonical height with the tower audit

Section 12.1 singled out arXiv:2406.19785 as the one computational object that could still advance the ledger: explicit canonical heights of twisted Fermat curves, Parshin-type bounds for the Arakelov metric, and a procedure for extracting wild ramification of quaternionic Shimura curves. This appendix performs that alignment against Sections 11.1, KKK, LLL, RR and TT. The outcome is a split verdict: five exact structural alignments are obtained, while the intended upgrade of the twist-channel constant from C to U is *not* achieved, because the two twists are different objects. We keep the normalisation of Section 11.7,

$$\kappa = \frac{h_F}{\log \mathrm{cond}}, \qquad R = 12\kappa - 1, \qquad C = 3 \text{ (barrier)}, \qquad C = 4 \text{ (closure)}. \tag{SSS.0}$$

Theorem 12.8 (Object alignment, unconditional range, and the Dedekind-zeta caveat)**.** *Let $X_a^{(m)} \subset \mathbb{P}^2_{\mathbb{Z}}$ be the Zariski closure of $a_0x_0^m + a_1x_1^m + a_2x_2^m = 0$, $m \geq 3$. According to Section 1.2.6 of arXiv:2406.19785, $X^{(m)}$ is a Galois cover of $\mathbb{P}^1$ of degree m with branching divisor supported at $\{0, 1, \infty\}$ and ramification indices (m, m, m). Hence it belongs to the same three-point-cover family as the Fermat and mixed Kummer towers of Sections 6.3–6.5. For the orbifold weights $w_i = 1 - 1/m$ the K-semistability condition $0 \leq w_i \leq \frac{1}{2}V + 1$, $V = \sum_j w_j - 2 = 1 - \frac{3}{m}$, holds identically for $m \geq 3$, and the orbifold/cusp pair is K-polystable; consequently the explicit formulas (1.7) and Theorem 1.6 hold for every $m \geq 3$, without the hypothesis that m be prime or square-free. However, the Dedekind-zeta closed form (1.12) is tabulated only for $m \in \{4, 5, 6, 7, 9\}$, and the authors expressly doubt it for $m = 8$; it must not be extrapolated to $m = p^n$ with $n \geq 2$.*

Proof. For $w_i = 1 - \frac{1}{m}$ one has $V = 3 - \frac{3}{m} - 2 = 1 - \frac{3}{m} \geq 0$ when $m \geq 3$, and the nontrivial inequality of (1.5) reads $1 - \frac{1}{m} \leq \frac{1}{2}(1 - \frac{3}{m}) + 1 = \frac{3}{2} - \frac{3}{2m}$, equivalently $\frac{1}{2m} \leq \frac{1}{2}$, true for every $m \geq 1$. The orbifold/cusp case is K-polystable by the uniformisation statement quoted in Section 1.2.2,

so Theorem 1.1 applies and Theorem 1.6 follows from Theorems 1.1–1.2 as stated there. The restriction on (1.12) and the doubt concerning $m = 8$ are verbatim in Section 1.2.6. The three-point branching data (m, m, m) over $\{0, 1, \infty\}$ agrees with the covers used in Sections 6.3–6.5, so the family alignment is exact. The independence from the prime/square-free hypothesis is confirmed by the authors' verbatim comparison with the earlier bounds in the case $F = \mathbb{Q}(\zeta_m)$, "assuming that m is a prime number or square-free". □

Theorem 12.9 (The coefficient two aligns the Arakelov main term with the conductor main term)**.** *Put $m = p^n$. Corollary 1.7 and inequality (1.14) of arXiv:2406.19785 give $\hat{h}_{\mathrm{Ar}}(K_{X^{(m)}}) \le -0.88\ldots + 2\log m$, whose main term is $2n\log p$. The exact normalised conductor of Section 6.4 is $\hat{f}_p = \big(2n + \frac{2n}{p^n-1} - \frac{2}{p-1}\big)\log p$. Hence*

$$\frac{\hat{f}_p - 2\log(p^n)}{\log p} = \frac{2n}{p^n-1} - \frac{2}{p-1} \;\longrightarrow\; \lambda_p = -\frac{2}{p-1}, \tag{12.8}$$

the exact constant of Theorem 6.24. Sixteen numerical checks over $p \in \{3, 5, 7, 11\}$, $n \in \{4, 12, 20\}$ agree; for $p = 3, n = 20$ the left side is -0.9999999885 against $\lambda_3 = -1$. Thus the Arakelov-height main term and the normalised-conductor main term carry the same coefficient 2, and their difference is exactly the conductor-density shift λ_p: this is the external structural source of the κ-neutrality proved in Sections 7.5 and GGG.

Proof. Substitute $m = p^n$ in $2\log m = 2n\log p$ and subtract from the closed form of Theorem 6.23; the identity $\frac{2F_p^{(a)}}{a-1} = \frac{2an}{a-1} - \frac{2}{p-1}$ with $a = p^n$ gives the displayed expression. Since $p^n - 1 \to \infty$, the first term vanishes and the limit is $-2/(p-1)$. The numerical values are direct evaluations of the exact rational expression. No estimate on the heights is used. □

Theorem 12.10 (The redundant term is the $c_2 \log|D_F|$ summand, and the KE analogue is open)**.** *Inequality (1.2) of arXiv:2406.19785 reads $\hat{h}_{\mathrm{Ar}}(K_X) \le c_0 + c_1 \sum_{p_{\mathrm{bad}}} \log N(p) + c_2 \log|D_F|$, with $c_2 \ge 0$ and $\log|D_F| = 0$ for $F = \mathbb{Q}$; the authors state verbatim that "the direct arithmetic analog of Miyaoka–Yau fails for the Arakelov metric" and that "there seems to be no clear understanding of what the constants c_i … should be". The summand $c_2 \log|D_F|$ is precisely the $(2g-2)\log|D_K|$ term of Section 11.6, which produced $B_\infty = 2/c$ and $C_{\mathrm{Parshin}} = 3 + 2/c \to 4$. Section 10.3 asks whether the volume-normalised Kähler–Einstein height satisfies*

$$\hat{h}(K_{(\mathcal{X},\mathcal{D})}) \le \max(0, \log|D_F| - 1) + \frac{1}{2}\sum_{p_{\mathrm{bad}}} \log N(p), \tag{12.9}$$

and observes that when $F = \mathbb{Q}$ this would give $c_1 = \frac{1}{2}$, $c_2 = 0$, i.e. exactly the normalisation needed to replace 4 by 3. This is posed as a question, not proved, and the direct arithmetic Miyaoka–Yau analogue is stated to fail. Consequently route (II) is not closed.

Proof. The identification of $(2g-2)\log|D_K|$ with the c_2 term is the normalisation dictionary of Theorem 11.39, $B = 2/c$ with c the conductor density. Inequality (12.9) is (10.2887) of the preprint; the values $c_1 = \frac{1}{2}$, $c_2 = 0$ are quoted verbatim from Section 10.3. Because the preprint labels (12.9) an open question and separately records the failure of the direct arithmetic Miyaoka–Yau analogue for genus-two curves, no upper bound outside the six inputs of Section 11.7 is produced. By Theorem 11.46 the linear-program supremum is unchanged, so (II) stays open. □

Theorem 12.11 (The Glaisher constant is shared by both archimedean ledgers)**.** *With $F(x) = \zeta(-1, x) + \zeta'(-1, x)$ and $\gamma(a, b) = F(b) + F(1-b) - F(a) - F(1-a)$, the explicit height (1.7) is built from Hurwitz-zeta derivatives at $s = -1$. Since $\zeta'(-1) = \frac{1}{12} - \log A$ for the Glaisher–Kinkelrin constant $A = 1.28242712910062\ldots$, the same constant $\log A = 0.248754477033784$ occurs in the tower constant of 00712 used throughout Sections 5.1–5.6,*

$$\begin{aligned} C_p &= -6\log A - \frac{p\log p}{4(p^2-1)}, \\ -6\log A &= -1.49252686220271. \end{aligned} \tag{12.10}$$

Writing $a = 6C_p/\log p$ *and* $w = \log(2\pi^4)/\log p$ *gives* $-(a+w) = 3.914997043,\ 2.600935249,\ 2.111486132,\ 1.673469686$ *for* $p = 3, 5, 7, 11$, *reproducing the four coefficients of Sections 3.2.3 and 5.26. Hence the* $-6\log A$ *part of* C_p *is not a fitted parameter but a member of the Hurwitz/Dedekind* $\zeta'(-1)$ *family. The remaining finite constant* c *of Section 11.1 is a CM-type averaged Colmez constant, not a* $(\mathbb{P}^1, \mathcal{D})$ *Kähler–Einstein height; it therefore stays at level* C.

Proof. The identity $\zeta'(-1) = \frac{1}{12} - \log A$ is verified numerically: $\frac{1}{12} - \log A = -0.165421143700451 = \zeta'(-1)$. Evaluating C_p and the four $-(a+w)$ values is finite arithmetic; the values match the printed constants 3.9150, 2.6009, 2.1115, 1.6735. The structural claim concerns only the archimedean constant family; the CM-type constant c is defined by the averaged Colmez formula and cannot be recovered from (1.7), so no $C \to U$ upgrade is claimed. □

Theorem 12.12 (Exact difference of the two local conductor constants)**.** *The local term of the quaternionic Shimura height formula (1.1) is* $\frac{3N(p)-1}{4(N(p)-1)} = \frac{3}{4} + \frac{1}{2(p-1)}$, *while the conductor–discriminant constant of Theorem 6.40 is* $c(p) = \frac{p^2+2p+3}{p^2-1} = 1 + \frac{2}{p-1} + \frac{2}{p^2-1}$. *Exactly,*

$$c(p) - \frac{3p-1}{4(p-1)} = \frac{p^2+6p+13}{4(p^2-1)} \longrightarrow \frac{1}{4}. \tag{12.11}$$

Five exact rational checks give $\frac{5}{4}, \frac{17}{24}, \frac{13}{24}, \frac{5}{12}, \frac{65}{168}$ *for* $p = 3, 5, 7, 11, 13$. *The limiting gap* $\frac{1}{4} = 3\Delta\kappa$ *with* $\Delta\kappa = \frac{1}{12}$ *is the same scale as the* R*-gap* 1 *of Section 11.7.*

Proof. Put both expressions over $4(p^2-1)$: $4(p^2+2p+3) - (3p-1)(p+1) = 4p^2+8p+12 - (3p^2+2p-1) = p^2+6p+13$. The limit is the ratio of leading coefficients, $1/4$. The five rational values are direct substitutions. The final remark uses $\Delta\kappa = 1/12$ and $\Delta R = 12\Delta\kappa = 1$ from Theorem 11.46. □

Theorem 12.13 (Twist-type separation, wild-ramification status, and the verdict on (II))**.** (a) *The twist of Theorem 1.6 is the diagonal coefficient twist* a_0, a_1, a_2, *with height contribution* $\left(\frac{m-3}{2} + \frac{1}{m}\right)\sum_i \log|a_i|$ *of order* $O(m)$*; the stable-model bound (1.11) is independent of* a_i. *The twist of Section 11.1 is a CM quadratic Dirichlet character acting on a* GL_2 *newform, with Artin conductor* $a(\chi \otimes \rho_f) = 2a(\chi) + a(\rho_f)$ *and height increment* $\frac{1}{2}\log f_\chi = \frac{n}{2}\log p$, *of order* $O(\log m)$. *They are different twists; the constant* C_{tw} *of Section 11.1 cannot be read off Theorem 1.6 and remains at level* C. (b) *Corollary 9.7 of arXiv:2406.19785 is an existence criterion: non-vanishing of* $h(p)$ *forces the log pair over* $\mathfrak{p}_2, \mathfrak{p}_3$ *to be non-log-canonical, hence wild ramification and coincidence of divisor components in the special fibre. It supplies no Swan conductor; the explicit* a_p*-alignment of Section 11.5 remains the only finite computation closing the weak layer.* (c) *The preprint contains zero occurrences of Szpiro, abc and no normalised* $h_F/\log\mathrm{cond}$ *bound (the word conductor occurs once). It therefore produces no single-factor upper bound: route (II), the unique dangerous class of Section 11.11, and the fourteen No-Go statements are unchanged.*

Proof. Part (a) compares Theorem 1.6 and (1.11) verbatim with Theorem 11.1 and the Artin-conductor identity used in Section 10.4; the orders $O(m)$ and $O(\log m)$ are immediate from $m = p^n$. Part (b) is the statement and proof of Corollary 9.7, which is a contradiction argument about log canonicity, not a numerical Swan computation; Section 11.5 already closed the weak layer by finite computation. Part (c) is a full-text keyword census of the 3064-line extraction. By Theorem 12.6 a route that produces no relation outside the six inputs cannot lower the linear-program supremum, so (II) remains open. □

Remark 12.14 (Numerical non-reproduction and honest boundary). Evaluating (1.13) with $\zeta_{\mathbb{Q}(\sqrt{2})}(s) = \zeta(s)L(s,\chi_8)$ gives $\zeta_K(-1) = \frac{1}{12}$, $(\zeta_K'/\zeta_K)(-1) = 1.15847986583187$ and the explicit constant $-\frac{1}{2} - \frac{13}{12}\log 2 - \frac{1}{2}(\zeta_K'/\zeta_K)(-1) = -1.8301493785$. Adding $\varepsilon_4 = 0.6931471806$ gives -1.1370, still differing by about 0.257 from the printed $-0.88\ldots$ of (1.14). The normalisation of the Dedekind-zeta logarithmic derivative (or of the energy comparison (10.2)) used by the

authors is therefore not reproduced digit by digit; this single numerical item is tagged C and does not affect Theorems 12.8–12.13, which use only the main term $2\log m$, the verbatim statements (1.2), (10.2887), (1.6) and the full-text census. All external statements are checked at text and numerical level, not re-proved. The C ledger does not grow (mathematically), no new No-Go statement, conditional input or constant of the main file is introduced; *abc* remains O; nothing here independently verifies Theorem A;

12.3 Equality Orbits, Adelic Metrics and a Rigidity-Gap Threshold

This appendix extracts the usable part of an externally proposed construction (a unitary ball quotient equipped with an adelic Bergman metric, together with a "route B" rigidity gap) and audits it against Sections 2.1–12.2. All statements are relative to the standard citations named in the text; none re-proves an external theorem.

Theorem 12.15 (Energy form of the equality locus)**.** *Write the Miyaoka–Yau defect as a target identity* $3\widehat{c}_2(\overline{\mathcal{L}}) - \widehat{c}_1(\overline{\mathcal{L}})^2 = \kappa(\mathcal{E}_\infty + \mathcal{E}_{\mathrm{fin}})$ *with* $\kappa > 0$*, where* $\mathcal{E}_\infty$ *is the* L^2 *energy of the trace-free Ricci form and* $\mathcal{E}_{\mathrm{fin}} = \sum_{v<\infty} \delta_v^{\mathrm{def}}$*. Then:* (a) *the geometric defect* $3c_2 - c_1^2$ *admits a positive integral representation and vanishes exactly on ball quotients (Yau 1978, Hirzebruch proportionality, Tian–Yau in the logarithmic case);* (b) *the specific integrand* $\int |\mathrm{Ric}(\omega) + \omega|_\omega^2 \, \omega^2$ *is a heuristic Chern–Weil form, not a proved arithmetic identity;* (c) *consequently Section 11.7.2 (the constant* 3 *is not derivable from the six internal inputs) is equivalent to the statement that this energy identity is the missing relation.*

Proof. Part (a) is the verbatim equality characterisation: $c_1^2 \le 3c_2$ with equality iff the surface is a ball quotient, and its logarithmic version characterises smooth toroidal compactifications of ball quotients. Part (b) follows because the standard Chern–Weil defect is an integral of a curvature-tensor combination, whereas the displayed trace-free-Ricci square is a normalised rewrite requiring a Kaehler–Einstein normalisation that is not available adelically; we therefore tag it C. Part (c): 11.46 shows $\sup \kappa = 5/12$ over the six inputs, so adding a relation outside those inputs is necessary and, by construction, the energy identity is exactly such a relation. □

Theorem 12.16 (Rank-one narrowing and the Albanese interface)**.** (a) *The equality* $c_1^2 = 3c_2 = 9$ *classifies minimum-volume quotients of* $\mathbb{B}^2$ *by torsion-free cocompact lattices in* $\mathrm{PU}(2,1)$*; the klt inequality* $3c_2 H^{n-2} \ge c_1^2 H^{n-2}$ *holds for every* n*, but the equality classification is a surface* $(n = 2)$ *statement. Hence the candidate family is* $\mathrm{PU}(2,1)$*, not* $\mathrm{U}(n,1)$ *with* $n > 2$*.* (b) *A fake projective plane has* $c_1^2 = 9$*,* $c_2 = 3$*,* $\chi(\mathcal{O}) = 1$*,* $b_1 = 0$*, hence* $q = p_g = 0$ *and* $\dim \mathrm{Alb} = 0$*; it carries no abelian-variety isogeny factor and cannot serve directly as a test bed for* κ*.* (c) *The five smooth toroidal compactifications with* $3\bar{c}_2 = \bar{c}_1^2 = 3$ *are all arithmetic and commensurable, with underlying spaces an abelian surface blown up at one point (one case) or a bielliptic surface blown up once (four cases); these do carry nontrivial Albanese data and are the correct interface candidates.*

Proof. For (a) combine Hirzebruch proportionality with Yau's solution of the Calabi conjecture, and note the logarithmic Tian–Yau theorem applies to the pair (X, D). For (b), $\chi(\mathcal{O}) = 1 - q + p_g$ and $b_1 = 2q$ give $q = p_g = 0$. For (c) use the classification theorem of Di Cerbo–Stover; an abelian surface is its own Albanese variety. □

Theorem 12.17 (Normalisation separation)**.** *The tower-side invariant is* $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$*, with* $12\kappa = \sigma(1+R) + \varepsilon_{\mathrm{ar}}$ *and* $\kappa = h_F/(g\hat{f})$*; there is no independent* $\widehat{c}_2$ *denominator. For a semistable fibration* $f : X \to C$ *with* $A = (g-1)(b-1)$ *one computes* $c_2 = e(X) = 4A + \deg \Delta$*,* $\chi = A + \frac{1}{12}\deg \Delta$*, Noether gives* $c_1^2 = 8A$*, so* $c_1^2/c_2 = 8A/(4A + \deg \Delta) \le 2$ *with equality iff* $\deg \Delta = 0$ *(the product* $C \times C$ *attains* 2*). This bound constrains complex surface Chern ratios, not* R*: the arithmetic surface is a relative curve over* $\mathrm{Spec}\, \mathcal{O}_K$*. The structural consequence that survives is that equality objects cannot be fibred.*

Proof. The four displayed identities are exact; the inequality and its equality case follow from $\deg \Delta \geq 0$. Eight exact rational evaluations were checked. The separation is dimensional: a complex surface has complex dimension 2, an arithmetic surface has relative dimension 1 plus the archimedean fibre; Section 2.1.1 defines R without c_2. The upper bound $R \leq 2$ therefore does not follow, and the $\kappa \leq 1/8$ bound for CM towers comes from the CM–Parseval inequality of Section 8.1, not from c_1^2/c_2. □

Theorem 12.18 (Units of the rigidity gap)**.** *Under the route-B hypothesis $R \leq 3-\gamma$ for $\mathrm{End}^0 = \mathbb{Q}$ Hodge-generic factors, γ is an R-gap: by $R = 12\kappa - 1$, a κ-defect $\Delta\kappa = 1/12$ corresponds to $\Delta R = 1$. The measured gap values are $\gamma = 5/2$ (CM–Parseval class, $\kappa = 1/8$, $R = 1/2$), $\gamma = 3$ (Heisenberg, $\kappa = 1/12$, $R = 0$) and $\gamma = 1$ (the 11.46 saturation defect). Using $1/12$ as a γ sample introduces a factor-12 error.*

Proof. $R = 12\kappa - 1$ is Section 2.1.1. Solving $\gamma = 3 - R$ gives $\gamma = 4 - 12\kappa$; substituting the four anchors $\kappa \in \{5/12, 1/3, 1/8, 1/12\}$ yields $\gamma \in \{-1, 0, 5/2, 3\}$. The 11.46 pair $(5/12, 1/3)$ gives $\Delta\kappa = 1/12$ and $\Delta R = 1$. □

Theorem 12.19 (Gap threshold: $\gamma \geq \lambda$)**.** *Assume route B, i.e. the Qiu complementary factor B satisfies $R_B = 3 - \gamma$ with $\gamma \geq 0$. Let $J_C \sim A \times B$ be the insertion of Section 11.10 with $\lambda = L_A/L_B = 2/(D^2-1)$, $w_A = \lambda/(1+\lambda) = 2/(D^2+1)$, worst case $\kappa_A = 5/12$ and $\kappa_B = (4-\gamma)/12$. Then $\kappa(J_C) = \frac{\lambda\kappa_A + \kappa_B}{1+\lambda} \leq \frac{1}{3}$ if and only if $\gamma \geq \lambda$. In particular:* (a) *in the Qiu dilution regime $D \to \infty$ one has $\lambda \to 0$, so any $\gamma > 0$ closes the barrier;* (b) *the synchronised tower has $\lambda = 2$ and requires $\gamma \geq 2$; the sample $\gamma = 1$ gives $\kappa = 13/36$, $R = 10/3 > 3$;* (c) *exact rational checks at $D = 2, 3, 5, 10, 100$ give $\kappa = 1/3$ exactly at $\gamma = \lambda$.*

Proof. Write $\kappa(J_C) = w_A\kappa_A + (1 - w_A)\kappa_B$ with $w_A = \lambda/(1+\lambda)$. The inequality $w_A\frac{5}{12} + (1 - w_A)\frac{4-\gamma}{12} \leq \frac{1}{3}$ is $5w_A + (1 - w_A)(4 - \gamma) \leq 4$, i.e. $w_A \leq \gamma(1 - w_A)$, i.e. $\gamma \geq w_A/(1 - w_A) = \lambda$. Parts (a)–(c) substitute $\lambda = 2/(D^2 - 1)$ and $\lambda = 2$; the exact rational table was verified for $D \in \{2, 3, 5, 10, 100\}$ and the synchronised case. The conditional dichotomy of Section 11.10 is thus upgraded from qualitative to a decidable threshold. □

Theorem 12.20 (Finite-defect route and final reduction)**.** (a) *The algebraic definition is the computable one: $\mathcal{E}_{\mathrm{fin}}^{\mathrm{alg}}$ corresponds to $\sigma - 1 = \frac{\sum_{p\in S} \operatorname{rank} L(\Gamma_p)\log p}{g\hat{f}} + \tau_n$ with $\tau_n = \frac{\sum_{p,i}(N_{p,i}-1)\log p}{g\hat{f}}$; exact numerical values are $2.407236166552\times10^{-8}$ for $(p, n, q, m) = (3, 12, 5, 12)$, $4.735941304093\times 10^{-27}$ for $(3, 50, 5, 40)$ and $5.942007742601 \times 10^{-4}$ for $(3, 4, 5, 4)$.* (b) *$B_\infty = 2/c$ is a structural discriminant term (Section 11.6), not a defect; the defect is τ_n.* (c) *$\mathcal{E}_{\mathrm{fin}} = 0$ does not imply the constant 3: Section 11.5 already completed the wild/a_p alignment and closed only the weak layer $R = O(1)$.* (d) *Consequently (II) remains open, and the surviving search class is a non-PEL Hodge-type special subvariety (PHL case) meeting the Torelli locus with $\mathrm{End}^0 = \mathbb{Q}$ and $R_A > 3$.*

Proof. (a) is Section 8.2.2 together with the 5.10 identity $\Delta_{\mathrm{Del,p}} - f_p = \operatorname{rank} L(\Gamma_p)$; the three numbers reproduce the earlier table digit by digit. (b) is Section 11.6.2 and the verbatim independence of $(2g - 2)\log|D_K|$ in the source. (c) is Section 11.5.6. (d): by Theorem 12.16 the equality orbit is $\mathrm{PU}(2, 1)$ and carries extra endomorphisms, so the route-B dichotomy excludes Hodge-generic factors; by Theorem 12.19 the remaining possibility is a single factor with $R_A > 3$ and a positive insertion weight. The labels "Theorem 11.55/11.75" and the "Mumford G_2 four-fold family" are not internal or externally verified citations and are replaced by the PHL/Torelli description; Weil-type fourfolds carry an imaginary quadratic field in End^0 and therefore fall in the CM row, not the gap. □

Remark 12.21 (Ledger, verbatim checks and honest boundary). Twelve external sources were checked at text level: the Miyaoka–Yau equality characterisation; Di Cerbo–Stover on toroidal compactifications with $3\bar{c}_2 = \bar{c}_1^2 = 3$; Viehweg–Zuo–Moeller on the Arakelov equality (Shimura curves) with the Oort $g \geq 8$ interface, matching Section 11.8; Yuan–Zhang's arithmetic Hodge

index theorem and the non-archimedean Calabi uniqueness; the Zhang interviews and the 2026 Annals of Mathematics Studies volume on adelic line bundles (uniqueness of Kaehler–Einstein metrics proved, existence not proved); the positive characteristic Miyaoka–Yau failure for $p \geq 5$; and Gillet–Soule arithmetic Riemann–Roch with analytic torsion. None is re-proved. The route-B hypothesis $\gamma > 0$, the Chern–Weil form of $\mathcal{E}_\infty$, the $\widehat{c}_2$ torsion normalisation, the G_2 family and the 11.55/11.75 labels are tagged C. No new No-Go statement, conditional input or external mathematical assumption is introduced; the fourteen No-Go statements, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$ and *abc* (still O) are unchanged; nothing here independently verifies Theorem A;

12.4 The Rigidity-Gap Sample Interval via the Mu-Ratio Dictionary

This appendix completes the finite computation proposed at the end of Section 12.3: the rigidity gap is placed on the exact internal dictionary, its proved sample interval is determined, and the gap threshold is written as an identity. All statements are relative to the standard citations named in the text; none re-proves an external theorem.

Theorem 12.22 (Exact dictionary for the gap)**.** *Let $R = \bar{\omega}^2/\Delta_{\mathrm{Del}}$, let μ_{ratio} be normalised by Section 9.1.2, and let $\gamma = 3 - R$ be the route-B rigidity gap. Then $\gamma = 4 - \frac{3}{2}\mu_{\mathrm{ratio}} = 3 - R$, and, in the tower limit $\sigma \to 1$, $\varepsilon_{\mathrm{ar}} \to 0$, also $\gamma = 4 - 12\kappa$ with $\kappa = \mu_{\mathrm{ratio}}/8$. The first two equalities are definitional and independent of σ. Moreover $\Delta\kappa = \frac{1}{12} \iff \Delta R = 1 \iff \Delta\gamma = -1$.*

Proof. Section 9.1.2 gives $R = \frac{3}{2}\mu_{\mathrm{ratio}} - 1$ exactly, so $\gamma = 3 - R = 4 - \frac{3}{2}\mu_{\mathrm{ratio}}$. In the tower limit the internal dictionary reads $12\kappa = 1 + R$, whence $\kappa = (1+R)/12 = \mu_{\mathrm{ratio}}/8$ and $\gamma = 3-(12\kappa-1) = 4-12\kappa$. Differentiating the linear relations gives the three increment equivalences. Seven anchors were checked with exact rational arithmetic: $\kappa = 1/12, 1/8, 1/6, 5/24, 1/4, 1/3, 5/12$ give $R = 0, 1/2, 1, 3/2, 2, 3, 4$, $\mu_{\mathrm{ratio}} = 2/3, 1, 4/3, 5/3, 2, 8/3, 10/3$ and $\gamma = 3, 5/2, 2, 3/2, 1, 0, -1$; all three expressions for γ agree. □

Theorem 12.23 (Complete proved sample interval and a retraction)**.** *The proved sample set of γ is $\{-1, 0, 1, \frac{3}{2}, 2, \frac{5}{2}, 3\}$. The six values $\{0, 1, \frac{3}{2}, 2, \frac{5}{2}, 3\}$ carry status U; the value $\gamma = -1$ comes only from the linear saturation model of Section 11.7.2 and is tagged C. Consequently the three-point sample list $\{1, \frac{5}{2}, 3\}$ recorded in Section 12.3 is retracted.*

Proof. By Theorem 12.22 the anchors listed there are in bijection with the seven values of μ_{ratio} fixed by the internal dictionaries of Sections 8.1, BBB, YY and MMM. The values $\gamma = 2$ (Mason, $\kappa = 1/6$), $\gamma = \frac{3}{2}$ (level $X(p^n)$, $\kappa = 5/24$), $\gamma = 0$ (modular tower, BMY equality) and $\gamma = -1$ (saturation) were omitted in Section 12.3; no new theorem supplies them, they are direct substitutions. The value $\gamma = 2$ is the critical value of the synchronous insertion weight $\lambda = 2$, which is why its omission mattered. □

Theorem 12.24 (Gap-dilution identity and the mu-unit criterion)**.** *In the Qiu insertion $J_C \sim A \times B$ with $\kappa_A = 5/12$, $R_A = 4$, $R_B = 3-\gamma$ and $\lambda = L_A/L_B = 2/(D^2-1)$, one has $\kappa(J_C) = \frac{5\lambda+4-\gamma}{12(1+\lambda)}$ and the exact identities $R_J - 3 = \frac{\lambda-\gamma}{1+\lambda}$ and $\mu_J - \frac{8}{3} = -\frac{2(\gamma-\lambda)}{3(1+\lambda)}$. Hence $\kappa(J_C) \leq \frac{1}{3} \iff \gamma \geq \lambda$, and equivalently $\mu_{\mathrm{ratio}}(B) \leq \frac{8}{3} - \frac{2\lambda}{3}$, with $R_B \leq 3 - \lambda$.*

Proof. The mass-weighted mean gives $\kappa(J_C) = (\lambda\kappa_A + \kappa_B)/(1+\lambda)$ with $\kappa_B = (4-\gamma)/12$, which is the displayed formula. Using $R = 12\kappa - 1$ in the tower limit, $R_J = (4\lambda + 3 - \gamma)/(1+\lambda) = 3 + (\lambda - \gamma)/(1+\lambda)$; multiplying by $2/3$ after $\mu = \frac{2}{3}(1+R)$ gives the μ_J identity. The inequality $\kappa(J_C) \leq 1/3$ is $5\lambda+4-\gamma \leq 4+4\lambda$, i.e. $\lambda \leq \gamma$. Substituting $\gamma = 4 - \frac{3}{2}\mu_B$ gives $\mu_B \leq (8-2\lambda)/3$. Nine exact rational checks: $\gamma = \lambda$ at $D = 2, 3, 5, 100$ gives $\kappa = 1/3$, $R = 3$, $\mu = 8/3$; $(\lambda, \gamma) = (2, 1)$ gives $\kappa = 13/36$, $R = 10/3$, $\mu = 26/9$; $(2, 2)$ gives exact equality; $(2/3, 3)$ gives $R_J = 8/5$, $\kappa = 13/60$; $(2/3, 0)$ gives $R_J = 17/5$, $\kappa = 11/30$. □

Theorem 12.25 (Weight spectrum and closure phase diagram). *For integral $D \geq 2$, $\lambda(D) = 2/(D^2-1)$ is decreasing with $\lambda_{\max} = 2/3$ at $D = 2$; the synchronous tower gives $\lambda \equiv 2$. Therefore: (a) on the legal fixed-A dilution tower, every proved anchor with $\gamma \geq 1$ and the whole band of Theorem 12.26 satisfy $\gamma \geq \lambda$; (b) only $\gamma = 0$ and $\gamma = -1$ fail, and both require a seed with $R_A = 4$; (c) on the synchronous tower only $\gamma \in \{2, \frac{5}{2}, 3\}$ close, with $\gamma = 2$ attaining $R_J = 3$ exactly, but that tower violates axiom (T3) by Section 11.11.2, so the equality is a curve-level identity, not tower evidence.*

Proof. Monotonicity of $2/(D^2-1)$ is immediate; the ten exact values 2/3, 1/4, 2/15, 1/12, 2/35, 2/63, 2/99, 2/143, 2/399, 2/9999 were verified for $D = 2$, 3, 4, 5, 6, 8, 10, 12, 20, 100. Combining with Theorem 12.24 gives the 10×7 closure matrix: the rows $\gamma = 1, \frac{3}{2}, 2, \frac{5}{2}, 3$ close on all six displayed weights, the rows $\gamma = 0, -1$ close on none, and the synchronous column closes exactly for $\gamma \geq 2$. The value $\kappa = 11/30$ at $D = 2$, $\gamma = 0$ is below the optimal closure 5/12 of Section 11.7.2, so no contradiction arises. The (T3) violation of the synchronous tower is verbatim from Section 11.11.2. □

Theorem 12.26 (Finite-layer gap and a proved countable band). *With $\hat{f}_p = (2n + \frac{2n}{p^n-1} - \frac{2}{p-1}) \log p$ from Section 6.4.2 and $w_p = \hat{f}_p/\hat{f}_{\text{tot}}$, $\gamma(n,m) = 4 - [(1+q^{-m})w_p + (1+p^{-n})w_q] + O_{p,q}(1/\hat{f})$. Consequently the joint limit gives $\gamma \to 3$, the one-sided limit $n \to \infty$ with m fixed gives $\gamma \to 3 - q^{-m}$, and for every prime $q \geq 3$ the band satisfies $\gamma \geq 8/3$. Hence the proved interval is $\{-1, 0, 1, \frac{3}{2}, 2, \frac{5}{2}, 3\} \cup [8/3, 3)$, the second part being a countable band of status U relative to Section 7.5.*

Proof. Section 7.5.4 gives $\mu_{\text{ratio}} = \frac{2}{3}[(1+q^{-m})w_p + (1+p^{-n})w_q] + O(1/\hat{f})$; applying $\gamma = 4 - \frac{3}{2}\mu_{\text{ratio}}$ from Theorem 12.22 yields the displayed formula, which is consistent with the one-sided $\kappa \to (1+q^{-m})/12$ of Section 7.5. For $p = 3$, $q = 5$ and $m = 1, 2, 4, 8$ the limits are $2.8, 2.96, 2.9984, 2.99999744$, and the finite-layer values at $n = 100$ differ by 2.9×10^{-3}, 1.05×10^{-3}, 8.39×10^{-5}, 2.62×10^{-7} respectively, monotonically decreasing; the joint limit at $(20, 20)$ is 3.000000000. Since $q^{-m} \leq 1/3$, the band lies in $[8/3, 3)$, and $8/3 > 2/3$ and $8/3 > 2$ exceed both the maximal legal and the synchronous insertion weight, so the band closes the barrier for every weight without invoking $\gamma > 0$. □

Theorem 12.27 (Calibration is a category error; final reduction). *$\Delta\kappa = 1/12$ corresponds to $\gamma = 1$, $R = 2$, $\kappa = 1/4$, i.e. the proved automorphic GL_2 anchor $X(N)$. It is not a sample of a Mumford–Tate corank gap for a generic $\mathrm{End}^0 = \mathbb{Q}$ Hodge factor. Calibrating γ(corank) on it underestimates the unique proved generic-class sample $\gamma = 5/2$ of the CM–Parseval bound by a factor 5/2. Finally, (II) is equivalent to the existence of an absolutely simple non-GL_2 factor with $\gamma < 0$.*

Proof. Substitute $\Delta\kappa = 1/12$ into $\gamma = 4 - 12\kappa$: $\gamma = 1$, and Section 9.1.2 gives $R = 2$, whose anchor is the $X(N)$ twist tower, already proved and squarely in the automorphic row. By contrast the CM–Parseval bound of Section 8.1 gives $\mu_{\text{ratio}} \leq 1$, hence $\gamma \geq 5/2$; the ratio of the two gaps is 5/2, and in κ units $\Delta\kappa_{\text{CM}} = 5/24$ against 1/12, again 5/2. This is the fourth separation of the same type in this volume, after Sections 9.1.1, 10.13 and the Chern-number comparison in Section 12.3. For the final claim, $\gamma < 0 \iff R > 3$ by definition, and Section 11.11.3 reduces a barrier counterexample to a single factor with $R_A > 3$; Section 11.8 restricts its class to non-PEL Hodge-type special subvarieties intersected with the Torelli locus and $\mathrm{End}^0 = \mathbb{Q}$. The hypothesis $\gamma > 0$ remains unproved. □

Remark 12.28 (Ledger, verbatim checks and honest boundary). No external source was consulted in this appendix; all inputs are the internal dictionaries of Sections 9.1, ZZ, YY, RR, PPP, QQQ, TTT and MMM, together with exact rational arithmetic and finite-layer numerics. Seven dictionary anchors, nine dilution identities, six mu-unit thresholds, ten weight values, a 10×7 closure matrix and twenty finite-layer values were recomputed. The quantity $\gamma > 0$ is an explicit

hypothesis, not a theorem; the normalisation of $\widehat{c}_2$ by analytic torsion, the Chern–Weil form of $\mathcal{E}_\infty$ and the enumerability of PHL–Torelli intersections remain tagged C and are absorbed in the existing Problem 10.12 item. The synchronous-tower equality is limited to curve level by (T3). No new No-Go statement, conditional input or external mathematical assumption is introduced; the fourteen No-Go statements, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$ and abc (still O) are unchanged; nothing here independently verifies Theorem A

12.5 Explicit Endomorphism and Mumford–Tate Determination of the Narrow Gap

This appendix performs the explicit determination proposed at the end of Section 12.4: the route-B narrow gap — an absolutely simple abelian variety with trivial endomorphism algebra but a properly small Mumford–Tate group, lying inside the Torelli locus — is identified, classified, and shown to collapse to a single dimension. All external results are cited in prose and used as text-level inputs; none is re-proved.

Theorem 12.29 (Criterion for the narrow gap)**.** *Let $A/\mathbb{C}$ be a polarised abelian variety of dimension g, let $V = H^1(A, \mathbb{Q})$, and let* $\mathrm{MT}(A)$ *be its Mumford–Tate group. Then*

$$\mathrm{MT}(A) = \mathbf{G}_m \cdot \mathrm{MT}^{\mathrm{der}}(A), \qquad \mathrm{End}^0(A) = \mathrm{End}_{\mathrm{MT}^{\mathrm{der}}(A)}(V).$$

Consequently $\mathrm{End}^0(A) = \mathbb{Q}$ *if and only if V is absolutely irreducible as a* $\mathrm{MT}^{\mathrm{der}}(A)$*-module. By Faltings' theorem the ℓ-adic Lie algebra and the Hodge Lie algebra have the same commutator algebra, and by Serre's open image theorem, if* $\mathrm{End}(A) \cong \mathbb{Z}$ *and $g \in \{2, 6\} \cup (2\mathbb{Z} + 1)$ then the ℓ-adic representation is surjective for all sufficiently large ℓ; hence a properly small Mumford–Tate group with trivial endomorphism algebra can occur only for even $g \notin \{2, 6\}$, subject to the Mumford–Tate conjecture.*

Proof. Weight one forces the homothety $\mathbf{G}_m$ inside $\mathrm{MT}(A)$; the commutant of $\mathrm{MT}(A)$ equals the commutant of its derived group because the centre acts by scalars. Deligne's Tannakian description of Hodge cycles on abelian varieties identifies $\mathrm{End}^0(A)$ with that commutant, and Faltings' theorem gives the equality of the ℓ-adic and Hodge commutants. Schur's lemma gives the irreducibility equivalence. The dimension statement is the quoted form of Serre's open image theorem as recorded in the 2026 preprint of Bouchet–Hanselman–Pieper–Schiavone; the Mumford–Tate conjecture for general Mumford-type fourfolds is not known, so the last implication is conditional on that conjecture. □

Theorem 12.30 (Representation-theoretic classification and two retractions)**.** *In Mumford's construction let F be totally real of degree d, let B/F be a quaternion algebra split at exactly one real place, and let $V \simeq 2^{\otimes d}$ be the corestriction representation of $G'_{\mathbb{R}} \simeq \mathrm{SL}_{2,\mathbb{R}} \times \mathrm{SU}_2(\mathbb{R})^{d-1}$. Then the tensor product carries an invariant form of sign $(-1)^d$: it is symplectic for odd d, and* $\dim \mathrm{End}_{\mathrm{SL}_2^d}(V) = 1$ *for every d. With $m = d$ for odd $d > 1$ and $m = d + 1$ otherwise, the corresponding abelian variety has dimension $g = 2^{m-1}$, and*

$$\mathrm{End}^0(X_\eta) = \mathbb{Q} \quad (m = d), \qquad \mathrm{End}^0(X_\eta) = \textit{a totally definite quaternion algebra} \quad (m = d + 1).$$

For $d = 1, 2$ the two Mumford-type Shimura curves are of PEL type. Hence the label "G_2 type" used in the external route-B note is a misnomer: the derived group is SL_2^d (type A_1^d), not the exceptional group G_2; and the assertion in Section 12.3 that no first-hand source for Mumford's four-dimensional family could be verified is retracted.

Proof. The sign of a tensor product of d symplectic two-dimensional forms is $(-1)^d$, since each factor contributes -1; for odd d the representation lands in Sp_{2^d}, giving $m = d$ and $g = 2^{d-1}$. For even d the form is orthogonal and one further symplectic factor is required, giving $m = d + 1$

and $g = 2^d$; this is exactly the dichotomy between $\mathrm{Mat}_{2^d}(\mathbb{Q})$ and $\mathrm{Mat}_{2^d}(L)$ with L an imaginary quadratic field. Each SL_2 factor acts irreducibly on its standard representation, so Schur's lemma gives $\mathrm{End}_{\mathrm{SL}_2^d}(2^{\otimes d}) = \bigotimes_d \mathbb{Q} = \mathbb{Q}$, of dimension one. The two cases and the $d = 1, 2$ PEL exceptions are quoted verbatim from Section 2.2 of Lu–Zuo (arXiv:1405.4751), which generalises Mumford's 1969 example; Galluzzi (2000) records $\mathrm{MT}(X)$ isogenous to $\mathrm{SL}(2)^3$ and strictly contained in $\mathrm{Sp}(8)$; Noot records $\mathrm{Lie}(G)_{\overline{\mathbb{Q}}} \cong \mathfrak{c} \oplus (\mathfrak{sl}_2)^3$ and $\mathrm{End}(X_\eta) = \mathbb{Z}$. Eleven exact integer triples (d, m, g) were checked. No G_2 factor occurs in this classification, and the seven-dimensional odd case $g = 7$ is excluded by Theorem 12.29's Serre condition. □

Theorem 12.31 (Collapse to a single dimension)**.** *Condition on the Serre dimension restriction of Theorem 12.29 and on Theorem A of Lu–Zuo (2014), which states that for $g > 11$ there is no Shimura curve of type I or II contained generically in the Torelli locus $\mathcal{T}_g$. Then the only one-dimensional special locus in the Torelli locus whose generic fibre is non-PEL, has* $\mathrm{End}^0 = \mathbb{Q}$ *and a properly small Mumford–Tate group, is the $g = 4$ Mumford-type family:*

$$d = 3 \implies g = 2^{d-1} = 4; \qquad d = 5, 7, 9, 11 \implies g = 16, 64, 256, 1024 > 11 \text{ excluded.}$$

For this survivor $\dim \mathcal{A}_4 = 10$, $\dim \mathcal{T}_4 = 3 \cdot 4 - 3 = 9$, $\dim S = 1$, $\mathrm{codim}_{\mathcal{A}_4} S = 9$ *and* $\mathrm{codim}_{\mathcal{T}_4} S = 8$.

Proof. By Theorem 12.30 the non-PEL narrow-gap fibres occur exactly in case (1), with $d \geq 3$ odd and $g = 2^{d-1} \in \{4, 16, 64, 256, 1024, \dots\}$. All these values are even and different from 2 and 6, so the Serre restriction removes none of them; Lu–Zuo's bound $g > 11$ removes every value except $g = 4$. The dimension counts are the standard $\dim \mathcal{A}_g = g(g+1)/2$ and $\dim \mathcal{M}_g = 3g - 3$, evaluated at $g = 4$; the Torelli morphism is injective. Moonen (2022) records that the Coleman–Oort conjecture must be taken with $g \geq 8$ because counterexamples exist for all smaller genera, in particular $g = 4$, which is consistent with the survivor. The statement is restricted to one-dimensional Shimura curves of Mumford type: positive-dimensional special subvarieties of higher dimension for $g \geq 8$ remain open and are tagged C. □

Theorem 12.32 (Explicit curve realisation and partial closure of L3)**.** *The $g = 4$ survivor is not merely an existence statement. Bouchet, Hanselman, Pieper and Schiavone (arXiv:2510.00093, 2026) give explicit equations for two Shimura families $C_{7,t}$ and $C_{9,t}$ of genus-4 curves whose Jacobians are abelian fourfolds of Mumford type; their Theorem 4.15 proves these Shimura families are of Mumford type, and their abstract records them as the first explicit examples over $\mathbb{Q}$ with geometric endomorphism algebra $\mathbb{Z}$ and additional Hodge tensors. Accordingly the "no explicit curve" part of the L3 gap recorded in Sections 11.10 and QQQ is closed; L3 now consists solely of an explicit* tower*, a positive limit weight, and the conductor–height data.*

Proof. The quoted statements are verbatim from the abstract and from Theorem 1.2 / Theorem 4.15 of the cited preprint. A Jacobian of a genus-4 curve is a point of $\mathcal{T}_4 \subset \mathcal{A}_4$, so the Shimura families lie inside the Torelli locus; generic $\mathrm{End}^0 = \mathbb{Q}$ makes the Jacobian an absolutely simple factor by itself. The preprint also records that degenerate fibres with extra automorphisms lie in zero-dimensional special subvarieties and therefore have complex multiplication, which is compatible with the generic trivial endomorphism algebra. No conductor, Faltings height or reduction data appear in the cited text, so the arithmetic part of L3 survives unchanged. □

Theorem 12.33 (Tower axiom failure and absence of convex dilution)**.** *Let A be a Mumford-type abelian fourfold. By Noot (Crelle 519, 2000; EMS, 2001), after a finite extension A has potentially good reduction at every non-archimedean place and the associated p-adic representations factor through $\widetilde{G}_{\overline{\mathbb{Q}}} \cong \mathbf{G}_m \times \mathrm{SL}_2^3$. Hence* $\log \mathrm{cond}(A)$ *is a constant independent of any congruence level, the spectrum-conservation bound of Section 11.11.3 gives $\kappa_n \leq \max_\rho \kappa_\rho = \kappa_A$ along every level tower, and the tower axiom (T3), $\hat{f}_n \to \infty$, fails. Moreover, for the curves of Theorem 12.32 the Jacobian* is *the single absolutely simple factor, so $w_{\mathrm{o}} = 1$, the Qiu insertion weight $\lambda = L_A/L_B$ has no complement B, and the criterion $\gamma \geq \lambda$ of Section 12.4 does not apply.*

Proof. Good reduction at all finite places after a finite extension bounds the conductor by a constant determined by the Shimura datum; a congruence subgroup changes only the level structure, not the isogeny class of A, so $\hat{f}_n$ is constant. By Section 11.11.3 the curvature ratio is a conductor-mass weighted average of the factor ratios and therefore cannot exceed the maximum, which is fixed. This is the second independent failure of (T3), after the synchronous-insertion tower of Section 11.11.2, and it now has a structural reason rather than a missing construction. Writing $J(C) \sim A$ with no complementary factor gives $w_{\mathrm{o}} = 1$ and $\lambda \to \infty$; the finite-D insertion weights $\lambda(D) = 2/(D^2 - 1) \leq 2/3$ of Section 11.10.2 are therefore not available as a diluting mechanism. □

Theorem 12.34 (Restatement of the remaining open problem)**.** *For the explicit $g = 4$ realisation, $\kappa(J_C) = \kappa_A$ and $R_J = R_A$, so with $\gamma_A = 3 - R_A = 4 - 12\kappa_A$ the barrier $\kappa \leq \frac{1}{3}$ is equivalent to $\gamma_A \geq 0$. Exact rational arithmetic gives*

$$\gamma_A = -1, 0, 1, 2, \tfrac{5}{2}, 3 \iff \kappa_A = \tfrac{5}{12}, \tfrac{1}{3}, \tfrac{1}{4}, \tfrac{1}{6}, \tfrac{1}{8}, \tfrac{1}{12},$$

the first value being precisely the linear saturation model of Section 11.7.2. Consequently problem (II) is equivalent to deciding the sign of one number, γ_A, for the explicit Mumford-type Jacobians; no value of γ_A is supplied by the literature consulted here.

Proof. Substitute $w_{\mathrm{o}} = 1$ into the convex-hull identity of Section 3.3.1 and the dictionary $\gamma = 4 - 12\kappa$ of Section 12.4.1. The six rational values follow from $\kappa = (4 - \gamma)/12$; the value $\gamma = -1$ gives $R_A = 4$ and $\kappa_A = 5/12$, matching the saturation model. Sections 12.4 and YY give a proved continuous band $[8/3, 3)$ for mixed Kummer towers, but those are a different class and carry no information about γ_A. □

Remark 12.35 (Status and next tests). Theorem 12.31 is conditional on the Serre dimension restriction and on Lu–Zuo's Theorem A for one-dimensional Shimura curves; higher-dimensional special subvarieties for $g \geq 8$, the Mumford–Tate conjecture for general Mumford-type fourfolds, and the value of γ_A remain open and are tagged C. All twelve external sources were verified at text level only; the 2026 preprint of Bouchet et al. is not peer reviewed. Three finite tests remain: compute $\log \mathrm{cond}$ and the Faltings height for a rational point of $C_{7,t}$ or $C_{9,t}$; compute the Chevalley–Weil Prym spectrum of a Galois cover of such a curve, following Sections 4.1 and EEE; and verify that no level tower restores (T3). None of this changes the fourteen No-Go reasons, the pointwise conclusions, or the status O of *abc*; Theorem A of the parent paper is not independently verified.

12.6 Arithmetic Invariants of the Explicit Mumford-Type Family $C_{7,t}$

This appendix carries out the first arithmetic computation proposed at the end of Section 12.5: take a rational point of the explicit genus-four family $C_{7,t}$ of Bouchet–Hanselman–Pieper–Schiavone (arXiv:2510.00093v2) and evaluate the conductor and Faltings height of its Jacobian, with the aim of obtaining the first non-anchor value of $\gamma_A = 3 - R_A$. The point value is *not* obtained, but the obstruction is structural rather than technical: the conductor is bounded explicitly, the tower axioms are shown to be mutually exclusive on the horizontal family, and the normalisation ambiguity of γ_A is quantified. All external results are cited in prose and used as text-level inputs; none is re-proved.

Theorem 12.36 (Exact discriminant and bad-prime locus)**.** *For $t \in \mathbb{Q} \setminus \{0, 1\}$ let $C_{7,t}$ be the hyperelliptic curve $y^2 = f_t(x)$ with*

$$f_t(x) = t\Big[\big(t - \tfrac{27}{16}\big)x^{10} - \tfrac{567}{64}x^9 - \tfrac{189}{4}tx^8 - \big(84t^2 + \tfrac{189}{4}t\big)x^7 - 189t^2x^6 - \tfrac{189}{2}t^2x^5 + 84t^3x^4 + 108t^3x^3 - 28t^4x\Big].$$

Then f_t has degree $10 = 2g + 2$ with $g = 4$ and a simple root at $x = 0$, and

$$\Delta(C_{7,t}) = 2^{16}\,\mathrm{disc}(f_t) = 2^{36}\,3^{36}\,7^{10}\,t^{54}(t - 1)^{12}.$$

Consequently the bad primes are contained in $\{2,3,7\} \cup \operatorname{supp}(t) \cup \operatorname{supp}(t-1)$.

Proof. The coefficient of x^{10} is $t(t-27/16)$ and the constant coefficient is zero, so $f_t = x \cdot h_t$ with $h_t(0) = -28t^5 \neq 0$; hence $x = 0$ is a simple root, $\deg f_t = 10$, and the smooth projective model has genus $(10-2)/2 = 4$. The discriminant identity was verified by exact rational arithmetic for thirteen values $t \in \{2, 3, 4, 7, 8, 9, 3/2, 7/4, -1, 1/2, 9/8, 49/48, 64/63\}$: writing $\operatorname{disc}(f_t)$ as an exact rational number, the quotient $2^{16} \operatorname{disc}(f_t)/\big(2^{36}3^{36}7^{10}t^{54}(t-1)^{12}\big)$ equals 1 in every case, with zero residual numerator. The bad-prime statement follows because a prime of good reduction cannot divide the discriminant of any integral Weierstrass model, and the constant prefactor is supported on $\{2,3,7\}$. □

Theorem 12.37 (Wild locus and conductor bound)**.** *Let* $A_t = \operatorname{Jac}(C_{7,t})$. *The semistable reductions of* $C_{7,t}$ *at* $t = 0, 1, \infty$ *are obtained after totally ramified extensions of degrees* 4, 7, 3 *respectively (Lemmata 4.8, 4.9, 4.10 of the cited preprint). Hence wild inertia can occur only at* $p \in \{2,3,7\}$, *and for every prime* $p \notin \{2,3,7\}$ *dividing* $t(t-1)$ *one has* $\operatorname{Sw}_p = 0$ *and*

$$f_p(A_t) = 2g - \dim V_\ell(A_t)^{I_p} \leq 2g = 8.$$

Therefore

$$\log \operatorname{cond}(A_t) \leq 8 \log \operatorname{rad}\big(t(t-1)\big) + C_{237}, \qquad \log \operatorname{cond}(A_t) \leq \log |\Delta(C_{7,t})|,$$

with C_{237} *independent of* t. *Along* $t \to \infty$ *the two upper bounds have slopes* 16 *and* 66 *in* $\log|t|$, *so the discriminant bound is looser by the exact factor* $66/16 = 33/8$.

Proof. A finite extension of degree e is tame at residue characteristic p whenever $p \nmid e$; the degrees $4, 7, 3$ have prime divisors $\{2\}, \{7\}, \{3\}$, so outside $\{2,3,7\}$ the inertia action is tame and the Swan conductor vanishes. The Artin conductor exponent of the ℓ-adic representation $V_\ell(A_t)$ (of dimension $2g$) equals $2g - \dim V_\ell^{I_p} + \operatorname{Sw}_p$; with tame inertia the last term is zero and $\dim V_\ell^{I_p} \geq 0$, giving $f_p \leq 8$. Summing $f_p \log p$ over $p \mid t(t-1)$ gives at most $8 \log \operatorname{rad}(t(t-1))$, and the finitely many primes $2, 3, 7$ contribute a t-independent constant. The bound $\log \operatorname{cond} \leq \log|\Delta|$ is the standard Ogg–Saito–Liu inequality $f_p \leq v_p(\Delta_{\min})$. Since $\log|\Delta| = 54\log|t| + 12\log|t-1| + O(1) = 66\log|t| + O(1)$ while $\log \operatorname{rad}(t(t-1)) \leq \log|t| + \log|t-1| = 2\log|t| + O(1)$, the slope ratio is $66/(8 \cdot 2) = 33/8$. □

Theorem 12.38 (Tower axioms and S-unit finiteness: retraction of 12.32)**.** *Assertion 12.32, that the horizontal family violates (T3) because a single Mumford-type seed has constant* log cond, *holds for a fixed seed but not for the family. If the bad set* S *is allowed to grow with* t, *then* $\log \operatorname{cond}(A_t) \to \infty$ *and (T3) holds. If (T2) is imposed, i.e.* S *is fixed, then* t *and* $1-t$ *are both* S*-units solving* $t + (1-t) = 1$, *and by the Evertse finiteness theorem there are only finitely many such* t; *hence no infinite tower exists. Consequently*

(T2) and *(T3) cannot hold simultaneously on the horizontal family.*

An exhaustive search of the exponent window $[-9, 9]$ *for* $S = \{2,3,7\}$ *produces* 75 *explicit parameters, all with bad set exactly* $\{2,3,7\}$*; this is a lower bound on the finite solution set, not a completeness claim.*

Proof. For fixed S, the condition $\operatorname{supp}(t(t-1)) \subseteq S$ makes t and $1-t$ elements of the finitely generated group of S-units. The S-unit equation $x + y = 1$ has finitely many solutions by Evertse's theorem; the explicit list was produced by enumerating $t = \pm 2^a 3^b 7^c$ with $a, b, c \in [-9, 9]$ and testing both t and $t-1$ for S-support, using exact integer arithmetic. Each such t has log cond bounded by the t-independent constant C_{237} of Theorem 12.37, so $\hat{f} = \log \operatorname{cond}/g$ is bounded and (T3) fails. Conversely, allowing new primes in $t(t-1)$ adds at most $8 \log p$ each, so choosing t with unboundedly many prime divisors makes $\hat{f} \to \infty$ and (T3) holds, at the cost of (T2). The statement that fixed genus $g = 4$ violates (T1) is unchanged and gives a third, independent exclusion. □

Theorem 12.39 (Normalisation ambiguity: no point value for γ_A)**.** *Let $\kappa = h_F/(\log \mathrm{cond})$ be the ratio used in this volume, and let h_F be measured in any of the three normalisations whose mutual differences are $c \cdot g$ with $c = 0.9189385332$ (Appendices DD and LL). On a single curve, unlike in a tower, the quotient $c \cdot g/\hat{f}\, g$ does not tend to zero, and*

$$\Delta\kappa = \frac{c\, g}{\log \mathrm{cond}}, \qquad \Delta\gamma_A = 12\,\Delta\kappa = \frac{48c}{\log \mathrm{cond}} = \frac{44.10905}{\log \mathrm{cond}}.$$

Since $\log \mathrm{cond}$ *is bounded on every (T2)-admissible family, the point value of γ_A carries an irreducible $O(1/\log \mathrm{cond})$ ambiguity. Fixing it requires two independent anchors as in Section 5.6; the two LL anchors belong to the Fermat tower and cannot be transported, and Table 1 of the cited preprint lists CM fields, discriminants and fields of definition (degree 3 and higher) but no rational t-values, so its CM points cannot supply a second anchor.*

Proof. By Appendices DD and LL the three standard Faltings-height normalisations differ by an additive $c \cdot g$ with $c = \frac{1}{2}\log(2\pi)$, the numerical value 0.9189385332 being reproduced by the two anchors of 00712. In a tower $\kappa = (h_F/g)/\hat{f}$ absorbs this as $c/\hat{f} \to 0$ because $\hat{f} \to \infty$. Here $\hat{f} = \log \mathrm{cond}\,/g$ is a finite positive number, so the shift in κ is $c\,g/\log \mathrm{cond}$, and by the dictionary $\gamma = 4 - 12\kappa$ of Section 12.4 the shift in γ_A is twelve times that. Table 1 of arXiv:2510.00093v2 records LMFDB labels, $\mathrm{disc}(K) = 7^4 q^4$ and fields such as 3.3.49.1, 6.2.31213.1, 9.3.164590951.1; no parameter t is tabulated and the fields of definition are not $\mathbb{Q}$, so no Colmez closed-form height can be attached to a listed point without additional data. □

Theorem 12.40 (The family is a Viehweg–Zuo equality family)**.** *For the two families of arXiv:2510.00093v2, the local contributions at $t = 0, 1, \infty$ are -6, $-9/7$, $+22/3$ and -6, $-29/9$, $+28/3$, giving*

$$\deg\big(e^*\omega_{A_7/X_7}\big) = \tfrac{1}{2}\Big(-6 - \tfrac{9}{7} + \tfrac{22}{3}\Big) = \tfrac{1}{42}, \qquad \deg\Omega^1_{X_7} = \tfrac{1}{84},$$

$$\deg\big(e^*\omega_{A_9/X_9}\big) = \tfrac{1}{2}\Big(-6 - \tfrac{29}{9} + \tfrac{28}{3}\Big) = \tfrac{1}{18}, \qquad \frac{1/42}{1/84} = \frac{1/18}{1/36} = 2 = \frac{g}{2} \quad (g = 4).$$

Hence the Arakelov inequality is attained with equality and the families are Shimura families; equivalently, $C_{7,t}$ is an equality-track family on the function-field side, which is precisely the geometric counterpart of the route-B equality orbit of Section 12.3. It is therefore a natural test object for (II), with the missing ingredient being purely arithmetic.

Proof. The three local contributions are quoted verbatim from Lemmata 4.8–4.10 and 4.11–4.13 of the preprint; the two displayed evaluations are exact rational computations. Proposition 4.14 of the preprint states that the extended families satisfy the Arakelov inequality with equality, whence they are Shimura families by the Viehweg–Zuo criterion. The ratio $\deg(e^*\omega)/\deg\Omega^1 = 2 = g/2$ is the equality case $\deg\omega = \frac{g}{2}\deg\Omega^1(\log S)$ for a Shimura curve of type I/II, not the surface inequality $c_1^2 = 3c_2$; the two are distinct objects, as noted in Section 12.3. □

Theorem 12.41 (Decidable interval for γ_A)**.** *Under the dictionary of Section 12.4 ($\sigma \to 1$, $\varepsilon_{\mathrm{ar}} \to 0$),*

$$\gamma_A \geq 0 \iff \kappa_A \leq \tfrac{1}{3} \iff h_F(A_t) \leq \tfrac{1}{3}\log \mathrm{cond}(A_t).$$

Combining with $\log \mathrm{cond} \leq \log|\Delta|$ *gives the sufficient condition $h_F > \frac{1}{3}\log|\Delta|$ for failure of the barrier. For $t \in \{2, 8, 3/2, 1/2, 101/100\}$ the thresholds $\frac{1}{3}\log|\Delta|$ are 40.4641, 73.2011, 32.5133, 12.7382, 9.7459, and the corresponding normalisation bands $\Delta\gamma_A$ are 0.3634, 0.2009, 0.4523, 1.1542, 1.5086. Conversely, the Wilms lower bound $\delta_{\mathrm{Fal}} > -2g\log(2\pi^4)$ permits $h_F < 0$, in which case $\kappa_A < 0$, $\gamma_A > 4$, and the barrier holds trivially.*

Proof. The equivalence is the UUU dictionary $\gamma = 4 - 12\kappa$ together with $\kappa = h_F/\log \mathrm{cond}$. The numerical thresholds follow by exact evaluation of $\log|\Delta(C_{7,t})|$ from Theorem 12.36, and the bands from Theorem 12.39. The Wilms bound is quoted as the constant $L_2 = 2\log(2\pi^4) = 10.5441$ of this volume; since h_F may be negative for a CM or small-height abelian fourfold, the trivial-barrier case is consistent with all established inequalities. No numerical value of $h_F(A_t)$ is claimed. □

Remark 12.42 (What remains to be computed). Three external ingredients are required to turn the interval of Theorem 12.41 into a point value: (i) a rational (or explicit algebraic) parameter t for one of the CM points of Table 1, to serve as a second height anchor; (ii) the conductor exponents f_2, f_3, f_7, which require a regular model or cluster picture at the three wild primes; (iii) the absolute Faltings height, requiring a numerical period matrix and theta constants with a fixed normalisation. The ten branch points of $C_{7,2}$ consist of two real roots (0 and $37.4528735151\ldots$) and four complex conjugate pairs, so the standard all-real cut construction does not apply. None of Magma, Sage or Pari/GP is available in the present environment, and no general algorithm for genus-four conductors is implemented in them. All statements above are therefore unconditional relative to the quoted inputs, and no proof of any cited theorem is reproduced here. The fourteen No-Go statements, the constants $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, and the status $abc =$ 'O' are unchanged

12.7 CM Anchor Parameters and the Point-Value Obstruction for γ_A

This appendix executes the unique data step proposed at the end of Section 12.6: recover the parameter values t of the putative CM points of the explicit Mumford-type families from the public repository accompanying Bouchet–Hanselman–Pieper–Schiavone (arXiv:2510.00093), and use them as a second height anchor to fix the normalisation constant in $\gamma_A = 3 - R_A$. The data are obtained in exact form, but the point value is *not*: every CM parameter is non-rational, and passing to its field of definition couples the normalisation ambiguity to the base-discriminant term that Section 11.6 identified as structural. All external inputs are used at text level only; none is re-proved.

Theorem 12.43 (The CM parameter data and non-rationality)**.** *The repository* `SamSchiavone/Mumford-type` *contains two files* `Shimura/Shimura_237_param.m` *(38 rows) and* `Shimura/Shimura_239_param.m` *(20 rows), committed 4 August 2025, whose rows have the shape*

$$|\ \textit{CM-field label}\ |\ P_K(x)\ |\ t = (1/D)\,Q(\mu)\ |, \qquad D \in \mathbb{Z}_{>0},\ Q \in \mathbb{Z}[\mu],$$

with P_K the minimal polynomial of μ. For all 58 rows: P_K is irreducible over $\mathbb{Q}$; $\deg P_K \in \{3,6,9,12,15,21\}$ with multiplicities $\{11,8,23,4,10,2\}$; $\deg Q \geq 1$ and $\deg Q < \deg P_K$; and the minimal polynomial of t over $\mathbb{Q}$, obtained by eliminating μ via $\mathrm{Res}_\mu(P_K(\mu), Dy - Q(\mu))$, has degree exactly $\deg P_K$. Consequently

$$\boxed{t \notin \mathbb{Q} \text{ for all } 58 \text{ rows}, \qquad \mathbb{Q}(t) = \mathbb{Q}(\mu) = K, \qquad [K:\mathbb{Q}] \in \{3,6,9,12,15,21\}.}$$

Proof. Each row was parsed and checked by exact symbolic computation. Irreducibility of P_K was verified by complete factorisation over $\mathbb{Q}$ for all 58 rows. Since P_K is monic irreducible of degree d and $0 < \deg Q < d$, the polynomial $Q(x) - c$ is nonzero of degree $< d$ and cannot be divisible by P_K; hence $Q(\mu) \neq c$ for every $c \in \mathbb{Q}$ and $t = Q(\mu)/D \notin \mathbb{Q}$. Eliminating μ gives a polynomial satisfied by t; its factorisation over $\mathbb{Q}$ contains an irreducible factor of degree d in every row, so $[\mathbb{Q}(t):\mathbb{Q}] = d$, and because $\mathbb{Q}(t) \subseteq K$ with equal degrees, $\mathbb{Q}(t) = K$. The paper's Tables 1 and 2 print only field labels and discriminants; the parameter values themselves reside in these two files. □

Theorem 12.44 (Exact norms and bounded height of the parameters)**.** *Let $d = \deg P_K$, let $Q_x(x) = Q(x)$ and let P_K be monic. Then*

$$N_{K/\mathbb{Q}}(t) = \frac{\mathrm{Res}_x(P_K, Q_x)}{D^d}, \qquad N_{K/\mathbb{Q}}(t-1) = \frac{\mathrm{Res}_x(P_K, Q_x - D)}{D^d},$$

both exact rational numbers. Across the 58 rows,

$$\log D \in [10.697,\, 331.375], \quad \log|N(t)| \in [-4.618,\, 12.887], \quad \log|N(t-1)| \in [-60.036,\, 7.705].$$

In particular the enormous denominators are cancelled by the resultants: every t is an algebraic number of height $O(1)$.

Proof. For monic P_K with roots $\mu_1, \dots, \mu_d$ in a splitting field, $\mathrm{Res}(P_K, Q_x) = \prod_{i=1}^{d} Q_x(\mu_i)$, and $t_i = Q_x(\mu_i)/D$ are the conjugates of t; hence $N(t) = \prod_i t_i = \mathrm{Res}(P_K, Q_x)/D^d$. Replacing Q_x by $Q_x - D$ gives the second identity. Both were evaluated as exact integers divided by exact integer powers, with zero floating-point error, for all 58 rows. The numerical ranges are the minima and maxima of the resulting exact logarithms. Since $\log|N(t)|$ stays uniformly bounded while $\log D$ reaches 331.375, the resultant numerator grows at the same exponential rate as D^d; equivalently the Weil height of t is $O(1)$. This explains why the CM parameters could be reconstructed by interpolation in the source paper: they remain in a compact region of the parameter line. □

Theorem 12.45 (Exact discriminant mass for the 237 family)**.** *For a 237-family point defined over K, the discriminant identity of Section 12.6, $\Delta(C_{7,t}) = 2^{36}3^{36}7^{10}t^{54}(t-1)^{12}$, gives*

$$\log N_{K/\mathbb{Q}}\big(\Delta(C_{7,t})\big) = d\, C_{237} + 54\log|N_{K/\mathbb{Q}}(t)| + 12\log|N_{K/\mathbb{Q}}(t-1)|,$$

with $C_{237} = 36\log 2 + 36\log 3 + 10\log 7 = 83.962442\ldots$. For $d = 1$, $t = 2$ this returns $83.962442 + 54\log 2 = 121.3924$, reproducing the value of $\log|\Delta(C_{7,2})|$ in Section 12.6. For the first row (`8.0.78340035470 41.1`), $d = 3$, $D = 601617706932059$, $N(t) = -7031543902913033/2910703085571511$ and $\log N(\Delta) = 253.75125$.

Proof. The norm is multiplicative and $N_{K/\mathbb{Q}}(a) = a^d$ for $a \in \mathbb{Q}$; applying $\log|\cdot|$ to $N(\Delta) = N(2^{36}3^{36}7^{10})\,N(t)^{54}N(t-1)^{12}$ gives the displayed formula. All 38 rows of the 237 file were evaluated exactly. The consistency check uses $N(t) = t$, $N(t-1) = 1$ at $d = 1$, $t = 2$: $C_{237} + 54\log 2 = 83.962442 + 37.429948 = 121.392390$, matching Section 12.6 to the quoted precision. The 239 family has a different discriminant formula and is not covered by this theorem; only its norms are tabulated. □

Theorem 12.46 (The normalisation band does not collapse)**.** *Let $c = \frac{1}{2}\log(2\pi) = 0.9189385332$ and $48c = 44.1090495938$, so that the normalisation ambiguity of γ_A is $\Delta\gamma_A = 48c/\log\mathrm{cond}$ (Section 12.6). For a curve over K the $\mathbb{Q}$-normalised conductor satisfies $\log\mathrm{cond}_{\mathrm{eff}} \le \frac{1}{d}\log N_{K/\mathbb{Q}}(\Delta)$, hence*

$$\Delta\gamma_A \;\ge\; \frac{48c}{\frac{1}{d}\log N_{K/\mathbb{Q}}(\Delta)}.$$

Over the 38 exact rows of the 237 family,

$$\tfrac{1}{d}\log N(\Delta) \in [9.221,\, 250.785], \qquad \Delta\gamma_A^{\min} \in [0.17588,\, 4.78351],$$

and the lower bound already exceeds 1 for 6 of the 38 rows. The rational point $t = 2$ gives 0.3634; most CM anchors therefore do not *improve the band and several make it substantially worse.*

Proof. The conductor of a curve over a number field, $\mathbb{Q}$-normalised by dividing the log norm by d, is bounded by the log norm of any integral discriminant (Section 12.6, Theorem 2). Applying Theorem 12.45 gives the upper bound on $\log\mathrm{cond}_{\mathrm{eff}}$, and inverting gives the lower bound on $\Delta\gamma_A$. All 38 values were computed from exact rationals. The failure of collapse has a structural cause: by Theorem 12.44 the parameters have height $O(1)$, so $\frac{1}{d}\log N(\Delta)$ is dominated by the fixed constant C_{237} plus bounded norm terms and does not grow with the discriminants of the CM fields listed in the paper's tables. □

Theorem 12.47 (Base-field coupling and the cyclic obstruction)**.** *Since $K \neq \mathbb{Q}$, the relative Parshin inequality over K contains the term $(2g-2)\log|D_K| = 6\log|D_K|$ for $g = 4$; this is precisely the structural term B of Section 11.6 ($B = 2/c$) which Appendix P5 proved cannot be absorbed into Δ_{Del}. Using $\mathrm{disc}(P_K) = \mathit{index}^2 \cdot D_K$ as an upper bound for $|D_K|$, the dimensionless ratio*

$$B_K^{\mathrm{up}} := \frac{6d \log|\mathrm{disc}(P_K)|}{\log N_{K/\mathbb{Q}}(\Delta)}$$

lies in $[0.0931, 13.0365]$ *over the 38 rows, exceeding* 1 *in* 27 *of them. Sampling the prime divisors of* $\mathrm{disc}(P_K)$ *at each degree gives* 13, 1399, 211, 3023, 757, 2447, 26261, 176887 *in addition to* $3, 7$*. Hence fixing the normalisation constant c from a CM anchor would require a theoretical value of κ, which over $K \neq \mathbb{Q}$ itself carries B_K: the two ambiguities are coupled, not eliminated.*

Proof. The relative Parshin–Arakelov inequality is the inequality quoted in Section 11.6; its K-dependence is exactly $(2g-2)\log|D_K|$ divided by the $\mathbb{Q}$-normalised conductor, giving $B = 2/c$ on the cyclotomic towers of that appendix. Appendix P5 showed that both possible conventions for absorbing this term yield coefficients $\frac{3+B}{1+B} \to 2$ or $3 + B \to 6$, neither equal to 3; the term is structural. Since $\mathrm{disc}(P_K)$ is an integer multiple of D_K by a square, $\log|D_K| \leq \log|\mathrm{disc}(P_K)|$, so B_K^{up} is an upper bound for the base-field contribution relative to the discriminant mass; its range was computed exactly from Theorems 12.44 and 12.45. The prime-divisor sample indicates that primes outside the fixed wild set $\{2, 3, 7\}$ can enter the relative conductor; because the index may contribute extra primes, this is a magnitude warning and not a statement about the exact prime divisors of D_K. A Colmez-type closed formula for h_F supplies a value of κ only after the CM type and the relative conductor normalisation are fixed, both of which involve D_K; thus the calibration is circular. □

Theorem 12.48 (Verdict)**.** *The 58 CM parameters are disjoint from the* 75 *rational S-unit parameters of Section 12.6 (all of the latter lie in $\mathbb{Q}$, none of the former does), so the mutual exclusivity of the tower axioms $(T2) \wedge (T3)$ on the horizontal family is unchanged. The criterion $\gamma_A \geq 0 \iff h_F \leq \frac{1}{3}\log \mathrm{cond}$ is unchanged, and γ_A remains an interval rather than a point value. Consequently Problem (II) — whether a single absolutely simple non-GL_2 factor with $R_A > 3$ exists — remains open; the data step is closed with a negative structural verdict, not with a missing computation. The fourteen No-Go statements, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$ and the status $abc =$ 'O' are unchanged, and*

Proof. Disjointness is immediate from Theorem 12.43: the 75 parameters of Section 12.6 are rational solutions of an S-unit equation, while every CM parameter generates a field of degree at least 3. The $(T2) \wedge (T3)$ argument of Section 12.6 used only rationality and fixed bad set, so it is unaffected. The γ_A criterion is an identity from $12\kappa = \sigma(1 + R) + \varepsilon_{\mathrm{ar}}$ and does not depend on the anchor. Theorems 12.46 and 12.47 show that no CM anchor fixes the normalisation constant c without a CM type and a relative conductor normalisation, i.e. without new analytic input; therefore γ_A cannot be promoted from an interval to a point value in this round. No new inequality, no new No-Go statement, and □

Remark 12.49 (Honest boundary). Theorems 12.43–12.48 are unconditional relative to: the verbatim contents and commit metadata of the two repository files; exact rational, resultant, discriminant and irreducibility computations; the discriminant identity and conductor bound of Section 12.6; and the structural term identification of Sections 11.6 and P5. The attribution of the 58 rows as *putative* CM points, the theorem chain of the source preprint, and the Colmez closed formula are not re-proved. The absolute Faltings height, the exact relative conductor exponents f_2, f_3, f_7, the CM type Φ of any single point, and the sign of γ_A remain open: no Magma, Sage or Pari/GP is available in the present environment and no general genus-four conductor algorithm is implemented here. B_K^{up} is an upper-bound ratio (it contains the index factor) and $\Delta\gamma_A^{\min}$ is a lower-bound band; the 239 family is not covered by the 237 discriminant formula. One

external data source (a public GitHub repository) is added at text level. This appendix is not an independent verification of Theorem A.

12.8 Recovering CM Types from the Repositories and the Colmez Disjunction

This appendix executes the data step proposed at the end of Appendix 12.7: recover a single CM type Φ (together with its reflex field and type discriminant) from the public repositories accompanying Bouchet–Hanselman–Pieper–Schiavone (arXiv:2510.00093), in order to feed the averaged Colmez formula and obtain a theoretical value of κ, hence of $\gamma_A = 3 - R_A$. The data shape is determined completely, but the point value is *not*: the CM-type database deliberately excludes the Galois group of the Mumford-type family, and, more fundamentally, the Colmez formula applies precisely to the class of factors that Appendix 8.1 already excluded from item (II). All external inputs are used at text level only; none is re-proved. Item (II) remains open.

Theorem 12.50 (Repository attribution and the shape of the CM-type data)**.** *The data described in Appendix 12.7 as belonging to one repository are in fact distributed over two. The repository* `SamSchiavone/Mumford-type` *(henceforth BHPS) has* 32 *entries on its* `main` *branch (recursive tree, not truncated) and contains neither a* `CM/` *directory nor* `CM-fields-deg-8.txt`*; its* `Shimura/Shimura_237_param.m` *and* `Shimura/Shimura_239_param.m` *supply the* 58 *parameter values* t*, its* `Data/mumford_jac.py` *a pool of* 529 *degree-eight CM field labels, and* `NootCMCube.m` *the* 8T24 *condition. The file* `CM-fields-deg-8.txt` *and the directory* `CM/` *with* `CMTypes.m`*,* `reflex.m`*,* `cm-g4.m`*,* `cm-script.m` *and* `data_processing.m` *belong instead to* `Andreas-Pieper/Genus-4-RM-CM` *(*133 *entries), to which the BHPS README points verbatim as the repository containing the CM-curve reconstruction code. Within that directory,* `cm-g4.m` *declares the output format*

$$\textit{label}\,|\,\textit{coeffs}\,|\,\textit{CM type}\,|\,\textit{ideal class rep}\,|\,\textit{polarising element}\,|\,\textit{complex conjugation}\,|\,|\text{Schottky}|,$$

and writes Φ *as a list of four complex embeddings at* 30*-digit precision. The intrinsics* `IsPrimitiveType` *(criterion* $\#G/\#H = 2d = 8$*),* `ReflexType` *and* `IdentifyCMType` *in* `CMTypes.m`*, and* `ReflexField`*,* `TypeNorm` *and* `ReflexTypeNorm` *in* `reflex.m`*, supply exactly the primitive CM type, its reflex field* K^r*, the reflex type and the type discriminant* d_Φ *required by the finite form of the averaged Colmez formula used in Appendices 5.4–5.8. Hence the* inputs *of the Colmez formula are aligned item by item with the code, but* Φ *is generated by a Magma computation from a CM field, its Galois group and the primitivity test; the repository stores no ready-made numerical CM-type table for the BHPS points (the older* `cm-fields-jacobians.txt` *records only an* `index in list of taus`*).*

Proof. Both repository trees were read through the public GitHub tree API with `recursive=1`; the BHPS tree returned `truncated=false` and contains exactly the files listed above. The BHPS README states verbatim that the code used to reconstruct the Shimura curves from CM points is not in that repository and points to three external repositories, the first being `Andreas-Pieper/Genus-4-RM-CM`. The format string, the precision reduction `ChangeUniverse(Phi, ComplexField(30))`, the substitution of the imaginary unit by `I`, the criterion $\#G/\#H = 2d$, the assertion $\#\Phi_{K^r} = \deg(K^r)/2$ in `ReflexField`, and the seven-field header of `data_processing.m` were read verbatim from the downloaded files. The Colmez finite form of this audit is $\mu(\Phi) = \sum_\chi m(\Phi,\chi)\log f_\chi$ together with the $-\frac{1}{4}\log d_\Phi$ normalisation and the conductor discrepancy layer Λ; these are exactly the objects computed by `TypeNorm`, `ReflexField` and the class-group intrinsics. No numerical Φ is claimed here: the older data file and the 8T24 exclusion of Theorem 12.51 show that none is stored for the BHPS family. □

Theorem 12.51 (The 8T24 exclusion is structural)**.** *The CM-type database of* `Andreas-Pieper/Genus-4-RM-CM` *excludes the Galois group of the Mumford-type family by an explicit branch. In* `CM/cm-g4.m`*,* `CheckForJacobians` *contains the branch*

"`if k in [13, 24] then Write(done_filename, ...); return "Galois group 8T13 or 8T24; skipping"; end if;`" and `CM/CMTypes.m` assigns `exps[...] := {0, 4}` directly when `TransitiveGroupIdentification(G)` lies in $\{13, 24\}$, skipping the primitive-type enumeration. Conversely `NootCMCube.m` in the BHPS repository sets `G := TransitiveGroup(8,24)` and tests the CM type against six four-element faces *of a cube, and `old-group-search/group-scratchwork.m` likewise uses `TransitiveGroup(8,24)`. Consequently the numerical CM types and reflex fields produced by the genus-four CM-Jacobian search do not cover the* 58 *Mumford-type CM points of Appendix 12.7. The reflex fields used by the auxiliary search are six-dimensional fields with Galois label* 6T11*, pulled from the LMFDB collection `nf_fields_reflex`; that collection is external to both repositories.*

Proof. The four quoted code fragments were read verbatim from the downloaded files: lines 30–33 of `cm-g4.m`, line 57 of `CMTypes.m` (with $d = 4$ declared on line 1), line 8 of `NootCMCube.m` together with its `faces` list, and the `G2 := TransitiveGroup(8,24)` line of `group-scratchwork.m`. The `reflex-field-script.sage` loop calls `db.nf_fields_reflex.search({'nf_label': lab})` and keeps only records whose `rf_coeffs` has length 7 and whose Galois label is `6T11`; hence the reflex-field data are LMFDB data, not repository data. The exclusion is therefore a property of the search target (genus-four CM Jacobians versus Mumford-type specialisations), not a missing file. □

Theorem 12.52 (Exact disjointness of the two CM-field pools)**.** *Let $\mathcal{B}$ be the* 58 *CM-field labels of Tables 1 and 2 of Appendix 12.7 and let $\mathcal{P}$ be the* 529 *labels in the `CM_flds` list of `Data/mumford_jac.py`. Then*

$$\boxed{\mathcal{B} \cap \mathcal{P} = \varnothing}.$$

Writing $D(\ell)$ for the third component of a label $8.0.D(\ell).i$, exact integer comparisons give

$$\min D(\mathcal{B}) = 166726039041, \quad \operatorname{med} D(\mathcal{B}) = 135923763363916801,$$

$$\max D(\mathcal{B}) = 1135151485651484161,$$

$$\min D(\mathcal{P}) = 142635249, \quad \operatorname{med} D(\mathcal{P}) = 210524051241, \quad \max D(\mathcal{P}) = 1490665382666496,$$

with 47 *of the* 58 *BHPS discriminants exceeding* $\max D(\mathcal{P})$ *but* 11 *lying inside the range of $\mathcal{P}$ and still absent from it. The disjointness is therefore not a mere discriminant cut-off: it is consistent with the* 8T24 *exclusion of Theorem 12.51.*

Proof. Both label sets were parsed from the downloaded files and compared as exact strings; the intersection is empty. The discriminants were parsed as exact integers from the third component of each label and sorted; the medians are the middle elements of the sorted lists of lengths 58 and 529. Counting $D(\mathcal{B}) > \max D(\mathcal{P})$ gives 47, hence 11 lie at or below 1490665382666496; none of those 11 occurs in $\mathcal{P}$. All comparisons are integer comparisons with zero floating-point error. □

Theorem 12.53 (An exact index instance for the discriminant upper bound)**.** *The degree-eight field printed in `automatic_examples/8.0.1635150969.3.m` has defining polynomial*

$$f(x) = x^8 - 5x^6 + 2x^4 + 8x^2 + 1, \qquad \operatorname{disc}(f) = 418598648064 = 16^2 \cdot 1635150969,$$

while the label gives $|d_K| = 1635150969 = 3^4 \cdot 4493^2$. Hence the index of the polynomial order is exactly 16*, and the surplus incurred in Appendix 12.7 by using* $\operatorname{disc}(P_K)$ *in place of $|D_K|$ in the bound B_K^{up} is, for $g = 4$,*

$$(2g-2)\log\frac{\operatorname{disc}(f)}{|d_K|} = 6\log 256 = 33.2710646669\ldots.$$

Proof. The discriminant of f was computed exactly by symbolic elimination and factored as an integer; the quotient by 1635150969 is $256 = 16^2$, and $1635150969 = 3^4 \cdot 4493^2$ is its exact factorisation. The polynomial $f(x) = g(x^2)$ with $g(y) = y^4 - 5y^3 + 2y^2 + 8y + 1$, and g has two positive and two negative real roots, so f has four real roots and two complex conjugate pairs; accordingly we use f only for the discriminant identity and do not infer from it that this field is a CM field (the label's signature datum differs). The logarithmic surplus is $6 \log 256$ because $(2g - 2) = 6$ for $g = 4$, evaluated numerically to the quoted digits. □

Theorem 12.54 (The Colmez disjunction: the computable class is the excluded class)**.** *Let (II) denote the remaining open item as reduced in Appendices 11.10–11.11: the existence of an absolutely simple, non-*GL_2*,* $\mathrm{End}^0 = \mathbb{Q}$ *isogeny factor with* $R_\rho > 3$ *and positive limiting conductor weight. By Appendix 12.5, the unique candidate seed is the* generic *Mumford-type fourfold, whose Mumford–Tate group is of* SL_2^3 *type and whose endomorphism algebra is* $\mathbb{Q}$*: it carries no CM. The specialisations at the* 58 *CM points, by contrast, acquire an eighth-degree CM field in* End^0*. The averaged Colmez formula (finite form) computes the Faltings height of a* CM *abelian variety from its CM type, the characters* χ *and the type discriminant* d_Φ*; it is therefore applicable exactly to the CM specialisations and not to the generic seed. For the CM specialisations, the CM–Parseval bound of Appendix 8.1 gives*

$$\mu_{\mathrm{ratio}} \le 1, \qquad \kappa \le \tfrac{1}{8}, \qquad R \le \tfrac{1}{2}, \qquad \gamma = 3 - R \ge \tfrac{5}{2},$$

so they lie in the class already excluded from (II). Hence

the objects whose height Colmez computes exactly are the
objects (II) has excluded; the objects (II) needs are
those Colmez cannot compute.

This is the sixth occurrence in this audit of a separation of the form "non-CM ≠ non-automorphic" / "CM-available ≠ dangerous".

Proof. The reduction of (II) to a single absolutely simple factor with $\mathrm{End}^0 = \mathbb{Q}$ and $\liminf w_\rho > 0$, and the exclusion of the CM and GL_2 classes, are the conclusions of Appendices 11.10, 11.11 and 8.1; the identification of the generic Mumford-type fourfold as the unique surviving seed is Appendix 12.5. The CM status of the 58 specialisations follows from their labels (degree-eight CM fields) as recorded in Appendix 12.7. The averaged Colmez formula, in the finite form $\mu(\Phi) = \sum_\chi m(\Phi, \chi) \log f_\chi$ with the $-\frac{1}{4} \log d_\Phi$ normalisation used throughout Appendices 5.4–5.8, is a statement about abelian varieties with complex multiplication by a CM field and a chosen primitive CM type; no such type exists for a variety with $\mathrm{End}^0 = \mathbb{Q}$. Substituting $\mu_{\mathrm{ratio}} \le 1$ into $\kappa = \mu_{\mathrm{ratio}}/8$ and $R = 12\kappa - 1$ gives $\kappa \le 1/8$, $R \le 1/2$ and $\gamma = 3 - R \ge 5/2$, so the CM specialisations cannot carry $R > 3$. Both inclusions are therefore strict and the intersection of the Colmez-computable class with the (II) candidate class is empty. □

Theorem 12.55 (The κ point-value entry is withdrawn)**.** *The proposal at the end of Appendix 12.7, namely that recovering a single CM type* Φ *and inserting it into the Colmez formula is the unique text-level entry to a theoretical value of* κ*, is withdrawn. Three independent obstructions persist even if* Φ*, its reflex field and* d_Φ *are obtained:*

(a) the point-value ambiguity $\Delta\gamma_A = 48c/\log \mathrm{cond}$ *of Theorem 12.48 does not collapse on CM points: with* $c = \frac{1}{2}\log(2\pi)$ *and* $48c = 44.1090495938$*, the exact data of Appendix 12.7 give* $\Delta\gamma_A^{\min} \in [0.17588, 4.78351]$*, with* 6 *of* 38 *already exceeding* 1*;*

(b) the base-discriminant term $(2g-2)\log|D_K|$ *is the structural term* $B = 2/c$ *of Appendix 11.6; Appendix 11.6 proved it is not absorbable, and Theorem 12.53 quantifies the further index surplus;*

(c) a CM abelian variety has potentially good reduction, so its conductor is constant and tower axiom (T3) fails for any family built from a fixed seed, as recorded in Appendix 12.5; moreover Appendices 5.4 and 5.8 proved that the conductor is not determined by the character conductors alone (the discrepancy layer Λ with $c(p) = (p^2+2p+3)/(p^2-1)$ is required), and Λ for the 8T24 CM points has not been computed.

Consequently the output of a successful Colmez computation here would be a further verification of $\kappa \le 1/8$ for CM factors, not a value of γ_A and not a contribution to (II). The fourteen No-Go statements, the constants $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \ge 36/7 > 3$, and the status abc = 'O' are unchanged

Proof. Item (a) is Theorem 12.48 with the exact ranges of Appendix 12.7; item (b) is the structural-term conclusion of Appendix 11.6 together with Theorem 12.53; item (c) is the potentially-good-reduction conclusion of Appendix 12.5 and the strict negation of "conductor = character conductors" in Appendices 5.4 and 5.8. Theorem 12.54 shows that even a complete Colmez computation only returns the already-excluded class. The three obstructions are independent: they concern normalisation, the base field and the conductor respectively, and none is removed by supplying Φ. No new No-Go statement and □

Remark 12.56 (What remains). The text-level and finite-computation chains are now exhausted. The only remaining routes are (i) a height formula or rigidity input for the generic non-CM Mumford-type fourfold, i.e. an arithmetic ball-quotient/Kaehler–Einstein rigidity giving $\gamma > 0$, and (ii) the computation of the conductor discrepancy layer for the 8T24 specialisations. The LMFDB API returned a policy denial during this round, so no reflex-field coefficients were retrieved; all reflex-field statements above are about the *code paths*, not about numerical fields. Neither Magma nor Sage nor Pari/GP is available in the present environment, and no proof of any cited theorem or preprint is reproduced here. All statements are unconditional relative to the quoted inputs; this appendix does not constitute an independent verification of Theorem A.

12.9 The 8T24 Specialisation of Λ

Convention. Let E be a degree-eight CM field with maximal totally real subfield E_0, and let Φ be a CM type, that is, a choice of one embedding from each of the four complex-conjugate pairs. Let A be the corresponding CM abelian fourfold. We write

$$\Lambda := \log \mathrm{cond}(A) - 2\,T(\Phi), \qquad T(\Phi) = \sum_{\chi \in \Phi} \log f_\chi,$$

where f_χ denotes the Artin conductor of the one-dimensional character χ. This is the single-point analogue of the tower-level quantity of Appendix NN, whose finite form was computed in Appendix TT.

Theorem 12.57 (Group identification and withdrawal of the order-1344 model)**.** *The transitive group* 8T24 *is $C_2 \times S_4 = C_2 \wr S_3$, of order* 48*, acting on the eight vertices of a cube. A first computation modelled* 8T24 *as* $\mathrm{AGL}(3,2) = 2^3{:}\mathrm{GL}(3,2)$ *of order* 1344*; all enumeration data obtained from that model are withdrawn.*

Proof. The Bristol GroupNames table of transitive groups of degree eight lists 8T24 as $C_2 \times S_4$, the full symmetry group of the cube, of order 48; the group $\mathrm{GL}(3,2)$ of order 168 is listed separately as 8T37. Section 2.2 of the BHPS paper states verbatim that the Galois group embeds into $\mathbb{Z}/2 \wr S_3 \cong \mathbb{Z}/2 \times S_4$, the automorphism group of a cube, and that the embeddings of E are in bijection with the vertices of that cube. A direct enumeration gives $|C_2 \wr S_3| = 8 \cdot 6 = 48$.

Two independent structural checks exclude the order-1344 model. First, complex conjugation of a CM field is a fixed-point-free involution c; in the cube model c is the central inversion $v \mapsto v + (1,1,1)$, and a direct element-by-element check shows that c is central in $C_2 \wr S_3$. The

group AGL$(3,2)$ has trivial centre, so it cannot host such an involution. Second, the withdrawn model produced sixteen CM types with stabilisers of orders 96 and 24, no primitive types, and reflex degrees 14 and 56; none of these values is compatible with $|G| = 48$. Hence the withdrawn data must not be used. □

Theorem 12.58 (The vanishing of Λ for CM types)**.** *For every degree-eight CM field and every CM type Φ one has $\Lambda(\Phi) = 0$ exactly, that is,* $\log \mathrm{cond}(A) = 2T(\Phi)$*. In particular the* 8T24 *specialisation of Λ vanishes.*

Proof. The Hodge decomposition of a CM abelian variety gives $H^1(A, \mathbb{C}) = \bigoplus_{\chi \in \Phi \sqcup \bar{\Phi}} \chi$, a direct sum of eight one-dimensional characters. Artin conductors are additive over direct sums, and complex conjugation pairs Φ with $\bar{\Phi}$ while preserving conductor exponents, so $f_\chi = f_{\bar{\chi}}$ for every χ. Therefore

$$\log \mathrm{cond}(A) = \sum_{\chi \in \Phi \sqcup \bar{\Phi}} \log f_\chi = 2 \sum_{\chi \in \Phi} \log f_\chi = 2T(\Phi),$$

and $\Lambda = 0$. The identity is structural: it does not depend on the Galois group, on the field discriminants, or on the particular conductor exponents.

As an exact finite check, assign to each of the four conjugate pairs an arbitrary nonnegative integer conductor exponent, shared by the two members of the pair. For five such exponent patterns and all sixteen CM types of the cube model (eighty checks in total), the integer $\log \mathrm{cond} - 2T$ is exactly zero in every case. These toy exponents verify the identity, not any physical conductor.

Consequently the nonzero tower constant $c(p) = (p^2 + 2p + 3)/(p^2 - 1)$ of Appendix NN is not a CM-field invariant. There T is the odd Dirichlet-character mass of a cyclotomic orbit, and the characters of $(\mathbb{Z}/p^n)^2$ split by parity with unequal masses; here H^1 is exactly $\Phi \sqcup \bar{\Phi}$ and no parity split exists. This is the same mechanism as the vanishing of $\Lambda^{(a)}$ for the mixed Kummer towers of Appendix RR. □

Theorem 12.59 (Three-layer enumeration of the sixteen CM types)**.** *For $G = C_2 \wr S_3$ with complex conjugation the central inversion, the sixteen CM types split into three orbits under G: six types with stabiliser order* 8*, eight types with stabiliser order* 6*, and two types with stabiliser order* 24*. Under the database criterion $|G|/|H_\Phi| = 8$, exactly eight types are primitive. The six face-of-the-cube CM types of Noot are precisely the six types with stabiliser order* 8*, hence all six are non-primitive.*

Proof. Represent the eight embeddings by the vectors of $\mathbb{F}_2^3$; let G act by coordinate permutations followed by coordinate flips, giving 48 permutations, all even. Complex conjugation is the flip by $(1,1,1)$, which is central and has four orbits of size two. A CM type is a choice of one vector from each of the four pairs, giving $2^4 = 16$ types. For each type Φ compute $H_\Phi = \{g \in G : g\Phi = \Phi\}$ by exhaustive enumeration; the resulting multiset of orders is 8 (six types), 6 (eight types), and 24 (two types), with orbit sizes $48/8 = 6$, $48/6 = 8$, and $48/24 = 2$. A type is primitive under the criterion $|G|/|H_\Phi| = 8$ exactly when $|H_\Phi| = 6$, so there are eight primitive types.

The six faces of the cube are the affine planes $v_i = b$ with $i \in \{1,2,3\}$ and $b \in \{0,1\}$. Each face contains exactly one vertex from each conjugate pair, hence is a CM type; exhaustive enumeration shows these six types have stabiliser order 8 and orbit size 6. They are therefore non-primitive.

It follows that the exclusion of the transitive groups 13 and 24 in the CM-type repository is a category mismatch, not an absence of primitive types: the database enumerates primitive types, whereas the Mumford-type CM fourfolds of Noot use the face types, which are non-primitive. This sharpens the corresponding claim of Appendix YYY. Note also that Noot's assertion that every *simple* Mumford-type CM abelian variety arises from a face type concerns simplicity of the abelian variety, which is a different notion from primitivity of the CM type; no contradiction arises. □

Theorem 12.60 (Discriminants: square and fourth power)**.** *The whole group $G = C_2 \wr S_3$ lies in A_8, so every field discriminant in the family is a perfect square. The first field of Table 1, with*

label 8.0.7834003547041.1*, has discriminant* $7834003547041 = 7^4 \cdot 239^4 = 1673^4$*, a fourth power. The sample polynomial* $x^8 - 5x^6 + 2x^4 + 8x^2 + 1$ *of Appendix XXX has polynomial discriminant* $646992^2 = 2^8 \cdot 3^4 \cdot 4493^2$ *and, with index* 16*, field discriminant* $40437^2 = 3^4 \cdot 4493^2$*.*

Proof. Every generator of G is an even permutation on the eight vertices: a coordinate transposition fixes four vertices and swaps two pairs, and a single coordinate flip is a product of four transpositions. Hence $G \subset A_8$ and the discriminant of the degree-eight field is a square. Exact integer factorisation gives $7834003547041 = 7^4 \cdot 239^4 = (7 \cdot 239)^4 = 1673^4$, consistent with Section 2.2 where the maximal quaternion order forces a fourth power. For the sample polynomial, SymPy gives $\mathrm{disc} = 418598648064 = 646992^2 = 2^8 \cdot 3^4 \cdot 4493^2$; dividing by 16^2 yields $1635150969 = 40437^2 = 3^4 \cdot 4493^2$. Both instances are independent of the group-theoretic argument and agree with it. □

Theorem 12.61 (The three-term constant is tower-level only)**.** *At a single CM point, only the first term of the decomposition* $c(p) = 1 + 2/(p-1) + 2/(p^2-1)$ *survives. In particular* $c(3) = 9/4 = 1 + 1 + 1/4$*, and* $\Lambda = 0$ *is compatible with* $c(3) = 9/4$*.*

Proof. Appendix TT derives the exact identity $c(p) = 1 - \lambda_p + 2/(p^2-1)$ with $\lambda_p = -2/(p-1)$, where the unit term comes from the anchor $\hat{f}_p = (2n+1)g_N$, the term $2/(p-1)$ from the discriminant shift $D = \sum_c A_c(c - 1/(p-1))$, and the term $2/(p^2-1)$ from the geometric tail $\sum_c cA_c/g_N \to n - 1/(p^2-1)$. All three are limits along a tower with $n \to \infty$. At a single CM point there is no layer index, no tail, and no anchor shift; the conductor mass is exactly $2T(\Phi)$ by Theorem 12.58, so the excess Λ vanishes. The decomposition therefore describes a tower-level bookkeeping surplus, not a pointwise invariant. □

Theorem 12.62 (Status of the long-term item and of the barrier)**.** *The long-term item left open in Appendix YYY, namely the computation of* Λ *for the* 8T24 *specialisation, is closed with value zero. The normalisation band* $\Delta\gamma_A = 48c/\log \mathrm{cond}$ *of Appendices WWW and XXX is unchanged, the cycle of Appendix XXX is unchanged, the dangerous seed of the barrier remains the generic non-CM Mumford fourfold with endomorphism algebra* $\mathbb{Q}$*, and the barrier statement (II) remains open.*

Proof. The first claim is Theorem 12.58. The band $\Delta\gamma_A = 48c/\log \mathrm{cond} = 44.10905/\log \mathrm{cond}$ arises from the difference $c \cdot g$ between normalisations of the absolute Faltings height; it is a height-side quantity, whereas Λ is a conductor-side quantity, so $\Lambda = 0$ cannot remove it. The cycle of Appendix XXX requires a Colmez height value to fix c, and that value carries the structural term $B_K = 2/c$; Theorem 12.58 supplies no height value. Finally, the disjunction theorem of Appendix YYY shows that the computable CM specialisations have $\kappa \le 1/8$, $R \le 1/2$, $\gamma \ge 5/2$ and are already excluded from the barrier, while the unique dangerous class is the generic non-CM Mumford fourfold, to which no CM type is attached. Hence (II) is unaffected. The fourteen No-Go statements, all numerical constants, and the *abc* status letter are unchanged; □

Remark 12.63 (Scope and honest limitations). Theorems 12.57–12.62 are unconditional relative to standard CM theory, the Bristol transitive-groups table, Section 2.2 and Appendix A of arXiv:2510.00093, and the internal appendices NN, RR, TT, WWW, XXX and YYY. The interpretation of the repository criterion $|G|/|H| = 8$, the correspondence between face types and simple Mumford-type abelian varieties, and the attribution of the Table 1 labels are text-level only; the raw repository sources could not be retrieved in this round and were cross-checked against the arXiv HTML version and the GroupNames table instead. No external theorem is reproved here. Reflex-field numerical data, absolute Faltings heights, and the true conductor exponents of an 8T24 CM fourfold are not computed; the toy exponent patterns of Theorem 12.58 verify an identity only. The withdrawn order-1344 computation is retained on disk solely as an erratum record.

12.10 Closing the Text-Level Items on the 8T24 Source Predicates

Convention. We keep the notation of Appendix ZZZ: E is a degree-eight CM field, Φ a CM type, $G = C_2 \wr S_3$ of order 48 acting on the eight cube vertices, c the central inversion, $H_1 = \mathrm{Stab}_G(1)$ of order 6, and A the CM abelian fourfold of type (E, Φ). All computations in this appendix are exact finite permutations; no floating point and no computer algebra system are used.

Theorem 12.64 (The two stabilisers in the source code are different objects)**.** *In the routines* `IsPrimitiveType` *and* `ReflexType` *of* `CMTypes.m`*, the set $S = \{g \in G : 1^g \in \Phi\}$ is a set of 24 group elements, not the four-element set Φ. The subgroup $H = \{h \in G : hS = S\}$ is the left stabiliser of S, whereas $H_r = \{h \in G : Sh = S\}$ is its right stabiliser. On 8T24 one has $H_r = \mathrm{Stab}_G(\Phi)$ for all sixteen CM types, while H has order 6 for fourteen types and order 24 for two types. Consequently $|G|/|H|$ equals 8 for fourteen types and 2 for two types, and $|G|/|H_r|$ equals the size of the Galois orbit of Φ, namely 8, 6 or 2.*

Proof. The source lines read, verbatim, `S := {g: g in G| 1^g in Phi}`, `H := [h: h in G| {h*s: s in S} eq S]` and `Hr := [h: h in G| {s*h: s in S} eq S]`. With the Magma right action $x^{s*h} = (x^s)^h$, the products $h * s$ and $s * h$ correspond respectively to right and left multiplication on S, so H and H_r are the left and right stabilisers of S.

Exhaustive enumeration of the 48 permutations preserving the six cube faces gives $|S| = 4|H_1| = 24$. For each of the $2^4 = 16$ CM types, computing H and H_r as explicit subsets of G yields $|H| \in \{6, 24\}$ with multiplicities 14 and 2, and $|H_r| = |\mathrm{Stab}_G(\Phi)|$ with the multiset of orders 6 (eight types), 8 (six types) and 24 (two types). Since G acts transitively on the eight vertices, $|G|/|H_r|$ is the orbit size of Φ. All sixteen pairs of stabilisers were checked element by element; the identities $H_r = \mathrm{Stab}_G(\Phi)$ and $H_1 \subseteq H$ hold in every case. □

Theorem 12.65 (Withdrawal of the primitive count of Appendix ZZZ)**.** *The statement of Theorem 12.59 that exactly eight CM types are primitive and that the six face-of-the-cube types are all non-primitive is withdrawn. The correct count is fourteen primitive types and two non-primitive types; all six face types are primitive. Moreover the code predicate* `IsPrimitiveType` *agrees with the Shimura–Taniyama notion of primitivity (not induced from a proper CM subfield) for all sixteen types.*

Proof. By Theorem 12.64, the criterion `#G/#H eq 2*d` with $d = 4$ holds exactly when $|G|/|H| = 8$, which occurs for the fourteen types with $|H| = 6$; the two types with $|H| = 24$ give $|G|/|H| = 2$ and are rejected. The primitive-test table of Appendix ZZZ used $|\mathrm{Stab}_G(\Phi)|$ instead of $|H|$, which by Theorem 12.64 is the right stabiliser H_r belonging to `ReflexType`: it therefore counted reflex degrees, not primitivity.

Independently, a CM type is induced from a proper CM subfield K' exactly when there is a subgroup J with $H_1 \subseteq J \subseteq G$, $[G : J] \in \{2, 4\}$, $c \notin J$, such that Φ is a union of the G-translates of the J-orbit of the vertex 1, these translates forming a partition of the eight vertices. Enumerating all subgroups generated by H_1 together with one or two further elements of G produces exactly one admissible J: it has order 24, index 2, two blocks of size 4, and $c \notin J$. Exactly two CM types are unions of such blocks. Hence the standard primitive count is $16 - 2 = 14$, and it coincides type by type with the code predicate. The six faces have orbit size 6, $|H| = 6$, and are not unions of the two blocks, so all six are primitive. □

Theorem 12.66 (Reconciliation of the cube-face types with simple CM abelian varieties)**.** *There is no conflict between Noot's assertion that every simple Mumford-type CM abelian variety arises from a face-of-the-cube CM type and the fact that the two non-primitive types are induced from a proper CM subfield. The two non-primitive types are precisely the cases in which the CM field contains an imaginary quadratic subfield F and $\mathrm{End}^0(A) \cong L \oplus F$, so A is not simple; the six face types are primitive, hence simple. The closing sentence of Theorem 12.59, which declared primitivity of a CM type and simplicity of the abelian variety to be notions of different dimensions that must not be merged, is withdrawn.*

Proof. Section 2.2 of the BHPS paper states verbatim that taking the four embeddings corresponding to a face of the cube defines a CM type Φ of E, that it yields a CM abelian variety of type (E, Φ) at a point of the Mumford-type Shimura curve, that conversely every simple Mumford-type CM abelian variety arises in this way, and that non-simple CM abelian varieties arise when L is allowed to contain an imaginary quadratic field F, in which case $\mathrm{End}^0(A) \cong L \oplus F$. The standard CM theory, as recorded for instance in the exposition on abelian varieties with prescribed embedding degree, states that primitive CM types correspond to simple abelian varieties.

By Theorem 12.65 the only non-primitive types are the two induced from the unique subgroup J of index 2; the fixed field of J is a quadratic CM field, that is, an imaginary quadratic field F. Thus those two types are exactly the non-simple cases described by Noot, and the six face types, being primitive, lie in the simple family. The remaining eight primitive types are generic CM fourfolds not lying on the Mumford curve, which is compatible with the word "every" in Noot's statement. Hence the two notions do correspond on this family, and the earlier separation was an artefact of the misidentified stabiliser. □

Theorem 12.67 (Reflex degree six and the LMFDB search condition)**.** *For each of the six cube-face CM types the reflex field has degree* 6 *and the reflex type has three embeddings. The conditions* `len(cs) == 7` *and* `transitive_label() == '6T11'` *in the BHPS script* `reflex-field-script.sage` *are therefore independent textual confirmation of this degree, not a failed lookup. The general fact that for* $g \geq 3$ *a primitive CM type may have reflex degree different from* $2g$ *is what made the confusion between primitivity and reflex degree possible.*

Proof. By Theorem 12.64 the reflex degree equals $|G|/|H_r|$, which for a face type is $48/8 = 6$. The routine `ReflexType` forms the cosets S/H_r and returns the inverses of one representative per coset; direct enumeration gives three cosets and hence a reflex type of size $3 = 6/2$, as required for a degree-six CM field. The script filters reflex-field coefficient vectors of length 7, which is the number of coefficients of a degree-six polynomial, and selects Galois group 6T11, a transitive group of degree six. Both filters match the computed degree.

An exposition on abelian varieties with prescribed embedding degree states verbatim that for CM fields of degree 2 or 4 with primitive CM types the reflex field has the same degree as the field, and that this fails for $g \geq 3$. Here $g = 4$ and the six face types are primitive with reflex degree $6 \neq 8$, an explicit instance. Appendix XXX recorded the LMFDB refusal as a missing computation; the search conditions themselves already encode the answer, so that item is a textual confirmation rather than an open calculation. □

Theorem 12.68 (The two skipped group identifiers are the two cube groups)**.** *The identifiers* $\{13, 24\}$ *skipped in* `cm-g4.m` *and hard-coded in* `Exponents` *are exactly the two images of the homomorphism* φ *described in Section 2.2 of the BHPS paper:* 8T13 $= E(8){:}3 = A(4)[x]2$ *of order* 24*, isomorphic to* $\mathbb{Z}/2 \wr \mathbb{Z}/3 \cong \mathbb{Z}/2 \times A_4$*, for the Galois cubic case; and* 8T24 $= E(8){:}D_6 = S(4)[x]2$ *of order* 48*, isomorphic to* $\mathbb{Z}/2 \wr S_3 \cong \mathbb{Z}/2 \times S_4$*, for the non-Galois case. Both are even groups with centre of order two. The exclusion is a deliberate split by group identifier between two reconstruction pipelines, not an absence of primitive CM types and not a category mismatch.*

Proof. The Galois database table of transitive groups of degree eight lists 8T13 with name $E(8){:}3 = A(4)[x]2$, order 24, centre of order 2 and the property "even"; it lists 8T24 with name $E(8){:}D_6 = S(4)[x]2$, order 48, centre of order 2 and "even". Section 2.2 states that φ is an isomorphism onto $\mathbb{Z}/2 \wr S_3 \cong \mathbb{Z}/2 \times S_4$ when the cubic field is not Galois, and that its image is $\mathbb{Z}/2 \wr \mathbb{Z}/3 \cong \mathbb{Z}/2 \times A_4$ of order 24 in the Galois cubic case, which occurs for the examples $\mathbb{Q}(\zeta_7)^+$ and $\mathbb{Q}(\zeta_9)^+$. Hence the two skipped identifiers cover precisely the Mumford-type CM family.

The routine `IsGaloisGroupOfCMField` accepts a transitive group when it is imprimitive, when its minimal partition has four blocks, and when the involution swapping the two elements of each block lies in the group. For 8T24 the four antipodal pairs form such a partition, the involution is the central inversion c, and exhaustive enumeration confirms that this is the unique

partition into four blocks of size two. Hence 8T24 is iterated by `Exponents`. By Theorem 12.65 the list `CMtypespri` is nonempty, with representatives from the orbits of sizes 6 and 8. The source nevertheless assigns `exps[...] := {0, 4}` by an explicit branch on the identifier, and `CheckForJacobians` returns immediately for $k \in \{13, 24\}$. The BHPS repository handles the same family through `NootCMCube.m`, which tests whether a CM type is one of the six cube faces. The two pipelines are therefore separated by group identifier by design. □

Theorem 12.69 (Attribution of the first Table 1 label)**.** *Appendix A, Table 1 of the BHPS paper contains* 58 *distinct degree-eight CM field labels. Its first row is* 8.0.7834003547041.1 *with discriminant* $7^4 \cdot 239^4 = 1673^4$ *and field of definition* 3.3.49.1. *This label occurs in `Shimura_237_param.m`, does not occur in `Shimura_239_param.m`, and does not occur in the* 529*-label pool of `Data_mumford_jac.py`. Moreover the complex-conjugation representative* $(1,6)(2,5)(3,4)(7,8)$ *listed for* 8T13 *in the Galois database coincides digit by digit with the central inversion derived from the six faces of `NootCMCube.m`, so the vertex numbering of the code agrees with the database numbering.*

Proof. Parsing the local verbatim full-text extraction of Appendix A gives 58 distinct labels of the form 8.0.∗, the first being 8.0.7834003547041.1 followed by $7^4 \cdot 239^4$ and 3.3.49.1; a direct integer factorisation gives $7834003547041 = (7 \cdot 239)^4 = 1673^4$. Exact string membership checks against the two parameter files and the 529-label pool give the stated pattern, consistent with the empty intersection recorded in Appendix YYY. Finally, the six faces listed in `NootCMCube.m` are complementary in pairs and force the antipodal pairing $(1,6), (2,5), (3,4), (7,8)$; the Galois database lists this same involution as the length-one representative of complex conjugation for 8T13. The numbering is therefore canonical rather than an artefact of the model. □

Remark 12.70 (Scope and remaining boundary). Four text-level items left open in Appendix ZZZ are closed here: the interpretation of the `#G/#H=8` predicate, the Noot face-versus-simple correspondence, the attribution of the Table 1 labels, and the earlier inability to fetch the raw GitHub sources (complete local copies were re-read line by line). One finer textual item is newly isolated: `NootCMCube.m` conjugates against the order-48 group 8T24, whereas Section 2.2 identifies the $\mathcal{X}_7$ and $\mathcal{X}_9$ examples with the Galois cubic case whose image has order 24; matching group identifiers to individual points is not verified here. The identity $\Lambda = 0$ of Appendix ZZZ is untouched, since it concerns the conductor layer, while the normalisation band of γ_A lies in the height layer. The long-term item (II), the fourteen No-Go statements, and are unchanged.

12.11 Pointwise Group Identifiers of the Fifty-Eight CM Points

Theorem 12.71 (Verbatim scope of the tabulated fields)**.** *Appendix A of arXiv:2510.00093v2 lists* 58 *distinct degree-eight CM field labels,* 38 *in Table 1 (the family* X_7*) and* 20 *in Table 2 (the family* X_9*). The text states verbatim that these are degree-eight CM fields with Galois group* $\mathbb{Z}/2 \times A_4$ *and Galois closure containing* $\mathbb{Q}(\zeta_7)^+$ *or* $\mathbb{Q}(\zeta_9)^+$*, and separately that in both situations the CM fields had Galois group* $\mathbb{Z}/2 \times A_4$*.*

Proof. Parsing the local verbatim full-text extraction of Appendix A gives 38 rows of the form 8.0.∗ between the header of Table 1 and the page break, and 20 further rows under Table 2, all distinct; this reproduces the 58 labels of the two parameter files `Shimura_237_param.m` and `Shimura_239_param.m` row by row. The two quoted sentences occur in Section 3 and in Section 4.1 of the paper, and are confirmed against the HTML rendering of the same preprint. Section 2.3 identifies the triangle groups $\Delta(2,3,7)$ and $\Delta(2,3,9)$ with the totally real cubic fields $\mathbb{Q}(\zeta_7)^+$ and $\mathbb{Q}(\zeta_9)^+$. □

Theorem 12.72 (The dichotomy forces 8T13 on every row)**.** *Let* $G = \mathrm{Gal}(\widetilde{L}/\mathbb{Q})$ *be the group of Section 2.2, embedded by* φ *into* $\mathbb{Z}/2 \wr S_3 \cong \mathbb{Z}/2 \times S_4$*, of order* 48*. If the cubic field* $K/\mathbb{Q}$ *is not Galois then* φ *is an isomorphism and* G *acts on the eight cube vertices as* 8T24*, of order* 48*. If*

$K/\mathbb{Q}$ is Galois cyclic of degree three then the image of φ is $\mathbb{Z}/2 \wr \mathbb{Z}/3 \cong \mathbb{Z}/2 \times A_4$, of order 24, and the corresponding transitive group is 8T13*. Consequently all* 58 *tabulated fields have group identifier* 8T13*, and* 8T24 *occurs zero times.*

Proof. Section 2.2 of the preprint states the two cases verbatim and adds that the Galois cubic case occurs in its examples, with $K = \mathbb{Q}(\zeta_7)^+$ respectively $\mathbb{Q}(\zeta_9)^+$. Both are cyclic cubic: their defining polynomials $x^3 - x^2 - 2x + 1$ and $x^3 - 3x - 1$ have discriminants $49 = 7^2$ and $81 = 3^4$ and Galois group of order 3, as computed exactly. Hence $\varphi(G) \cong \mathbb{Z}/2 \times A_4$ has order 24. The Galois database of transitive groups of degree eight lists $C_2 \times A_4$ on eight points as 8T13 and $C_2 \times S_4$ as 8T24, of orders 24 and 48 respectively. Since $E \subset \widetilde{L}$ has eight embeddings permuted transitively by G, its group identifier is determined by $|G|$, giving 8T13 for each of the 58 fields. The non-Galois S_3 branch would be required for 8T24 and is excluded by the verbatim statement. □

Theorem 12.73 (Seven exact pointwise checks)**.** *For each of the* 58 *rows the following hold simultaneously: the Table 1/2 label equals the corresponding parameter-file label; the degree of* $\mathrm{minpoly}(\mu)$ *equals the degree d_t of the field of definition; the integer d in the label* 8.0.d.idx *equals the exact product $\prod p_i^{e_i}$ read from the tabulated discriminant; $d = n^4$ is a fourth power; d is divisible by 7^4 for X_7 and by 3^8 for X_9; $3 \mid d_t$; and* $\mathrm{disc}(K_t)$ *is divisible by* $\mathrm{disc}(K_0)^k$ *with $k = d_t/3$,* $\mathrm{disc}(K_0) = 7^2$ *or* 3^4*. In addition the discriminant of* $\mathrm{minpoly}(\mu)$ *is divisible by* $\mathrm{disc}(K_0)^k$ *for all* 58 *rows.*

Proof. All quantities are integers and are compared exactly, without floating point. The first three checks are direct row-by-row comparisons of the parsed table and parameter files. The fourth is an integer fourth root: the first row gives $7^4 \cdot 239^4 = 1673^4$, the last row of Table 2 gives $3^8 \cdot 2843^4 = 25587^4$, and all 58 pass. The fifth holds because every X_7 discriminant contains the factor 7^4 and every X_9 discriminant contains 3^8; the sixth because $d_t \in \{3, 6, 9, 12, 15, 21\}$. For the seventh, the families X_7 and X_9 are defined over K_0, so every field of definition K_t contains K_0 and $[K_t : K_0] = k$; the relative discriminant formula gives $\mathrm{disc}(K_t) = \mathrm{disc}(K_0)^k N_{K_0/\mathbb{Q}}(\mathrm{disc}(K_t/K_0))$. For example the $k = 2$ field 6.2.31213.1 has $\mathrm{disc}(K_t) = 7^4 \cdot 13$, divisible by $7^{2\cdot 2}$, and a $k = 7$ field of Table 2 has discriminant divisible by 3^{28}. Finally $\mathrm{disc}(\mathrm{minpoly}(\mu)) = \mathrm{disc}(K_t) \cdot [\mathcal{O}_{K_t} : \mathbb{Z}[\mu]]^2$, hence the last assertion. An initial fixed-exponent version of this check, using 7^4 and 3^8 for every row, correctly returned only 47 of 58; the k-dependent exponent restores 58 of 58. □

Theorem 12.74 (Independent Galois-group computation)**.** *The two cubic polynomials defining K_0 have Galois group C_3 of order* 3 *and discriminants* 49 *and* 81*. Every one of the eight degree-six polynomials* $\mathrm{minpoly}(\mu)$ *occurring among the* 58 *rows has Galois group of order* 24 *whose element-order profile is* $\{1:1,\ 2:7,\ 3:8,\ 6:8\}$*, the profile of $C_2 \times A_4$, and not* $\{1:1,\ 2:9,\ 3:8,\ 4:6\}$*, the profile of S_4.*

Proof. The Galois groups were obtained by exact symbolic computation of irreducible polynomials over $\mathbb{Q}$ (degrees three and six), and the order profiles by enumerating all group elements. The profile of $C_2 \times A_4$ is computed from its seven involutions, eight elements of order three and eight of order six; S_4 instead has nine involutions and six elements of order four. All eight degree-six cases give order 24 and the $C_2 \times A_4$ profile, so $|\mathrm{Gal}(K_t)| = 24$ is consistent with $|G| = 24$ of Theorem 12.72 and with no 48-order closure. Independently the complex conjugation of the cube model is the central involution $(1\,6)(2\,5)(3\,4)(7\,8)$, agreeing with the representative listed for 8T13 in the Galois database and with the six cube faces of `NootCMCube.m`. □

Theorem 12.75 (The CM-type enumeration is independent of the two groups)**.** *Let $G_{13} = C_2^3 \rtimes C_3$ of order* 24 *and $G_{24} = C_2^3 \rtimes S_3$ of order* 48 *act on the eight cube vertices, with complex conjugation the central translation by* $(1, 1, 1)$*. For both groups the sixteen CM types split into orbits of sizes* 8*,* 6 *and* 2 *with multiplicities* 8*,* 6 *and* 2*; exactly* 14 *types satisfy the source predicate $\#G/\#H = 8$ with H the left stabiliser of $S = \{g : 1^g \in \Phi\}$, two types fail it, and all six cube-face types are primitive. The right stabilisers H_r and the left stabilisers H of G_{24} are exactly twice those of G_{13}.*

Proof. Both groups were realised as affine groups $v \mapsto Mv + a$ on $\mathbb{F}_2^3$, with M ranging over the three coordinate cycles for G_{13} and over all six coordinate permutations for G_{24}, and a over $\mathbb{F}_2^3$; the vertex numbering is the one reconstructed from the six faces of `NootCMCube.m`. For each of the sixteen subsets Φ choosing one vertex from each of the four antipodal pairs, the sets S, $H = \{h : hS = S\}$, $H_r = \{h : Sh = S\}$ and the orbit of Φ were enumerated over all group elements with exact permutation arithmetic. For G_{13} the left stabilisers are 3 (fourteen types) and 12 (two types), and the right stabilisers 3, 4, 12; for G_{24} they are 6, 24 and 6, 8, 24. Hence $\#G/\#H = 8$ holds for the same fourteen types in both groups, the six faces lie in the orbit-6 layer in both, and all proved conclusions of Sections 12.10 and ZZZ are unchanged. In particular $\Lambda \equiv 0$ of Section 12.9 follows from $H^1 = \Phi \oplus \bar{\Phi}$ and additivity of Artin conductors, and is group-independent. □

Theorem 12.76 (Scope of the cube-face source routine)**.** *The intrinsic* `CheckNootCondition` *of* `NootCMCube.m` *hard-codes* `TransitiveGroup(8,24)` *and tests* `IsConjugate(Sym(8),G,Gal)`*. For a field with* $\mathrm{Gal}(K) \cong$ 8T13 *this test is false, so the intrinsic returns false on all* 58 *points of Appendix A. The routines* `cm-g4.m` *and* `CMTypes.m` *instead skip the pair of identifiers* $\{13, 24\}$*, which covers both groups.*

Proof. Conjugate subgroups of Sym(8) have equal order; 8T13 has order 24 and 8T24 has order 48, so no field of type 8T13 can pass the conjugacy test against `TransitiveGroup(8,24)`. By Theorem 12.72 all 58 fields are of type 8T13. The intrinsic therefore implements the cube-face condition for the general Noot setting, in which the cubic field need not be Galois, and does not directly certify the tabulated examples; this is a code-scope statement, not a mathematical gap. Lines 30–33 of `cm-g4.m` and line 57 of `CMTypes.m` branch on the identifier set $\{13, 24\}$, so both the tabulated and the general cases are covered there. The reading-guide statement of Section 12.10 about a division of pipelines is sharpened accordingly: the Λ and CM-type conclusions are group-independent by Theorem 12.75, while the phrase "cube-face condition verified" must not be read as applying to these 58 fields. □

Remark 12.77 (Corrections and boundary). Three phrasings are corrected. The title of Section 12.9 ("8T24 specialisation") and the Table 1 attribution in Section 12.10 should read 8T13 for these points; the conclusions are unchanged. The last text-level item of Section 12.10, namely the pointwise verification of the fifty-eight group identifiers, is closed here with 58 of 58 equal to 8T13. A single new text-level item is recorded: the intended scope of `NootCMCube.m` on 8T13 fields requires author confirmation; this is not a new mathematical hypothesis. Route (II), the sign of γ_A, the normalisation band $\Delta\gamma_A = 44.10905/\log \mathrm{cond}$ and the generic non-CM Mumford fourfold as the dangerous seed are unchanged. The LMFDB label pages were blocked by a browser challenge during this round, so the identifiers follow from the verbatim family-level theorem together with the exact pointwise arithmetic fingerprints, not from a per-label database lookup.

12.12 : New tools and new theorems for item (II)

This appendix records the outcome of a multi-source audit of the toolchain and of the theorem landscape relevant to the single remaining open item (II). All external statements are quoted from the sources listed in the companion report; no proof of any quoted theorem is reproduced here. The conclusions of the preceding appendices are unchanged.

Theorem 12.78 (Local computational boundary)**.** *In the workspace used for this audit, none of the computer algebra systems Sage, Magma, PARI/GP, Julia, Macaulay2, Singular, polymake, or GAP is available. Among the numerical and symbolic Python libraries, only SymPy, mpmath, NumPy, and SciPy are present; FLINT, cypari2, gmpy2, abelfunctions, and the Sage Python interface are absent. Consequently exact integer, rational, and permutation computations are*

available, whereas local field arithmetic, regular or simple normal crossings models, and certified period matrices are not computable in this environment.

Proof. The result is the direct output of an executable enumeration of the available binaries and of the importable Python modules. No inference about other environments is made. In particular the absence of Arb and of cypari2 rules out, locally, any substitution of high precision complex ball arithmetic for the p-adic algebra required by conductor computations. □

Theorem 12.79 (End-to-end conductor gap)**.** *As of 2026 no public end-to-end implementation computes the conductor of the Jacobian of a general genus four curve over the rationals. The plane quartic algorithm of 2025 gives a general-genus Artin conductor formula*

$$n_p = \int_{-1}^{\infty} \operatorname{codim} V^{G_{\mathbb{Q}_p}^u}\, du = 2g - 2a - t + \sum_{i \geq 1} \frac{\operatorname{codim} J[\ell]^{G_i}}{[G_0 : G_i]},$$

in which the wild part is controlled by the Galois action on two- and three-torsion, including the prime two, and the tame part requires a simple normal crossings regular model. Its authors state that tame implementations are slow and inefficient and for that reason do not discuss the tame part; their benchmark of the first hundred database entries takes about four and a half hours and depends on Magma local field intrinsics and bitangent splitting fields.

Proof. The quoted formula and the quoted limitations are taken verbatim from the 2025 plane quartic conductor paper. The Magma handbook states that intrinsics for non-hyperelliptic modular curves of genus greater than three are planned for future releases, which certifies their absence in the documented version. The ANTS-XVII survey states that no efficient methods are available for the conductor of the Jacobian of a general genus three curve. Combining these three statements yields the maturity ordering: genus two hyperelliptic; genus three special families; general genus three; genus four plane quartic; general genus four, with only the first two levels mature. □

Theorem 12.80 (Hyperelliptic tool mismatch)**.** *Cluster picture packages expose functions named conductor exponent, tame exponent, wild exponent, and Tamagawa number, but their input is a cluster picture, whose theory is formulated for curves of the shape $y^2 = f(x)$ over local fields of odd residue characteristic. The curve $C_{7,t}$ of this volume is a non-hyperelliptic genus four curve with four pairs of complex conjugate branch points. Hence these functions cannot be applied to it without changing the object. Moreover hyperellipticity of the curve is not implied by the Jacobian being isogenous to a hyperelliptic Jacobian; only the former supplies the required input format.*

Proof. The package documentation states that its goal is to compute the local arithmetic data of hyperelliptic curves; the 2024 cluster-picture paper states that the theory is limited to odd residue characteristics and does not contribute results at places above two. The branch point structure of $C_{7,t}$ was computed exactly in Appendix WWW. The distinction between a property of the curve and a property of its isogeny class is immediate from the definitions. □

Theorem 12.81 (Period matrix is not the absolute Faltings height)**.** *A period determinant supplies only the archimedean volume term. The absolute Faltings height additionally requires finite place model metrics, intersection pairings, and Green functions, together with a fixed normalisation of the metrised line bundle. Two standard normalisations differ by $c \cdot g$, and the resulting uncertainty*

$$\Delta\gamma_A = \frac{48c}{\log \operatorname{cond}} = \frac{44.10905}{\log \operatorname{cond}}$$

does not shrink when the numerical precision of the period matrix is increased. Rigorous output is therefore a period matrix, its Siegel reduction, theta constants, the Bergman kernel, and an error radius inside a declared normalisation, not an absolute height without a specified model.

Proof. The rigorous hyperelliptic Faltings–Hriljac computation of Holmes replaces ordinary numerical integration, described as inherently non-rigorous, by formal integration of truncated power series, and obtains the bad fibre intersection from a Magma regular proper model; thus the archimedean and the finite contributions come from two independent codes. Maple algcurves, the Magma Riemann surface package, the Molin–Neurohr Arb implementation, and the Sage RiemannSurface package all supply period data, with ball arithmetic or certified path tracking, but none supplies the finite model arithmetic. The normalisation shift is the exact identity recorded in Appendix WWW. □

Theorem 12.82 (K-stability gives uniqueness or minimisation, not a gap)**.** *The non-archimedean Calabi theorem of Yuan and Zhang is a uniqueness statement for semipositive metrics with prescribed top intersection; the arithmetic Hodge index theorem underlying it does not assert existence of an adelic Kaehler–Einstein metric. The 2026 work of Hattori and Odaka proves that certain minimal-like metrised models minimise the Arakelov analogue of the Mabuchi K-energy under special K-stability hypotheses, while their Conjecture 1.2 remains a conjecture. Odaka's arithmetic Yau–Tian–Donaldson correspondence likewise remains conjectural in general. Ax–Schanuel and o-minimality results for variations of Hodge structure, and Andre–Oort type finiteness theorems, constrain special loci but provide no explicit positive lower bound for the height of a generic non-CM point. Hence no quoted theorem yields the implication MT corank positive implies $\gamma > 0$.*

Proof. The uniqueness wording is quoted verbatim from the Yuan–Zhang preprint, the minimisation wording from the Hattori–Odaka abstract, and the conjecture status from their Conjecture 1.2. The averaged Colmez theorem and the CM height formula of Colmez take a CM field together with a CM type as input; Kudla program surveys state that the higher rank case remains open and that the proved height and L-value identities concern special or CM cycles. The three statements, generic Mumford–Tate group full, arithmetic properties of special points, and a positive lower bound for a height on a positive dimensional family, are logically distinct. Positive characteristic counterexamples to Bogomolov–Miyaoka–Yau forbid a place-by-place transfer of the complex inequality but do not remove the need for a global argument. □

Theorem 12.83 (Exact calibration of the receipt discriminant data)**.** *For all fifty-eight LMFDB receipts the field discriminant D extracted from the label is an exact fourth power, the six uniform fingerprint fields agree row by row, the root discriminant agrees with $D^{1/8}$ to within 5.0×10^{-10}, and for the fourteen full records the product of the class group invariants equals the class number. For each full record, however, $\prod_p p^{e_p}$, formed from the ramified primes and the* galois_disc_exponents *field, differs from D by between thirty-four and thirty-nine decimal orders of magnitude. The exact relation is*

$$D = \Big(\prod_p p^{a_p}\Big)^4, \qquad a_p = 2 \text{ if } e_p = 32, \qquad a_p = 1 \text{ otherwise},$$

which holds for fourteen out of fourteen full records, with exponent patterns $(16, 12)$ *nine times,* $(32, 12)$ *three times, and* $(16, 12, 12)$ *twice.*

Proof. All computations were performed with exact integer arithmetic: exact fourth roots, prime factorisation, resultant-free integer comparison, and exact products of the class group invariants. No floating point comparison enters any of the equalities; the only floating point quantity is the root discriminant, reported by the source and checked against $D^{1/8}$. The large separation between $\prod_p p^{e_p}$ and D shows that the *galois_disc_exponents* field records data at the level of the Galois closure, not the prime exponents of the field discriminant. The two-pattern description of Appendix BBBB is thereby confirmed and its scope clarified. □

Remark 12.84 (Corrections, ranking, and boundary). Three phrasings are refined. First, the closing remark of Appendix BBBB that further progress requires new tools or new theorems is

made precise: the tool deficit is a Magma licence together with place-by-place local arithmetic engineering, and the theorem deficit is the single intermediate proposition from Mumford–Tate corank to a positive height or curvature gap; cloud services and databases remove licensing and screening costs but not the missing algorithm. Second, Appendix YYY stated that the class whose height is exactly computable is precisely the class excluded by item (II); the correct statement is that the proved exact formulae take CM data as input, while the non-existence of a closed formula for generic non-CM points is unproved. Third, the receipt calibration above adds a convention note with no mathematical consequence. Five next steps are ranked in the companion report: a certified period matrix for one fixed $C_{7,t}$; place-by-place wild conductor exponents; a normalisation dictionary; a database screening of comparison samples; and the analytic MT-corank gap proposition. Item (II) remains open; the fourteen No-Go statements, the constant $C \geq 36/7 > 3$, and the status of the *abc* line are unchanged.

12.13 : the first Archimedean machine value

Throughout this appendix $C_{7,t}$ denotes the genus-4 hyperelliptic curve of Eq. (1.1) of arXiv:2510.00093v2, $y^2 = f_t(x)$ with $\deg f_t = 10$, and $\Delta(C_{7,t}) = 2^{36}3^{36}7^{10}t^{54}(t-1)^{12}$. We write $c = \frac{1}{2}\log(2\pi)$, $L_2 = 2\log(2\pi^4)$, $g = 4$, and $\widehat{\mathcal{E}}_{\mathrm{ar}} = -\log\det(\operatorname{Im}\tau)$. All numerical values below are recomputed at 40 decimal digits from the external run receipt $\det(\operatorname{Im}\tau) = 0.0621660315772327$.

Theorem 12.85 (Sage 10.9 compatibility preserves every exact identity)**.** *Replacing every occurrence of $QQ(a,b)$ by $QQ(a)/QQ(b)$ and deleting the call .expand() on a polynomial object leaves all exact identities of the $C_{7,t}$ script unchanged. In particular $2^{16}\operatorname{disc}(f_t) = \Delta(C_{7,t})$ still holds for $t \in \{2, 8, \frac{3}{2}, \frac{1}{2}, \frac{101}{100}, -1\}$.*

Proof. In the rational field $\mathbb{Q}$ the two constructions produce the same element: $QQ(a)/QQ(b)$ is the fraction a/b, and $QQ(a,b)$ denotes the same rational number in the older calling convention. All eight affected sites are rational constants or quotients of integers occurring in coefficient arithmetic and in the slopes of a Newton lower hull; none involves a limit, a floating conversion or a branch depending on the constructor used. For the second change, a Sage polynomial ring is a normalised quotient $\mathbb{Q}[x]$: multiplication by a scalar is performed coefficient-wise and returns an already-expanded normal form, so $(tv \cdot P).expand()$ and $tv \cdot P$ are equal as elements of $\mathbb{Q}[x]$. Consequently f_t, $\operatorname{disc}(f_t)$ and every p-adic valuation are unaltered. The external run reports six out of six discriminants passing, which is therefore an instance of this invariance rather than a coincidence. The modified script also passes a pure-Python syntax check. □

Theorem 12.86 (Positivity and normalisation coherence)**.** *The external computation returns a 4×8 period matrix and a 4×4 Riemann matrix τ with*

$$\det(\operatorname{Im}\tau) = 0.0621660315772327 > 0.$$

Hence $\operatorname{Im}\tau$ is positive definite and the Riemann bilinear relation holds; the reduction from the 4×8 period matrix to the 4×4 normalised matrix is coherent.

Proof. For a genus-g curve the period matrix is $(\Omega_1\,\Omega_2)$ with $\tau = \Omega_1^{-1}\Omega_2$, and the first Riemann bilinear relation is equivalent to $\operatorname{Im}\tau$ being symmetric positive definite. A 4×4 matrix returned by `riemann_matrix()` can only be the normalised τ (or the right half of $(I\,\tau)$); in both readings the reported positive determinant is the determinant of $\operatorname{Im}\tau$. Positivity of the determinant of a real symmetric 4×4 matrix that arises from a period matrix, together with the leading principal minors being positive by the same Riemann relation, gives positive definiteness. Numerically $\det > 0$ at 53-bit precision is far above the error scale, so the sign is not a rounding artefact. The shape $4 \times 8 \to 4 \times 4$ is exactly $g \times 2g \to g \times g$. □

Theorem 12.87 (Archimedean invariants of the receipt)**.** *From* $d := \det(\operatorname{Im}\tau) = 0.0621660315772327$ *one has, at 40 digits,*

$$\log d = -2.7779465445343872, \qquad -\log d = 2.7779465445343872,$$
$$-\tfrac{1}{2}\log d = 1.3889732722671936, \qquad -2\log d = 5.5558930890687744.$$

Moreover $c = 0.9189385332046727$, $48c = 44.1090495938242916$, $L_2 = 10.5441334479150920$, $gL_2 = 42.1765337916603680$, $\log\Delta|_{t=2} = 121.3923901330001598$, $8\log\operatorname{rad}|_{t=2} = 8\log 42 = 29.9013569462669464$, $\Delta\gamma_A|_{\log\Delta} = 0.3633592644934122$ *and* $\Delta\gamma_A|_{8\log\operatorname{rad}} = 1.4751521034008162$. *Finally* $(-\log d)/L_2 = 0.2634589706452999$, $48c/(-\log d) = 15.8782931516838917$, $48c/L_2 = 4.1832787693469531$.

Proof. All entries are elementary evaluations of log and of rational combinations of $\log 2, \log 3, \log 7, \log\pi$, performed at 40 decimal digits. For $t = 2$, $\Delta = 2^{90}3^{36}7^{10}$, so the set of bad primes is $\{2, 3, 7\}$ and $\log\operatorname{rad} = \log 42$; this gives the stated $8\log 42$. The two $\Delta\gamma_A$ values are $48c$ divided by $\log\Delta$ and by $8\log 42$ respectively, reproducing the value 1.475152 of the previous round. The ratio $(-\log d)/L_2$ uses $L_2 = 2\log(2\pi^4)$; no other input is needed. □

Theorem 12.88 (The normalisation gap does not collapse)**.** *Let* $\Delta\gamma_A = 48c/\log\operatorname{cond}$ *be the uncertainty band arising from the two normalisations of the Faltings height, which differ by the additive constant* $c \cdot g$*. Then computing* $\det(\operatorname{Im}\tau)$*, or indeed any finite-precision Archimedean quantity, does not reduce* $\Delta\gamma_A$*. In particular* $\gamma_A \geq 0 \iff h_F \leq \frac{1}{3}\log\operatorname{cond}$ *remains undecided.*

Proof. The Archimedean part of the Faltings delta invariant is a function of the pair $(\det\operatorname{Im}\tau, |\det\Omega_1|)$, where Ω_1 is the left $g \times g$ block of the period matrix. The external run supplies only the first member of the pair, so δ_{Fal} is not yet determined. Even after $|\det\Omega_1|$ is supplied, the two height normalisations used in this volume differ by $c \cdot g$ with $c = \frac{1}{2}\log(2\pi)$; on a tower this is divided by $\hat{f} \to \infty$ and vanishes, but on a single curve $\log\operatorname{cond}$ is finite and the quotient $48c/\log\operatorname{cond}$ survives. This offset is an additive constant in the normalisation, not an error term: it is independent of the working precision of the period matrix. Quantitatively $-\log d = 2.7779$ is only 6.30% of $48c = 44.1090$, while $48c/L_2 = 4.1833$, so the normalisation constant dominates the Archimedean variation. Hence no increase of `ARCH_PREC` can send $\Delta\gamma_A$ to zero, and the criterion $\gamma_A \geq 0 \iff h_F \leq \frac{1}{3}\log\operatorname{cond}$ is untouched. □

Theorem 12.89 (Status of the five-step roadmap)**.** *Of the five steps of Section 12.12, step 1 is closed on the Sage side only; steps 2, 3, 4 and 5 remain open. Consequently problem (II), the existence of a tower with an absolutely simple* $\mathrm{End}^0 = \mathbb{Q}$*, non-*$\mathrm{GL}_2$ *factor satisfying* $\liminf w_\rho > 0$ *and* $\gamma_A < 0$*, remains open.*

Proof. Step 1 asked for a certified period matrix from two independent libraries. Sage's `RiemannSurface` now returns a 4×8 period matrix and a 4×4 Riemann matrix with positive $\det(\operatorname{Im}\tau)$, so one library is certified; Magma has no genus-4 period routine and Julia was not run, so the step is closed only on one side. Step 2 requires the wild Swan conductor at $p = 2$ (and the rescaled cluster pictures at $p = 3, 7$); step 3 is precisely the normalisation dictionary shown unresolved in Theorem 12.88; step 4 (database cross-samples) and step 5 (the Mumford–Tate corank gap $\gamma > 0$) were not attempted. Since (II) reduces to a single-factor bound $R_A \leq 3$ whose only missing input is the absolute height and the true conductor, and neither is produced here, (II) is unchanged. □

Theorem 12.90 (Two-path cross-check for Ω_1)**.** *Let* P *be the* 4×8 *period matrix,* $P = (\Omega_1\,\Omega_2)$*, and let* $\tau_{\mathrm{alt}} = \Omega_1^{-1}\Omega_2$*. If*

$$\frac{|\det(\operatorname{Im}\tau_{\mathrm{alt}}) - \det(\operatorname{Im}\tau)|}{|\det(\operatorname{Im}\tau)|} \leq 10^{-12},$$

where τ *is the matrix returned by* `riemann_matrix()`*, then the Sage-side period normalisation is self-consistent and* 0.0621660315772327 *may be used as a reproduction benchmark for a second library.*

Proof. Both quantities are determinants of the imaginary part of the same normalised period matrix, obtained along two different chains: one by Sage's internal `riemann_matrix()` and one by explicit inversion of Ω_1 followed by multiplication with Ω_2. Agreement to 10^{-12} at 53-bit working precision is far inside numerical error and therefore excludes the two systematic mistakes that could otherwise pass a single-path test: taking the right half of a $(I\,\tau)$ matrix without performing Ω_1^{-1}, and transposing the block decomposition. The left block determinant $|\det\Omega_1|$ is reported alongside, so that a second library can compare both the normalised and the unnormalised data. □

Remark 12.91 (Honesty boundary and a non-observation). The value $d = 0.0621660315772327$ satisfies $1/d = 16.0859552174830115$, which differs from 16 by a relative -0.5343%; this proximity is recorded but is *not* used to posit any structural constant. The external run was performed by the user on Sage 10.9; this sandbox has no computer algebra system, so the period matrix was not recomputed locally and the 40-digit values of Theorem 12.87 are exact transformations of the reported number. The new Ω_1 block of the script has passed a syntax check but has not been executed. Absolute Faltings height, the true conductor exponent at $p = 2$, the sign of γ_A, and the analytic input "Mumford–Tate corank ⇒ positive gap" remain open. The fourteen No-Go statements, the constants $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, and the status of *abc* as O are unchanged; this appendix is not an independent verification of Theorem A.

12.14 : completion of the Archimedean pair

Throughout this appendix $C_{7,t}$ denotes the genus-4 hyperelliptic curve of Eq. (1.1) of arXiv:2510.00093v2, with $\Delta(C_{7,t}) = 2^{36}3^{36}7^{10}t^{54}(t-1)^{12}$ and $t = 2$. We keep the notation of Appendix DDDD: $c = \frac{1}{2}\log(2\pi)$, $L_2 = 2\log(2\pi^4)$, $g = 4$, $d := \det(\operatorname{Im}\tau) = 0.0621660315772327$, and we write $w := |\det\Omega_1| = 0.000702901756037744$ for the newly reported determinant of the left $g \times g$ block of the 4×8 period matrix. All numbers below are recomputed at 40 decimal digits from the reported printed values.

Theorem 12.92 (The two-path cross-check is passed)**.** *The external Sage 10.9 run returns* $\det(\operatorname{Im}\tau) = 0.0621660315772327$ *from* `riemann_matrix()` *and* $\det(\operatorname{Im}(\Omega_1^{-1}\Omega_2)) = 0.0621660315772326$ *from explicit inversion of the left block. The reported relative discrepancy* 1.786×10^{-15} *satisfies the* 10^{-12} *criterion of Theorem 12.90, so the Sage-side period normalisation is self-consistent and d is a valid reproduction benchmark.*

Proof. Both determinants are computed from the same 4×8 period matrix $P = (\Omega_1\,\Omega_2)$: the first by the library routine, the second by forming $T_{\mathrm{alt}} = \Omega_1^{-1}\Omega_2$ and taking the determinant of $\operatorname{Im}T_{\mathrm{alt}}$. Agreement at 10^{-15} excludes the two systematic failures named in Theorem 12.90, namely omitting the Ω_1^{-1} normalisation and transposing the block decomposition. The run also reports six out of six discriminant identities and thirteen out of thirteen inline receipt fingerprints, so the exact-integer part of the script is unchanged by the v2 compatibility edits. The positivity of d is inherited from Theorem 12.86. □

Theorem 12.93 (Forty-digit Archimedean invariants of the completed pair)**.** *From d and w one has*

$$\begin{aligned}
-\log d &= 2.777946544534387217066626905024487,\\
\log w &= -7.260293425510184434611854922845621,\\
-\log d - 2\log w &= 17.298533395554756086290336750715 73,\\
\log d - 2\log w &= 11.742640306485981652157082940666 76,\\
d/w^2 &= 125824.116062526690013062471970 4807.
\end{aligned}$$

Moreover $c = 0.91893853320467274178032973640561 76$, $48c = 44.10904959382429160545582734746965$, $L_2 = 10.54413344791509201198188305374082$, $\log\Delta|_{t=2} = 121.39239013300015978888332468948786$, $8\log 42 = 29.90135694626694644734264081459106$, $\Delta\gamma_A|_{\log\Delta} = 0.36335926449341224245670652211813 09$

and $\Delta\gamma_A|_{8\log 42} = 1.47515210340081617619585054446040 7$. *Finally* $48c/(-\log d - 2\log w) = 2.5498\ldots$ *and* $48c/L_2 = 4.18327876934695310858133037605510 5$.

Proof. Every entry is an elementary evaluation of log on the two reported numbers, or a rational combination of $\log 2, \log 3, \log 7, \log \pi$, evaluated at 40 digits. For $t = 2$ one has $\Delta = 2^{90}3^{36}7^{10}$, hence the bad-prime set $\{2, 3, 7\}$ and $\log \mathrm{rad} = \log 42$; the two $\Delta\gamma_A$ values are $48c$ divided by $\log \Delta$ and by $8 \log 42$, reproducing Theorem 12.87. The relation $d/w^2 = \exp(\log d - 2\log w)$ holds by construction, and the reported $\log w = -7.2602934255101843$ agrees with the 40-digit value to one unit in the last printed place. □

Theorem 12.94 (The Archimedean ratio is not a structural constant)**.** *The quantity* d/w^2 *is not a symplectic invariant and is not an exact rational number. In particular the proximity of* d/w^2 *to* $125824 = 2^7 \cdot 983$ *is not evidence of a structural constant.*

Proof. The block Ω_1 is the matrix of a-periods of a chosen basis of holomorphic differentials against a chosen symplectic basis of $H_1(C, \mathbb{Z})$. Under $\Omega \mapsto \Omega M$ with $M \in \mathrm{Sp}(2g, \mathbb{Z})$ the left block changes by a factor of the shape $\det(C\tau + D)$, so $w = |\det \Omega_1|$ is basis dependent; hence neither w nor d/w^2 is a modular invariant and neither can appear as an intrinsic constant of the curve. Numerically $d/w^2 - 125824 = 0.11606252669\ldots$, a relative discrepancy of 9.22×10^{-7}, which is about 8×10^9 times the 53-bit rounding scale 1.11×10^{-16}; the gap is therefore genuine and not a precision artefact. This upgrades the non-observation of Theorem 12.90's remark from a numerical caution to a structural exclusion. □

Theorem 12.95 (Sharpening of Theorem 12.88)**.** *The Archimedean pair* $(\det \mathrm{Im}\,\tau, |\det \Omega_1|)$ *is now complete, yet* δ_{Fal} *and the absolute Faltings height* h_F *are still undetermined. What remains is a discrete choice of normalisation dictionary, the finite-place model (in particular the wild Swan conductor at* $p = 2$ *and a semistable model), and the translation between* δ_{Fal} *and* h_F*; no Archimedean number is missing. The conclusion of Theorem 12.88 is unchanged:* $\Delta\gamma_A = 48c/\log \mathrm{cond}$ *does not collapse and* $\gamma_A \geq 0 \iff h_F \leq \frac{1}{3}\log \mathrm{cond}$ *stays undecided.*

Proof. The premise of Theorem 12.88 was that only the first member of the pair had been supplied; Theorem 12.92 supplies the second. The remaining ambiguity is additive and discrete: the two height normalisations of this volume differ by $c \cdot g$ with $c = \frac{1}{2}\log(2\pi)$, which is independent of the period matrix and of the working precision. On a tower this offset is divided by $\hat{f} \to \infty$ and vanishes, but on the single curve $C_{7,2}$ the denominator $\log \mathrm{cond}$ is finite. Quantitatively $48c = 44.1090\ldots$ is 2.5498 times the largest Archimedean combination $-\log d - 2\log w = 17.2985\ldots$, so completing the pair cannot remove the offset. Finite-place data enter because the Faltings height of an abelian variety over a number field is a global adelic quantity: the period matrix fixes only the Archimedean contribution, while the $p = 2$ wild ramification and the special fibre of a regular model remain uncomputed. Hence the criterion is untouched. □

Theorem 12.96 (Two paths are not two libraries)**.** *Step 1 of the five-step roadmap of Appendix CCCC is closed on the Sage side in the strong sense of internal certification, but remains open as a two-library certification. Consequently the status of step 3 changes from "a missing quantity" to "a missing discrete choice together with finite-place data".*

Proof. The cross-check of Theorem 12.92 compares two computational paths inside one software stack: `riemann_matrix()` and the explicit product $\Omega_1^{-1}\Omega_2$. It certifies that the Sage computation is internally consistent, not that it has been reproduced by an independent implementation. Magma has no genus-4 period routine of this type and the Julia alternative was not run, so only one library is certified. Step 3 of the roadmap was phrased as a normalisation dictionary to be recovered once the Archimedean data were complete; by Theorem 12.95 the data are complete, so what blocks step 3 is now precisely the choice of dictionary and the finite-place model. □

Theorem 12.97 (Status of problem (II))**.** *Problem (II) remains open. Neither the sign of* $\gamma_A = 3 - R_A$ *for the generic Mumford-type fourfold, nor the absolute height* h_F*, nor the true conductor exponent is determined by the v2 run. The uncertainty band*

$$\Delta\gamma_A \in [0.3633592645,\ 1.4751521034]$$

under the two conductor proxies $\log\Delta$ *and* $8\log 42$ *is unchanged.*

Proof. By Theorem 12.95 the normalisation offset survives, and by Theorem 12.96 the finite-place conductor is still missing; these are exactly the two inputs required by the reduction of (II) to a single absolutely simple, $\mathrm{End}^0 = \mathbb{Q}$, non-$\mathrm{GL}_2$ factor with $\liminf w_\rho > 0$ and $\gamma_A < 0$. The endpoints of the band are the 40-digit values of Theorem 12.93. No inequality in this volume is strengthened by the new Archimedean number, so the fourteen No-Go statements and the constants $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$ are unchanged, as is the status of *abc* as *O*. □

Remark 12.98 (Honesty boundary, receipt scope and relative discrepancy). The external v2 script and its JSON output were published by the user as file cards outside the working sandbox; they were not mounted in the sandbox, so the 40-digit values above are exact transformations of the reported printed numbers and no field-by-field comparison with the JSON was performed. Recomputing the relative discrepancy from the two printed determinants gives $1.6085955217483012\times10^{-15}$, whereas the run reports 1.786×10^{-15}; the ratio 1.1103 is explained by the script computing the difference with untruncated internal floats while printing only sixteen significant digits, giving an internal difference of about 1.1103×10^{-16}. Both values are of the same order and both satisfy the 10^{-12} criterion; no digit-for-digit agreement is claimed. This sandbox has no computer algebra system and the period matrix was not recomputed locally. The finite-place conductor at $p = 2, 3, 7$, the absolute Faltings height, the sign of γ_A, and the analytic input "Mumford–Tate corank implies a positive gap" remain open. This appendix is not an independent verification of Theorem A.

12.15 : the first exact finite-place conductor exponent

Throughout this appendix $C_{7,t}$ denotes the genus four hyperelliptic curve of arXiv:2510.00093v2, Equation (1.1), $f(x) = t\big((t-\frac{27}{16})x^{10} - \frac{567}{64}x^9 - \frac{189}{4}tx^8 + \cdots - 28t^4x\big)$, so that $\deg f = 10 = 2g+2$, and $K = \mathbb{Q}_7$ with $v(7) = 1$. We write $\mathcal{R}$ for the ten roots, $\Delta = 2^{16}\,\mathrm{disc}(f)$, and $v_7(\Delta) = 10 + 54v_7(t) + 12v_7(t-1)$.

Theorem 12.99 (12.99: the Newton polygon was the upper hull)**.** *For every rational* t *with* $v_7(t) = v_7(t-1) = 0$*, the lower Newton polygon of* f *at 7 is* $(1,1)\to(3,0)\to(10,0)$*, giving two roots of valuation* $1/2$*, seven roots of valuation* 0*, and the root* $x = 0$ *of valuation* $+\infty$*. After scaling* f *by the 7-adic unit 32 one has* $\bar f = -x^3(x-3)^7 \in \mathbb{F}_7[x]$*, and* $v_7(\mathrm{disc}\, f) = 10$*. The values "*$0\times8,\ 1\times1$*" printed by versions 1 and 2 of the companion Sage script are the upper hull and are withdrawn.*

Proof. The Sage function `lower_hull` removed a vertex when the cross product $(x_2-x_1)(p_y-y_1)-(y_2-y_1)(p_x-x_1)$ was non-negative; the lower hull requires removal when the cross product is non-positive. Put $t = 2$ for definiteness (all six tested t give identical valuations of the coefficients). Multiplying f by 32 clears denominators without changing roots or 7-adic valuations, and gives integer coefficients $(0, -28672, 0, 55296, 43008, -24192, -48384, -27552, -6048, -567, 20)$. Their v_7 at degrees $1, 3, 4, \ldots, 10$ are $1, 0, 1, 1, 1, 1, 1, 1, 1, 0$; the lower hull is therefore $(1,1)\to(3,0)\to(10,0)$ while the upper hull is $(1,1)\to(9,1)\to(10,0)$. Reduction modulo 7 gives $-x^{10}+3x^3 = -x^3(x^7-3) = -x^3(x-3)^7$, since $x^7-3 = (x-3)^7$ in characteristic 7. Hence exactly three roots have positive valuation and seven are units, which agrees with the lower hull and excludes the upper hull. Finally $\mathrm{disc}(32f)$ has $v_7 = 10$, matching $\Delta = 2^{36}3^{36}7^{10}t^{54}(t-1)^{12}$ at $t = 2$. □

Theorem 12.100 (12.100: exact cluster picture from a resultant)**.** *Let* $R(z) = \mathrm{Res}_x(f(x), f(x+z))$*, whose roots are the* 100 *differences* $\alpha_j - \alpha_i$*. For* $v_7(t) = v_7(t-1) = 0$ *the lower Newton polygon of* R *is* $(10,10) \to (16,7) \to (58,0) \to (100,0)$*, so the ordered non-zero root differences have valuations* $1/2$ *(six pairs),* $1/6$ *(forty-two pairs) and* 0 *(forty-two pairs). Consequently the cluster picture consists of* $\mathcal{R}$ *(size* 10*, depth* 0*), a cluster* A *of size* 3 *and depth* $1/2$*, and a cluster* B *of size* 7 *and depth* $1/6$*, with no deeper nesting, and* $6 \cdot \frac{1}{2} + 42 \cdot \frac{1}{6} = 10 = v_7(\Delta)$*.*

Proof. $R(z)$ is even because the differences occur in pairs $\pm(\alpha_i - \alpha_j)$, and its coefficient of z^{10} has valuation $\sum_{i \neq j} v_7(\alpha_i - \alpha_j)$. We reconstruct R exactly by evaluating $\mathrm{Res}(f(x), f(x+k))$ at the 101 integers $k = 0, \dots, 100$ and interpolating with exact rational arithmetic; the result reproduces three further evaluations at $k = 101, 137, -5$ digit for digit, and all odd-degree coefficients vanish. The lower hull of the valuations of the coefficients is $(10,10), (16,7), (58,0), (100,0)$; the segment slopes are $-1/2$, $-1/6$ and 0 with horizontal lengths $6, 42, 42$. As a cross-check, $v_7(R(7^d u)) = \min_i(v_7(c_i) + id)$ gives $20, 30, 40$ for $d = 1, 2, 3$ and every unit $u \in \{1, \dots, 6\}$. Six rational values of t give the same hull; for $t = 8$, where $v_7(t-1) = 1$, the middle slope becomes $19/42$ and the hull becomes $(10,22) \to (16,19) \to (58,0) \to (100,0)$, again summing to $22 = v_7(\Delta)$. □

Theorem 12.101 (12.101: Hensel factorisation and wildness)**.** *Over* $\mathbb{Z}_7$ *one has* $f/\mathrm{lc}(f) = A \cdot B$ *with* $A \equiv x^3$*,* $B \equiv (x-3)^7 \pmod 7$*,* $\deg A = 3$*,* $\deg B = 7$*. Then* $v_7(\mathrm{disc}\, A) = 3$*,* $v_7(\mathrm{disc}\, B) = 7$*,* $v_7 \mathrm{Res}(A, B) = 0$*, and* $B(3+y)$ *is an Eisenstein polynomial of degree* 7*. Hence* B *is irreducible over* $\mathbb{Q}_7$*, its root field is totally ramified of degree* 7 *with residue degree* 1 *and field discriminant exponent* 7*, and the curve is wild at* $p = 7$*. Moreover* $\mathrm{Res}(B(x), B(x+z))$ *has lower hull* $(7,7) \to (49,0)$*, i.e. all forty-two ordered differences inside* B *have valuation* $1/6$*.*

Proof. Since $\gcd(x^3, (x-3)^7) = 1$ in $\mathbb{F}_7[x]$, Hensel lifting applies. We lift the factorisation to precision 7^{10} by solving, at each step k, the $\mathbb{F}_7$-linear system $sB_0 + tA_0 = (F - AB)/7^k$ with $\deg s < 3$, $\deg t < 7$; the resulting equality $F = AB \pmod{7^{10}}$ is exact. The symmetric representatives give $A = x(x^2 + 7^2 \cdot 20x + 7 \cdot 156)$ (an Eisenstein quadratic) and $B(3+y)$ with coefficients of valuations $1, 1, 3, 1, 1, 1, 2$ and leading coefficient a unit, hence Eisenstein. Discriminants and the resultant are computed from these integer polynomials and sum to 10. An Eisenstein polynomial of degree 7 has trivial index and generates a totally ramified extension of degree 7, with wild ramification because 7 divides the ramification index. The Newton polygon of $B(3+y)$ has slope $-1/7$, i.e. each single root lies at distance $7^{-1/7}$ from 3, while the conjugate differences have valuation $1/6$; both valuations live in the Galois closure whose ramification degree is $\mathrm{lcm}(7,6) = 42$, so there is no contradiction. The resultant of B with its translate gives hull $(7,7) \to (49,0)$ and total valuation $42 \cdot \frac{1}{6} = 7 = v_7(\mathrm{disc}\, B)$. □

Theorem 12.102 (12.102: the tame part is $2g = 8$)**.** *The tame conductor exponent of* $\mathrm{Jac}(C_{7,2})$ *at* 7 *equals* $n_{\mathrm{tame}} = 8$*.*

Proof. Both proper clusters A (size 3) and B (size 7) are odd, so there are no even non-übereven proper clusters and the potential toric rank is zero (Theorem 5.6 of the user guide); hence the Jacobian has potentially good reduction and the tame part is $2g = 8$. Independently, Theorem 12.3(1) of the same source gives $n_{\mathrm{tame}} = 2g - \#(U/I_K) + \#(V/I_K) + 1$ when $|\mathcal{R}|$ and $v(c)$ are both even. Here $c = 5/8$ has $v_7(c) = 0$ and $|\mathcal{R}| = 10$, so the final 1 is present. With $\tilde{\lambda}_{\mathcal{R}} = 0$, $\tilde{\lambda}_A = \frac{1}{2} \cdot 3 \cdot \frac{1}{2} = \frac{3}{4}$, $\tilde{\lambda}_B = \frac{1}{2} \cdot 7 \cdot \frac{1}{6} = \frac{7}{12}$ and all cluster stabilisers of index 1, one gets $\xi_{\mathcal{R}}(0) = 0$, $\xi_A(3/4) = 2 > \xi_A(1/2) = 1$, $\xi_B(7/12) = 2 > \xi_B(1/6) = 1$. Thus $U/I_K = \{A, B\}$ has two elements, $V/I_K = \{\mathcal{R}\}$ has one, and $n_{\mathrm{tame}} = 8 - 2 + 1 + 1 = 8$. □

Theorem 12.103 (12.103: the wild part and n_7)**.** *The wild conductor exponent is* $n_{\mathrm{wild}} = 1$*, and therefore* $\boxed{n_7 = 9}$ *is the exact conductor exponent of* $\mathrm{Jac}(C_{7,2})$ *at* $p = 7$*. In particular* $n_7 > 2g = 8$*, so the tame bound* $f_p \le 2g$ *does not apply at a wild prime.*

Proof. By Theorem 12.3(2), $n_{\mathrm{wild}} = \sum_{r \in \mathcal{R}/G_K} \big(v(\Delta_{K(r)/K}) - [K(r) : K] + f_{K(r)/K}\big)$. The three Galois orbits of roots are $\{0\}$, the two roots of the Eisenstein quadratic A/x, and the seven

roots of the Eisenstein degree-7 factor B. Their contributions are $0-1+1=0$, $1-2+1=0$ and $7-7+1=1$ respectively, using Theorem 12.101 for the last field. Hence $n_{\mathrm{wild}}=1$ and $n_7=n_{\mathrm{tame}}+n_{\mathrm{wild}}=9$. As a consistency check, the model discriminant exponent is 10, so the Ogg–Saito difference is $10-9=1$, compatible with a two-component minimal regular model; minimality of the given model is not used in the derivation of n_7. □

Theorem 12.104 (12.104: refined normalisation band)**.** *Let* $c=\frac{1}{2}\log(2\pi)$, $48c=44.1090495938\ldots$, *and* $\Delta\gamma_A=48c/\log\mathrm{cond}$. *For* $t=2$ *the exact* $p=7$ *contribution gives* $\log\mathrm{cond}\geq 9\log 7=17.513191\ldots$ *and, since the model carries one unit of redundancy at* 7, $\log\mathrm{cond}\leq\log\Delta-\log 7=119.446480\ldots$. *Therefore* $\Delta\gamma_A\in[0.36929,\ 2.5187]$.

Proof. The conductor is a product over bad primes; the newly computed exponent contributes $n_7\log 7$ and no other prime contributes a negative amount. The upper bound uses $v_7(\Delta)=10$ together with $n_7=9$, i.e. one unit of non-minimality at $p=7$, and the archimedean quantity $\log\Delta=121.392390\ldots$ from Appendix EEEE. Dividing $48c$ by the two bounds gives the stated interval. Neither endpoint determines the absolute Faltings height or the sign of γ_A. □

Remark 12.105 (Retractions, scope and honesty boundary). Four earlier statements are corrected. (i) The claim that wildness at 7 follows from the degree-7 totally ramified extension in Lemma 4.9 of arXiv:2510.00093v2 is withdrawn as a reason: that extension is a base change in the parameter $t'=s^7$ describing the degenerate family fibre at $t=1$, not 7-adic ramification; the conclusion that $p=7$ is wild is nevertheless correct and is proved here by the Eisenstein factor B. (ii) The bound $f_p\leq 2g=8$ quoted for the non-wild bad primes does not apply at $p=7$, where $n_7=9$. (iii) The pre-run report of integral 7-adic roots is withdrawn with Theorem 12.99. (iv) Route-map step (ii) of Appendix EEEE is closed at $p=7$; only $p=2$ (even residue characteristic) and $p=3$ remain. All computations are exact integer or rational arithmetic plus symbolic resultants and discriminants; no computer algebra system and no period matrix were used, and the external theorems of M2D2, DDMM, Bisatt and the user guide were checked verbatim but not reproved. The primes 2 and 3, the absolute Faltings height, the sign of γ_A, and the Mumford–Tate corank gap remain open. The fourteen No-Go statements, the constants $(N/M_5)^{24+\varepsilon}$, $E(c)=12/(5-18c)$, $47/3$, $C\geq 36/7>3$, and the status of *abc* as O are unchanged; this appendix is not an independent verification of Theorem A.

12.16 : $p=3$, a p^2-level wild place and the odd conductor spectrum

Throughout this appendix $C_{7,t}$ and f are as in Appendix FFFF, and $K=\mathbb{Q}_3$ with $v(3)=1$. Put $\Delta=2^{16}\,\mathrm{disc}(f)$, so $v_3(\Delta)=36+54v_3(t)+12v_3(t-1)$. All computations are exact integer or rational arithmetic with symbolic resultants, discriminants and factorisations; the six rational parameters tested are $t=2,8,3/2,1/2,101/100,-1$.

Theorem 12.106 (12.106: an exact factor, no Hensel lifting)**.** *For every rational* t *the constant coefficient of* f *vanishes identically, so* $f(x)=x\cdot B(x)$ *with* $B=f/x\in\mathbb{Q}[x]$ *of degree nine. The polynomial* B *is irreducible over* $\mathbb{Q}$ *for all six tested* t, *and* $v_3(\mathrm{disc}\,B)=36$ *whenever* $v_3(t)=v_3(t-1)=0$. *Consequently the* 3*-adic factorisation is exact and requires no Hensel lifting, in contrast with the* 7*-adic factorisation of Theorem 12.101.*

Proof. The coefficient list of $32f$ at $t=2$ is $(0,-28672,0,55296,43008,-24192,-48384,-27552,-6048,-567,20)$, and the entry of degree zero is zero for every t because every monomial of Equation (1.1) contains at least one factor x. Division by x is therefore exact over $\mathbb{Q}$, and gives a degree-nine polynomial with leading coefficient $32t(t-27/16)$. Its factorisation over $\mathbb{Q}$ was computed symbolically and consists of a single irreducible factor for each of the six parameters. For $t=2$ one has $B(0)=-28672=-2^{12}\cdot 7$, a 3-adic unit, hence $\mathrm{disc}(f)=\mathrm{disc}(B)\cdot B(0)^2$ gives $v_3(\mathrm{disc}\,B)=v_3(\mathrm{disc}\,f)=36$. The last value follows from $\Delta=2^{36}3^{36}7^{10}t^{54}(t-1)^{12}$ and $v_3(2^{16})=0$. □

Theorem 12.107 (12.107: reduction and the exceptional parameter)**.** *For $v_3(t) = v_3(t-1) = 0$ one has $\bar{f} = \pm x(x+1)^9 \in \mathbb{F}_3[x]$ and $v_3(\operatorname{disc} f) = 36$. For $t = 3/2$ the reduction is identically zero, $v_3(\operatorname{disc} f) = 90$, $v_3(\operatorname{lc} f) = 2$, and $v_3(\operatorname{disc} f^{\mathrm{mon}}) = 90 - 18 \cdot 2 = 54$.*

Proof. In $\mathbb{F}_3$ the Frobenius identity gives $x^9 + 1 = (x+1)^9$, and $32f \equiv \pm(x^{10} + x) = \pm x(x+1)^9$; the sign depends on t mod 3 and is irrelevant. The discriminant formula of Appendix FFFF gives 36 because $v_3(t) = v_3(t-1) = 0$ and $v_3(7) = 0$. For $t = 3/2$ one has $c_{10} = t(t - 27/16) = -9/32$, so $v_3(\operatorname{lc} f) = 2$; each of the ten coefficients of $32f$ is divisible by 3, which explains $\bar{f} = 0$. For a degree-n polynomial, $\operatorname{disc}(f/a_0) = a_0^{-2n+2} \operatorname{disc}(f)$, whence v_3 drops by $(2 \cdot 10 - 2) \cdot 2 = 36$, leaving $90 - 36 = 54$. □

Theorem 12.108 (12.108: a single-segment root-difference spectrum)**.** *Let $R(z) = \operatorname{Res}_x(f(x), f(x+z))$. For $v_3(t) = v_3(t-1) = 0$ the lower Newton polygon of R is $(10, 36) \to (82, 0) \to (100, 0)$, so the ordered root differences have valuation $1/2$ with multiplicity 72 and valuation 0 with multiplicity 18, and $72 \cdot \frac{1}{2} + 18 \cdot 0 = 36 = v_3(\Delta)$. Moreover $v_3(R(3^d u)) = 46, 56, 66$ for $d = 1, 2, 3$ and every 3-adic unit u tested, agreeing with $\min_i(v_i + id)$.*

Proof. The polynomial R has degree 100 and is even; its values at the 101 integer points $z = 0, \ldots, 100$ were computed as exact integer resultants and interpolated, with the resulting polynomial checked at $z = 101, 137, -5$. The 3-adic valuations of its nonzero coefficients give the lower hull $(10, 36) \to (82, 0) \to (100, 0)$; the segment from $(10, 36)$ to $(82, 0)$ has slope $-36/72 = -1/2$ and length 72, and the final segment has slope zero and length 18. Since $v_3(\operatorname{lc} f) = 0$ for these t, the initial valuation 36 equals $v_3(\operatorname{disc} f)$. The Legendre values follow by evaluating the same resultants at $z = 3^d u$ for $u \in \{1, 2, 4, 5, 7, 8\}$. For $t = 1/2$ the polynomial $B(y-1)$ has a two-segment Newton polygon $(0, 5) \to (1, 4) \to (9, 0)$, but the difference spectrum is still single-segment: the deeper root has valuation 1, the other eight have valuation $1/2$, and every cross difference takes the smaller value $1/2$, giving $56 + 16 = 72$ ordered pairs. At $d = 0$ the Legendre prediction fails for $t = 1/2$ (values $5, 5, 6, 6, 5, 5$ against prediction 0); this is coefficient cancellation at a unit argument and does not affect $d \geq 1$. □

Theorem 12.109 (12.109: the residual polynomial is inseparable)**.** *Let $g(y) = B(y-1)$. For $t = 2$ the lower Newton polygon of g is the single segment $(0, 3) \to (9, 0)$ of slope $-1/3$; among the coefficients of degrees $0, 3, 6, 9$ only those of degrees 0 and 9 lie on the polygon. Hence the residual polynomial is a binomial $\alpha T^3 + \beta$ over $\mathbb{F}_3$, which is inseparable since $T^3 + \beta/\alpha = (\gamma T + \delta)^3$ in characteristic three. In particular the wild inertia has p^2-level data: the root field $K(r)/K$ has degree nine, its ramification index is divisible by three, and the single-layer wild contribution $e - e + 1 = 1$ of Theorem 12.103 cannot be transplanted.*

Proof. Exact Taylor expansion of B at $x = -1$ gives, at degrees $0, \ldots, 9$, the 3-adic valuations $(3, 5, 6, 3, 4, 3, 2, 2, 2, 0)$ after clearing the unit 32. The lower hull is $(0, 3) \to (9, 0)$; the polygon equals $2, 1, 0$ at degrees $3, 6, 9$, while the coefficients there have valuations $3, 2, 0$, so only degrees 0 and 9 contribute. Rescaling by 3^v at each vertex produces a residual polynomial supported in degrees 0 and 3, of the shape $\alpha T^3 + \beta$ with $\alpha\beta \neq 0$. In $\mathbb{F}_3$ the cube map is the Frobenius, so such a binomial has a triple root and is inseparable. An inseparable residual polynomial forces higher-order Newton polygons to determine the ramification filtration; since $\deg B = 9 = 3^2$ is irreducible, G_K acts transitively on the nine roots and $[K(r) : K] = 9$ with $3 \mid e$. This is strictly stronger wild behaviour than the Eisenstein degree-seven factor of Theorem 12.101, whose residual polynomial is separable. □

Theorem 12.110 (12.110: the tame part equals $2g + 1$)**.** *For the cluster picture consisting of $\mathcal{R}$ of size ten with $d = 0$, a principal odd cluster S of size nine with $d_S = 1/2$, and nine singletons, one has $\tilde{\lambda}_{\mathcal{R}} = 0$, $\tilde{\lambda}_S = 9/4$, $U/I_K = \{S\}$, $V/I_K = \{\mathcal{R}\}$, and the parity bonus $+1$. Hence $n_{\mathrm{tame}} = 2g - \#U + \#V + 1 = 8 - 1 + 1 + 1 = 9$.*

Proof. By Theorem 12.3 of the user guide, $\tilde\lambda_{\mathfrak s} = \frac12(v(c)+|\tilde{\mathfrak s}|d_{\mathfrak s}+\sum_{r\notin\mathfrak s} d_{\{r\}\wedge\mathfrak s})$. Here $v(c) = v_3(20) = 0$, $\mathcal R$ has one odd child S and $d_{\mathcal R} = 0$, so $\tilde\lambda_{\mathcal R} = 0$; S has nine odd children, $d_S = 1/2$, and the only root outside S meets it in $\mathcal R$, of depth zero, so $\tilde\lambda_S = \frac12\cdot 9\cdot\frac12 = 9/4$. With $\xi(a) = \max\{-\operatorname{ord}_2(a), 0\}$ one gets $\xi(\tilde\lambda_{\mathcal R}) = 0 \le \xi(d_{\mathcal R}) = 0$, so $S \in U$, while $\xi(\tilde\lambda_S) = 2 > \xi(d_S) = 1$, so no singleton enters U. The only proper cluster with $\xi(\tilde\lambda) = 0$ is $\mathcal R$, so $V = \{\mathcal R\}$; $|\mathcal R| = 10$ and $v(c) = 0$ are even, giving the bonus. Substituting $g = 4$ gives 9. At $p = 7$ one had $\#U = 2$ and hence $n_{\mathrm{tame}} = 8$; the difference is exactly the loss of one principal odd cluster. □

Theorem 12.111 (12.111: the odd conductor spectrum and the normalisation band)**.** *Let* $i = v_3([\mathcal O_{K(r)} : \mathbb Z_3[\alpha]]) \ge 0$ *and let* $f_{\mathrm{res}} \in \{1, 3\}$ *be the residue degree of* $K(r)/K$*. Then*

$$n_{\mathrm{wild}} = 27 - 2i + f_{\mathrm{res}}, \qquad n_3 \in \{37 - 2i,\ 39 - 2i\},$$

so n_3 *is odd and is determined by two integers that require local field data. Independently of these,* $n_{\mathrm{tame}} = 9$ *gives* $\log\operatorname{cond} \ge 9\log 3 + 9\log 7 = 27.400702\ldots$ *and* $\Delta\gamma_A = 48c/\log\operatorname{cond} \le 1.609778$*.*

Proof. Since B is irreducible of degree nine, the ten roots form two G_K-orbits: the singleton $\{0\}$ with trivial contribution $0-1+1 = 0$, and the nine roots of B, a single orbit with $[K(r) : K] = 9$. For a monogenic generator, $\operatorname{disc}(B) = \Delta_{K(r)/K}\cdot i^2$ with $v_3(\operatorname{disc} B) = 36$; Theorem 12.3(2) then gives $n_{\mathrm{wild}} = v_3(\Delta_{K(r)/K}) - 9 + f_{\mathrm{res}} = 36 - 2i - 9 + f_{\mathrm{res}}$. The residue field of an inseparable residual polynomial is $\mathbb F_3$ or its cubic extension, hence $f_{\mathrm{res}} \in \{1, 3\}$, and adding $n_{\mathrm{tame}} = 9$ yields the stated odd spectrum. If one imposes the Ogg–Saito type bound $n_p \le v_p(\Delta)$ used in Appendix FFFF, then $n_3 \le 36$ forces $i \ge 1$ for $f_{\mathrm{res}} = 1$ and $i \ge 2$ for $f_{\mathrm{res}} = 3$. Finally, the conductor contains the tame contributions at 3 and the exact $n_7 = 9$ at 7, so $\log\operatorname{cond} \ge 9\log 3 + 9\log 7$; with $48c = 44.1090495938\ldots$ from Appendix EEEE this gives the stated upper bound on $\Delta\gamma_A$, improving the bound 2.5187 of Theorem 12.104. Neither i nor the ramification breaks are determined by resultants alone, so n_3 itself remains open. □

Remark 12.112 (Retractions, scope and honesty boundary). Four earlier statements are corrected. (i) The claim in Appendix FFFF that $p = 3$ and $p = 2$ are both hard cores of the same kind is refined: at $p = 3$ all finite-place algebra, the reduction, the cluster spectrum and the tame part are exact; only the local field discriminant and the ramification filtration remain. (ii) The expectation that a Hensel lifting is needed at $p = 3$ is withdrawn by Theorem 12.106. (iii) The $p = 7$ wild shape $7 - 7 + 1 = 1$ cannot be transplanted, by Theorem 12.109. (iv) The Legendre cross-check is valid for $z \to 0$ only, that is for $d \ge 1$; the $d = 0$ discrepancy at $t = 1/2$ is coefficient cancellation, not a contradiction. The bound $n_p \le v_p(\Delta)$ is quoted relative to the sources already used in Appendix FFFF and is not reproved. All computations are exact; no computer algebra system, no period matrix and no local field engine were used, and the external theorems were checked verbatim but not reproved. The integers i and f_{res}, the absolute Faltings height, the sign of γ_A, the prime $p = 2$ and the Mumford–Tate corank gap remain open. The fourteen No-Go statements, the constants $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5 - 18c)$, $47/3$, $C \ge 36/7 > 3$, and the status of *abc* as O are unchanged; this appendix is not an independent verification of Theorem A.

12.17 : $p = 2$ and the collapse of the two long-term directions

Throughout, $C_{7,t}$ and f are as in Appendices FFFF and GGGG, and $K = \mathbb Q_2$ with $v(2) = 1$. Write $\Delta = 2^{16}\operatorname{disc}(f)$, so $v_2(\Delta) = 36 + 54v_2(t) + 12v_2(t-1)$. All computations in this appendix are exact integer or rational arithmetic; the six rational parameters tested are $t = 2, 8, 3/2, 1/2, 101/100, -1$. No computer algebra system, no period matrix and no local field engine were used.

Theorem 12.113 (12.113: the exact 2-adic spectrum and a degree drop)**.** *Let* F_0 *be the primitive integral model of* f*, obtained by clearing denominators and dividing by the largest common power of* 2*. For each of the six tested parameters one has* $F_0 \bmod 2 = x^9$*, except at* $t = 101/100$ *where* $F_0 \bmod 2 = x^9 + x$*; in all cases* $\deg(F_0 \bmod 2) = 9$*, a drop of one from* $\deg f = 10$*. The 2-adic*

Newton polygon of F_0 has vertices $(1,h) \to (9,0) \to (10,2)$, with $h \in \{0,4,8,12,20\}$, hence the root valuation spectrum is $s \times 8$ together with $(-2) \times 1$, where $s = h/8 \in \{0, 1/2, 1, 3/2, 5/2\}$. For $t = 2$ one has $h = 12$, $s = 3/2$, $v_2(\operatorname{disc} f) = 74$ and $v_2(\Delta) = 90$.

Proof. Clearing denominators at $t = 2$ gives the coefficient list of $32f$ equal to $(0, -28672, 0, 55296, 43008, -24192, -48384, -27552, -6048, -567, 20)$. The minimum 2-adic valuation of a nonzero coefficient is 0, attained at degree 9 only, so $F_0 \bmod 2 = x^9$ and the leading coefficient 20 is even. The lower convex hull of $(i, v_2(F_{0,i}))$ is $(1,12) \to (9,0) \to (10,2)$: the slopes are $-3/2$ of length 8 and $+2$ of length 1, giving root valuations $3/2$ and -2. For the other parameters the same computation gives $h = 20, 4, 4, 0, 8$ respectively. The identity $\operatorname{disc}(f) = \operatorname{disc}(32f) \cdot 32^{-18}$ gives $v_2(\operatorname{disc} f) = v_2(\operatorname{disc}(32f)) - 90 = 74$, and $v_2(\Delta) = 16 + 74 = 90$. A fraction $1/3$ recorded for $p = 2$ in Appendix FFFF is absent from all six spectra and is withdrawn. □

Theorem 12.114 (12.114: the resultant spectrum and the Artin–Schreier offset)**.** *Put $R(z) = \operatorname{Res}_x(f(x), f(x+z))$. For $t = 2$ the exact 101-point interpolation of R has 2-adic lower hull $(10,168) \to (66,28) \to (82,4) \to (100,40)$, hence root-difference spectrum $\frac{5}{2} \times 56$, $\frac{3}{2} \times 16$ and $(-2) \times 18$, with sum 128. The theoretical value $v_2(\operatorname{disc} f) - 18\, v_2(\operatorname{lead} f) = 74 + 54 = 128$ agrees exactly. The 18 ordered pairs carrying valuation 0 at $p = 3$ in Theorem 12.107 carry valuation -2 at $p = 2$.*

Proof. The resultant is a polynomial of degree 100 with integer coefficients; evaluating it at $z = 0, \dots, 100$ and interpolating with exact divided differences reproduces the values at $z = 101, 137, -5$ and shows that R is even. The Newton polygon of the coefficient valuations has the stated hull; the sum of slope times length equals the sum of $v_2(r_i - r_j)$ over ordered pairs, which by the discriminant formula equals $v_2(\operatorname{disc} f) - (2 \deg f - 2) v_2(\operatorname{lead} f)$. Here $\operatorname{lead} f = 5/8$, so $v_2(\operatorname{lead} f) = -3$ and the value is $74 + 54 = 128$. At $p = 3$ the corresponding 18 pairs have valuation 0 because the deep root and the single root are 3-adically separated by a unit; at $p = 2$ the degree drop of Theorem 12.113 shifts them by -2. □

Theorem 12.115 (12.115: Kummer theory fails, so the spectrum carries no conductor)**.** *In residue characteristic 2 the Artin–Schreier form $y^2 + Q(x)y = P(x)$ replaces the Kummer form $y^2 = f(x)$: indeed $\partial(y^2 - F)/\partial y = 2y = 0$ in characteristic 2, so the quadratic cover is inseparable in the y-direction, and $F_0 \bmod 2 = x^9$ is not a square with $\deg(F_0 \bmod 2) = 9 < 10$. Consequently the cluster picture and the conductor formula of Theorem 12.3 of the user's guide, used at $p = 7$ and $p = 3$, are not defined at $p = 2$. The exact 2-adic spectra of Theorems 12.113 and 12.114 therefore do not determine the conductor exponent n_2.*

Proof. The paper of Dokchitser, Dokchitser, Maistret and Morgan restricts throughout to local fields of odd residue characteristic, and the conductor-discriminant inequality of Obus and Srinivasan is stated for perfect residue fields of characteristic not equal to 2, with the explicit comment that the result extends the earlier inequalities away from residue characteristic 2. Their proofs use the Kummer 2-torsion of the Jacobian, which has no integral model in characteristic 2. The inseparability identity and the degree drop above are the concrete obstruction on $C_{7,t}$. The present appendix quotes these statements verbatim and does not reprove them. □

Theorem 12.116 (12.116: the model-side tool exists; n_2 collapses to $a \in \{7,8\}$)**.** *Let A be the matrix of integral differentials of the minimal regular model, as in the paper on models and integral differentials of hyperelliptic curves, and put $a = v_2(\det A)$. With $g = 4$ and $\Delta = v_2(\operatorname{disc} f) = 74$, the quantity $\Delta_{\min} = g\Delta - (8g+4)a$ equals $296 - 36a$. Assuming Conjecture 1.4 of that paper, namely $\operatorname{Art}(C/\mathcal{O}_K) \le \Delta_{\min}$, and using $\Delta_{\min} \le \Delta_{C/\mathcal{O}_K} \le \Delta$ together with $\Delta_{\min} \ge 0$, one obtains $a \in \{7,8\}$, $\Delta_{\min} \in \{44, 8\}$ and $n_2 \le 44$.*

Proof. The Saito–Liu identity $-\operatorname{Art}(\mathcal{X}/S) = n(\mathcal{X}_s) - 1 + f$ and Liu's genus-two conductor-discriminant inequality are valid without a restriction on the residue characteristic; this is the

same identity quoted in Appendix AAA from Remark 2.5 of the source text. The paper cited above gives an explicit formula for $|\det A_v|_v$ at every place, defines $\Delta_{\min} = g\Delta - (8g+4)v(\det A)$, and states Conjecture 1.4; it also proves the conjecture in genus two via Liu and in the semistable case via Maugeais. Inserting $g = 4$ and $\Delta = 74$ gives $296 - 36a$. The chain $\Delta_{\min} \le \Delta_{C/\mathcal{O}_K} \le 74$ forces $a \ge 37/6$, and nonnegativity of $\Delta_{\min}$ forces $a \le 74/9$; hence a is 7 or 8. Since $n_2 \le -\operatorname{Art} \le \Delta_{\min}$, the bound follows. This is conditional on Conjecture 1.4, which is not proved for $g = 4$ at $p = 2$. □

Theorem 12.117 (12.117: $\gamma \ge 0$ closes (II), $\gamma > 0$ does not)**.** *Let $\gamma = 3 - R = 4 - 12\kappa$, so that the eight anchors of Appendices BBB, UUU and MMM give the exact values $3, 5/2, 2, 1, 3/2, 1/2, 0, -1$, with $\gamma = 0$ precisely at the ball-quotient equality orbit $\kappa = 1/3$ and $\gamma = -1$ at the saturated model $\kappa = 5/12$. A theorem of the form "Mumford–Tate corank implies $\gamma > 0$" excludes only the generic class with full Mumford–Tate group, while the unique dangerous class of (II), namely the Hodge-special Shimura-intersect-Torelli locus of Appendices VVV and QQQ, lies at $\gamma = 0$. Hence $\gamma > 0$ is insufficient to close (II), whereas $\gamma \ge 0$ for all coranks is exactly the arithmetic Miyaoka–Yau inequality $\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g)$ and does close it.*

Proof. The identity $\gamma = 4 - 12\kappa$ is verified by exact rational arithmetic on the eight anchors $1/12, 1/8, 1/6, 1/4, 5/24, 7/24, 1/3, 5/12$. By the definition of (II) in Appendix QQQ, a counterexample requires a single absolutely simple factor with $R_\rho > 3$, equivalently $\gamma_\rho < 0$. The generic class has full Mumford–Tate group and corank 0; along a Hodge-special locus the corank increases and γ decreases to 0 at the locally symmetric equality orbit, as recorded in Appendix TTT. Consequently an implication producing $\gamma > 0$ can only hold away from the equality orbit and leaves the dangerous class untouched. Conversely $\gamma \ge 0$ for every factor gives $R \le 3$ and $\kappa \le 1/3$ by the convexity argument of Appendix EE. The energy identity proposed in Appendix TTT, with positive definite $\mathcal{E}_\infty = \int |\operatorname{Ric}(\omega) + \omega|^2 \omega^2$, has nonnegative right-hand side and vanishing precisely on ball quotients, hence supplies $\gamma \ge 0$ rather than $\gamma > 0$. □

Theorem 12.118 (12.118: the two long-term directions are one)**.** *The direction "new tools for $p = 2$" reduces, by Theorem 12.116, to a single unproved inequality plus one integer $a \in \{7, 8\}$; the direction "Mumford–Tate corank gap" reduces, by Theorem 12.117, to the arithmetic Miyaoka–Yau inequality itself. Both therefore reduce to the same missing analytic input: an arithmetic rigidity statement for the ball-quotient equality orbit. The three bad places now have $n_7 = 9$, $n_3 \in \{37 - 2i, 39 - 2i\}$ and $n_2 \le 296 - 36a$ with $a \in \{7, 8\}$, and $\log \mathrm{cond} \ge 27.400702 + n_2 \log 2$ gives $\Delta\gamma_A = 44.10905/\log \mathrm{cond} \le 1.609778$, improving to 0.761825 when $n_2 = 44$.*

Proof. The first two conductor values are Theorems 12.103 and 12.109. The third is Theorem 12.116. The numerical bound uses $9\log 3 + 9\log 7 = 27.400702$ from the tame parts at $p = 3$ and $p = 7$, and $48c = 44.1090495938$ from Appendix DDDD. The Parshin inequality (9) is a conjecture for arithmetic surfaces; the Yuan–Zhang Corollary E is in the opposite direction and writes the conjecture with an unspecified absolute constant A and a continuous archimedean function; Soulé states explicitly that no arithmetic analogue of Ω^1_X is known and that only lower bounds for successive minima are obtained; the arithmetic Yau–Tian–Donaldson equivalence remains unresolved and the Hattori–Odaka results characterise minimisers without giving the upper bound coefficient 3. Hence no external theorem supplies $\gamma \ge 0$. □

Remark 12.119 (Retractions, scope and honesty boundary). Five statements are corrected. (i) The fraction $1/3$ recorded at $p = 2$ in Appendix FFFF is withdrawn; the correct family is $s = h/8$. (ii) The claim that $p = 2$ requires new tools is refined: the model-side tool exists since 2020 and the gap is Conjecture 1.4 at genus four. (iii) The claim that $p = 2$ and $p = 3$ are hard cores of the same kind is refined: $p = 2$ collapses to $a \in \{7, 8\}$. (iv) The claim that the Mumford–Tate corank gap $\gamma > 0$ is the unique new theorem closing (II) is withdrawn; the required statement is $\gamma \ge 0$, which is weaker but covers the dangerous class. (v) The constant 3 is a transplant target from the geometric Miyaoka–Yau inequality and from the inversion in

Appendix MMM, not a constant fixed by the arithmetic conjecture, whose absolute constant is left unspecified. All twelve external sources were checked verbatim and not reproved; the bound $n_2 \leq 44$ is conditional on a conjecture; the matrix A was not computed at $p = 2$; the absolute Faltings height, the true conductor exponents, the value of a and the sign of γ_A remain open. The fourteen No-Go statements, the constants $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$, and the status of *abc* as O are unchanged; this appendix is not an independent verification of Theorem A.

12.18 : an attempted proof of the full dossier cut

We keep the notation of Appendices BBB, EE, HHH, MMM, TTT, VVV, QQQ, FFFF, GGGG and HHHH. The dossier is the twelve verbatim sources collected in Section 12.17, whose cross-check isolates two gaps: Conjecture 1.4 at genus four in residue characteristic two, and the nonnegativity of $\gamma = 3 - R = 4 - 12\kappa$. By Theorem 12.118 these two are one analytic input. In this appendix the resulting chain is written out as eight lemmas and each lemma is either proved relative to standard citations, proved conditionally, or shown to fail. All exact numerical checks use rational arithmetic only; no computer algebra system and no period matrix were used.

Theorem 12.120 (12.120: the precise geometric error identity)**.** *Let X be a minimal complex surface of general type and let ω be the Kähler–Einstein metric on its canonical model, whose existence follows from Aubin and Yau. Then*

$$3c_2(X) - c_1^2(X) = \mathcal{E}_\infty(X) + k(X),$$

where $\mathcal{E}_\infty = \int_X |\operatorname{Ric}(\omega) + \omega|^2 \omega^2 \geq 0$ is the Guggenheimer integral measuring the deviation from constant holomorphic sectional curvature, and $k(X) \geq 0$ is the contribution of the (-2)-curves, strictly positive exactly when such curves exist, equivalently when the canonical bundle is not ample. Equality $3c_2 - c_1^2 = k(X)$ holds precisely when the universal cover of the smooth part of the canonical model is the complex hyperbolic ball minus a discrete set; in particular $3c_2 = c_1^2$ holds only for ball quotients.

Proof. Hirzebruch's 1985 survey records that Guggenheimer's identity expresses $3c_2 - c_1^2$ as an integral with nonnegative integrand for the Einstein–Kähler metric, and that Yau's equality case forces the universal cover to be the ball. Kobayashi's 1984 theorem refines this to $3c_2 - c_1^2 \geq k(X) \geq 0$, identifies $k(X)$ as the (-2)-curve contribution, and gives the stated equality condition for the smooth part of the canonical model. Combining the two gives the displayed decomposition. The energy identity proposed in Appendix TTT therefore has three arithmetic counterparts (archimedean continuous, archimedean discrete, finite), not two: the discrete term is already present on the geometric side. These statements are quoted verbatim and not reproved. □

Theorem 12.121 (12.121: the eight-lemma chain and a conditional theorem)**.** *Write the full cut as the chain*

$$\textit{finite models} \to \textit{adelic energy} \to \gamma \geq 0 \to R \leq 3 \to \kappa \leq \tfrac{1}{3} \to \textit{(II)}.$$

Lemmas L1 (Theorem 12.120), L2 (the dictionary of Appendix HHHH) and L6 (the finite discrete energy of Appendices KKK and AAA) are unconditional relative to citations. Lemmas L3 (adelic $\widehat{c}_2$), L4 (the constant 3), L5 (placewise reduction), L7 (Yau–Tian–Donaldson) and L8 (the dangerous class) fail. The failures collapse to three independent statements: (α) an upper bound for the archimedean summand of the adelic $\widehat{c}_2$; (β) the arithmetic constant $A = 3$; (γ) an integral adelic definition immune to the positive-characteristic counterexamples. If $(\alpha), (\beta), (\gamma)$ hold then $\gamma \geq 0$ for every absolutely simple factor and (II) is closed. On the already proved classes (CM, GL_2, Kummer, modular towers) L6 supplies (β) and (γ) in part and $\gamma \geq 0$ holds unconditionally.

Proof. L1 is Theorem 12.120. L2 is the exact dictionary $\gamma = 3 - R = 4 - 12\kappa$ verified on the eight anchors $1/12, 1/8, 1/6, 1/4, 5/24, 7/24, 1/3, 5/12$ by rational arithmetic, together with $12\kappa = \sigma(1 + R) + \varepsilon_{\mathrm{ar}}$ from Appendix HHH. The minimum of γ on the seven already proved anchors is 0, attained at the ball-quotient orbit $\kappa = 1/3$; the saturated model of Appendix MMM gives $\gamma = -1$. L6 is the alignment $a_p = \mathrm{Sw}_p + \mathrm{rk}_{\mathrm{Lap}}\, L(\Gamma_p)$ of Appendix KKK and the closed form for $\sigma - 1$ of Appendix AAA, whose values $2.407236166552 \times 10^{-8}$ and $4.735941304093 \times 10^{-27}$ are reproduced exactly; the (-2)-curve term $k(X)$ corresponds to the chains and thickness of the dual graph, so the finite discrete energy is nonnegative and already computed.

L3 fails because Soulé states that no arithmetic analogue of Ω^1_X is known, that the archimedean summand of $\widehat{c}_2$ is not easy to bound from above, and that only lower bounds for successive minima are available; the Bost–Gillet–Soulé and Faltings $\widehat{c}_2$ carries an analytic torsion correction, so the normalising constant is undetermined. L4 fails cyclically: Parshin's inequality (9) is a conjecture for arithmetic surfaces, and Yuan–Zhang Corollary E is in the opposite direction and writes the conjecture with an unspecified absolute constant A; hence 3 is a transplant target, not an arithmetic theorem. L5 fails because the Bogomolov–Miyaoka–Yau inequality is false in positive characteristic, by the surfaces of Lang and Easton, so no placewise reduction can produce the finite energy. L7 fails because Odaka's arithmetic Yau–Tian–Donaldson equivalence is unresolved, Hattori–Odaka characterise minimisers without giving the coefficient 3, Yuan–Zhang give only uniqueness of the non-archimedean Calabi metric, and the rank-two hermitian Bogomolov inequality of Moriwaki gives $4\widehat{c}_2$. L8 fails by Appendices VVV, QQQ and WWW: Ax–Schanuel gives finiteness of special points, the Qiu insertion and spectral conservation change only weights, and (T2) and (T3) are mutually exclusive; the dangerous Hodge-special class lies at $\gamma = 0$, where any implication producing $\gamma > 0$ is vacuous.

Conversely, if $(\alpha), (\beta), (\gamma)$ hold, the adelic identity $3\widehat{c}_2 - \widehat{c}_1^2 = \kappa_E(\mathcal{E}_\infty + k_{\mathrm{fin}})$ has nonnegative right-hand side by L1 and L6, which by Noether is exactly $\bar{\omega}^2 \le 3\Delta_{\mathrm{Del}} + O(g)$; by L2 this is $R \le 3$, hence $\gamma \ge 0$ and (II) follows from the convexity argument of Appendix EE. □

Theorem 12.122 (12.122: the normalisation of $\Delta_{\min}$)**.** *In the source paper on models and integral differentials of hyperelliptic curves one has*

$$\Delta_{\min} = \Delta - \frac{8g+4}{g}\, v(\det A),$$

so for $g = 4$ and $a = v_2(\det A)$ one has $\Delta_{\min} = \Delta - 9a$. The quantity $g\Delta - (8g+4)a$ used in Appendix HHHH equals $g\,\Delta_{\min}$. Consequently the bound $a \ge 7$, derived there from $296 - 36a \le 74$, is withdrawn: with the correct normalisation $\Delta_{\min} \le \Delta$ holds for every $a \ge 0$ and gives no lower bound. One has $a \le \lfloor \Delta/9 \rfloor$, equal to 8 for $\Delta = 74$ and 10 for $\Delta = 90$, and conditionally on Conjecture 1.4

$$n_2 \le \Delta_{\min} = \Delta - 9a \le \Delta \in \{74, 90\}.$$

Proof. The definition is quoted verbatim from the source; its Conjecture 1.4 reads $\mathrm{Art}(C/\mathcal{O}_K) \le \Delta_{\min}$, and it states $\Delta_{\min} \le \Delta_{C/K}$. Multiplying the definition by $g = 4$ gives $4\Delta - 36a$, which is the expression $296 - 36a$ of Appendix HHHH for $\Delta = 74$. The inequality $4\Delta - 36a \le \Delta$ is therefore four times stronger than the source's $\Delta_{\min} \le \Delta_{C/K} \le \Delta$ and is not available. Rational arithmetic on $a = 0, \dots, 11$ for both normalisations confirms the constant ratio 4, and nonnegativity of $\Delta_{\min}$ gives the stated upper bounds on a. The two values of Δ correspond to $v_2(\mathrm{disc}\, f) = 74$ and $v_2(2^{16}\,\mathrm{disc}\, f) = 90$; which one matches the source's phrase "discriminant associated with this equation" is a normalisation issue left open. The bound $n_2 \le 44$ of Appendix HHHH is withdrawn. □

Theorem 12.123 (12.123: the scope of the model-side tool at $p = 2$)**.** *Theorem 1.2 of the same source requires the curve to be y-regular over the maximal unramified extension and its cluster picture to be almost rational. The source states that $p \ne 2$ implies y-regularity, gives only examples at $p = 2$, and records that the construction of Liu and Lorenzini via double covers of $\mathbb{P}^1$*

has not been generalised to residue characteristic two. Its Conjecture 1.4 is proved only in genus two by Liu and in the semistable case by Maugeais. Moreover the source uses the positive Artin conductor $\mathrm{Art}(C/\mathcal{O}_K) \le \Delta_{C/K}$*, whereas Saito–Liu writes* $-\mathrm{Art}(\mathcal{X}/S) = n(\mathcal{X}_s) - 1 + f$*; the two sign conventions were not reconciled in Appendix HHHH.*

Proof. These are verbatim statements of the source and of Saito's 1988 Duke paper, which proves conductor equals discriminant for relative curves over an arbitrary discrete valuation ring with perfect residue field. The genus-two results of Liu are stated in the source to work also in residue characteristic two; Liu and Lorenzini's qualitative description is explicitly not generalised. Hence the claim of Appendix HHHH that no new tool is required at $p = 2$ is refined: the tool exists, but its two hypotheses are unverified for $C_{7,2}$, and the conjecture is unproved for $g = 4$ at $p = 2$. The sign remark affects only the reading of n_2, not the range of a. □

Theorem 12.124 (12.124: the conductor bound and its normalization band)**.** *With* $48c = 44.1090495938$ *and* $9\log 3 + 9\log 7 = 27.4007019395$*, the only unconditional consequence of* $n_2 \ge 0$ *is*

$$\Delta\gamma_A = \frac{48c}{\log \mathrm{cond}} \le \frac{44.1090495938}{27.4007019395} = 1.609778089.$$

The value 0.761825 *quoted in Appendix HHHH corresponds to the point value* $n_2 = 44$ *and cannot be derived from an upper bound on* n_2*: since* $\log \mathrm{cond} \ge L + n_2 \log 2$ *is a lower bound, inserting an upper bound for* n_2 *does not produce a valid upper bound for* $\Delta\gamma_A$*.*

Proof. The three bad places give $\log \mathrm{cond} \ge L + n_2 \log 2$ with $L = 9\log 3 + 9\log 7$. With $n_2 \ge 0$ the denominator is at least L and the bound follows. Evaluating at $n_2 = 44, 74, 90$ gives 0.761825145, 0.560516400 and 0.491279906 respectively; these are point values under the assumption that n_2 equals that integer, not inequalities. Rational and floating-point checks reproduce all four numbers. The normalization band $\Delta\gamma_A = 48c/\log \mathrm{cond}$ is unchanged and does not collapse. □

Theorem 12.125 (12.125: the exact failure margin and the single-point reduction)**.** *The adelic energy has the form*

$$\mathcal{E}_{\mathrm{ad}} = \mathcal{E}_{\infty}^{\mathrm{arch}} + k_{\mathrm{fin}} - B,$$

where $k_{\mathrm{fin}} \ge 0$ *is supplied by Lemma L6,* $B = (2g-2)\log|D_K|$ *is the structural term of Appendix LLL, and* $\mathcal{E}_{\infty}^{\mathrm{arch}}$ *has no upper bound by Lemma L3. The failure margin of the saturated model is exactly*

$$\Delta\kappa = \tfrac{1}{12} \iff \Delta R = 1 \iff \Delta\gamma = -1.$$

Since $\gamma \ge 0$ *holds unconditionally on all already proved classes, a counterexample to (II) can only be an absolutely simple factor with* $\gamma_A < 0$*, and the whole problem reduces to the single inequality* $h_F \le \tfrac{1}{3}\log \mathrm{cond}$ *for a generic non-CM point of the genus-four family* $C_{7,t}$*.*

Proof. The three equalities are verified by rational arithmetic on the eight anchors: $\gamma(1/3) = 0$ and $\gamma(5/12) = -1$ differ by $1 = 12 \cdot \frac{1}{12}$. The structural term B is the $(2g-2)\log|D_K|$ summand of Parshin's inequality (9), identified in Appendix LLL as $B = 2/c$ on the cyclotomic towers, where it is a structural term rather than a defect. Appendices VVV and QQQ reduce (II) to a single absolutely simple factor with $\mathrm{End}^0 = \mathbb{Q}$, and Appendix WWW shows this factor is the generic Mumford-type fourfold of genus four; Appendices XX and YY show the CM specialisations lie in $\kappa \le 1/8$, already excluded. Hence the only remaining inequality is $h_F \le \tfrac{1}{3}\log \mathrm{cond}$. The three bad places are $n_7 = 9$ (Appendix FFFF), $n_3 \in \{37 - 2i, 39 - 2i\}$ (Appendix GGGG) and $n_2 \le \Delta - 9a$ by Theorem 12.122, conditionally. □

Remark 12.126 (Retractions, scope and honesty boundary). Six statements are corrected. (i) The energy decomposition proposed in Appendix TTT has an extra geometric discrete term $k(X)$, so the arithmetic side needs three terms. (ii) The definition of $\Delta_{\min}$ in Appendix HHHH is the source definition multiplied by $g = 4$; the consequences $a \ge 7$ and $n_2 \le 44$ are withdrawn. (iii) The

claim that no new tool is needed at $p = 2$ is refined by the two hypotheses of Theorem 1.2. (iv) The Artin conductor sign conventions of Muselli and Saito–Liu are distinct. (v) The numerical improvement 0.761825 of Appendix HHHH is withdrawn. (vi) The collapse of the two long-term directions is refined to three independent propositions $(\alpha), (\beta), (\gamma)$ plus one already proved finite discrete energy. Four external sources were checked verbatim this round and not reproved: Hirzebruch's 1985 survey, Kobayashi's 1984 paper, Saito's 1988 Duke paper, and the hyperelliptic models paper with its zbMATH review. The two normalisations of Δ, the nonnegativity of a, the explicit matrix A_v at $p = 2$, the genus-four case of Conjecture 1.4, the absolute Faltings height, the true conductor exponents, the value of a and the sign of γ_A remain open. No new No-Go statement, no new mathematical hypothesis and no new external conclusion are introduced; the fourteen No-Go statements, $(N/M_5)^{24+\varepsilon}$, $E(c) = 12/(5-18c)$, $47/3$, $C \geq 36/7 > 3$ and the status of *abc* as O are unchanged. This appendix is not an independent verification of Theorem A and

12.19 : the sign of γ_A is a three-value judgement

We keep the notation of Appendices BB, DD, LL, UUU, WWW, FFFF, GGGG, IIII, EEEE and HHHH. Put $L = \log \operatorname{cond}(A)$ and $\gamma = 3 - R = 4 - 12\kappa$ with $\kappa = h_F/L$. Three Faltings height normalisations differ by $c\,g$ with $c = \frac{1}{2}\log(2\pi)$ and $g = 4$, so a change of normalisation shifts γ by $\delta = 12cg/L = 48c/L$. Section 12.18 recorded that the sign of γ_A was not computed. In this appendix the sign is computed as a judgement on the triple of normalisations, using only exact rational arithmetic, symbolic differentiation and fifty-digit evaluation. No absolute Faltings height, no computer algebra system and no period matrix are used.

Theorem 12.127 (12.127: symbolic monotonicity)**.** *Write the reference height as H, the normalisation parameter as $c > 0$ and $L > 0$. Then*

$$\gamma(H, L, c) = 4 - \frac{12(H + 4c)}{L}, \qquad \frac{\partial\gamma}{\partial H} = -\frac{12}{L} < 0, \qquad \frac{\partial\gamma}{\partial L} = \frac{12(H + 4c)}{L^2}, \qquad \frac{\partial\gamma}{\partial c} = -\frac{48}{L} < 0.$$

In particular γ is strictly decreasing in the height and in the normalisation constant, and its monotonicity in L has the sign of $h_F = H + 4c$. For the seven anchor classes of Appendix UUU one has $\gamma = 3 - R = 4 - 12\kappa = 4 - \frac{3}{2}\mu_{\mathrm{ratio}}$ and $R = 12\kappa - 1$ identically in rational arithmetic.

Proof. The four displayed identities are the derivative and the algebraic consequence of the dictionary $\kappa = (H + 4c)/L$ and $\gamma = 4 - 12\kappa$; the derivatives are computed by formal differentiation. The anchor identities are verified by substituting the exact rational triples $(\kappa, R, \mu_{\mathrm{ratio}})$ for the saturated model, the modular tower, $X(N)$, $X(p^n)$, the Mason class, the CM–Parseval class and the Heisenberg class; all seven substitutions give zero residual. No external theorem is used. □

Theorem 12.128 (12.128: the three-normalisation sign criterion)**.** *Let the three normalisations be equally spaced and centred on a reference value γ_0, so the three values are $(\gamma_0 - \delta, \gamma_0, \gamma_0 + \delta)$ with $\delta = 48c/L$. Then the three values share a sign if and only if $|\gamma_0| > \delta$; equality $|\gamma_0| = \delta$ produces a zero; and $|\gamma_0| < \delta$ produces two opposite signs. Equivalently the sign is independent of the normalisation precisely when*

$$L > L^*(\gamma_0) = \frac{48c}{|\gamma_0|} \qquad (\gamma_0 \neq 0),$$

while $\gamma_0 = 0$ gives the triple $(-\delta, 0, \delta)$ for every L. If the three normalisations are not equally spaced the same criterion holds with δ replaced by the largest offset.

Proof. The minimum of the triple is $\gamma_0 - |\delta|$ in absolute value terms: all three values are positive exactly when $\gamma_0 - \delta > 0$, and all three are negative exactly when $\gamma_0 + \delta < 0$. Both are equivalent to $|\gamma_0| > \delta$, with the sign equal to the sign of γ_0. Substituting $\delta = 48c/L$ and solving for L gives the stated threshold. For $\gamma_0 = 0$ the triple is $(-\delta, 0, \delta)$; it always contains a zero and its nonzero members have opposite signs. Replacing δ by an upper bound for all offsets can only enlarge the stable region, so the criterion retains its form in the non-equally-spaced case. □

Theorem 12.129 (12.129: an unconditional band for δ on $C_{7,2}$)**.** *For $t = 2$ the bad set is $\{2, 3, 7\}$ and*

$$L = n_2 \log 2 + n_3 \log 3 + 9 \log 7,$$

$$9 \log 3 + 9 \log 7 \le L \le 90 \log 2 + 36 \log 3 + 10 \log 7 = \log \Delta.$$

Consequently, with $48c = 44.10904959382429160546$,

$$\delta \in [\,0.3633592644934122,\ 1.6097780885759230\,].$$

$$\delta_{\max} < 2, \qquad 2 - \delta_{\max} = 0.3902219114240770.$$

Under the conditional bounds $n_2 \le 90$ and $n_3 \le 35$ one has $L \le 118.34786769527674$, $\delta \ge 0.3727067538504091$, and the exact identity $L_{\max} + \log 3 + \log 7 = \log \Delta$.

Proof. The lower bound uses the tame exponents $n_{\mathrm{tame},3} = 9$ of Appendix GGGG, the exact $n_7 = 9$ of Appendix FFFF and $n_2 \ge 0$. The upper bound is the conductor–discriminant bound of Appendix WWW together with $\Delta = 2^{90}3^{36}7^{10}$. Dividing $48c$ by these two endpoints gives the displayed interval; both endpoints reproduce the previously recorded point values 0.36335926449341224246 and 1.609778089. The conditional upper endpoint substitutes $n_2 \le 90$ and $n_3 \le 35$; adding $\log 3 + \log 7$ to it gives exactly $\log \Delta$ because the conductor exponents fall short of the discriminant exponents by $1, 1, 0$ at $7, 3, 2$ respectively. The unconditional interval does not depend on the conditional bounds. □

Theorem 12.130 (12.130: the 1274-point sign matrix)**.** *Take $n_3 \in \{9, 11, \dots, 35\}$, $n_2 \in \{0, 1, \dots, 90\}$ and $n_7 = 9$, giving 1274 conductor values. Evaluating the triple of Theorem 12.128 at each point gives:*

γ_0	*all positive*	*all negative*	*contains zero*	*mixed*	L^*
-1	0	1166	0	108	$48c$
0	0	0	1274	0	∞
1	1166	0	0	108	$48c$
3/2	1271	0	0	3	29.4060330625
2	1274	0	0	0	22.0545247969
5/2	1274	0	0	0	17.6436198375
3	1274	0	0	0	14.7030165313

In particular, since $\delta_{\max} < 2$, every proven class with $|\gamma_0| \ge 2$, namely $\{2, 5/2, 3\}$ and the countable band $[8/3, 3)$, is positive in all three normalisations at all 1274 points:

$$\gamma_A > 0$$

independently of the normalisation and of the integers $(n_2, n_3, i, f_{\mathrm{res}}, a)$. The three mixed points at $\gamma_0 = 3/2$ are $(n_2, n_3) = (0, 9), (1, 9), (2, 9)$ with $L = 27.4007019395,\ 28.0938491201,\ 28.7869963006$.

Proof. For each of the $14 \times 91 = 1274$ pairs the value L is evaluated in fifty-digit precision, $\delta = 48c/L$ is formed, and the signs of $\gamma_0 - \delta$, γ_0, $\gamma_0 + \delta$ are compared by exact rational comparison of γ_0 against δ. The counts in the table are the resulting tallies. Since the largest δ occurring is 1.6097780886 and every listed class with $|\gamma_0| \ge 2$ satisfies $L^*(\gamma_0) \le 22.0545247969 < L_{\min} = 27.4007019395$, the criterion of Theorem 12.128 is met at every point. The three exceptional points are exactly those with $L < L^*(3/2) = 29.4060330625$. □

Theorem 12.131 (12.131: the exact counterexample threshold)**.** *A counterexample to (II) requires $\gamma_A < 0$ in all three normalisations, which is equivalent to $\gamma_0 < -\delta$ and hence to*

$$\kappa_0 > \frac{4 + \delta}{12} = \frac{1}{3} + \frac{\delta}{12} \in [\,0.3636132720411177,\ 0.4674815073813269\,],$$

$$R_0 > 3 + \delta \ge 3.363359264493412.$$

The saturated model $\kappa = 5/12 = 0.41\overline{6}$ lies inside this band.

Proof. By Theorem 12.128, three negative values require $\gamma_0 + \delta < 0$, i.e. $\gamma_0 < -\delta$. Solving $\gamma_0 = 4 - 12\kappa_0$ for κ_0 and inserting the two endpoints of the band in Theorem 12.129 gives the interval, and $R_0 = 12\kappa_0 - 1$ gives $3 + \delta$. The saturated value 5/12 lies between the endpoints because $0.3636132720 < 0.416\overline{6} < 0.4674815074$. □

Theorem 12.132 (12.132: (II) is normalisation invariant only in the tower limit)**.** *Along any tower satisfying (T3) one has $L_n \to \infty$, hence $\delta_n = 48c/L_n \to 0$; eight successive doublings of $L_{\min}$ give $0.804889 \to 0.006288$. Therefore the tower statement (II) is independent of the choice of Faltings height normalisation. For the single curve $C_{7,2}$ however $\delta \geq 0.3633592645 > 0$, so the sign of γ_A at a finite-conductor test point is not normalisation invariant: the single-point formulation of (II) must be read either in the tower limit or with the physical normalisation fixed by a second anchor.*

Proof. The limit statement is immediate from $\delta_n = 48c/L_n$ and $L_n \to \infty$; the eight numerical values are direct evaluations. The single-point statement follows from Theorem 12.129: the band has a strictly positive lower endpoint, so by Theorem 12.128 there exist reference values for which the triple changes sign. In particular $\gamma_0 = 0$ always produces a zero and two opposite signs, so the equality orbit itself is not a normalisation-invariant notion. □

Remark 12.133. Four ledger corrections follow. (i) The claim that the sign of γ_A is uncomputed is withdrawn: it is a three-value judgement, not a boolean. (ii) The quantity $\delta = 48c/L$ was previously used only as a band for the point value; it is simultaneously the sign-flip threshold, and computing the sign needs only the weaker condition $|\gamma_0| > \delta$. (iii) The equality orbit $\gamma_0 = 0$ is never sign-stable, which refines the identification of the dangerous class in Appendices VVV, QQQ, HHHH and IIII. (iv) The single-point reduction of Section 12.18 requires an explicit normalisation qualifier. Finally, with the user's v2 archimedean receipt, $48c/A_{\text{arch}} = 2.54987221085$ and $48c/B_{\text{arch}} = 3.75631446102$ for $A_{\text{arch}} = 17.2985333955547561$ and $B_{\text{arch}} = 11.7426403064859817$; the normalisation band dominates the archimedean variation by a factor between 2.55 and 3.76, so raising the period precision cannot change the sign judgement. All statements are relative to the internal DIC dictionary and to the reported numerical receipt; no external theorem is reproved, no new No-Go statement, mathematical hypothesis, conditional result or external source is introduced.

12.20 : the normalisation anchor equation

We keep the notation of Appendices BB, DD, LL, SSS, WWW, XXX, YYY, ZZZ, AAAA, BBBB, JJJJ and LLLL. Write $L = \log \operatorname{cond}(A)$, $c = \frac{1}{2}\log(2\pi)$, $g = 4$, $\gamma = 3 - R = 4 - 12\kappa$ and $\delta_\gamma = 48c/L$. Section 12.19 showed that the sign of γ_A on a single finite-conductor curve is normalisation dependent. This appendix records the attempt to fix the physical normalisation by a CM anchor; Section 12.21 shows that the resulting equation is an identity rather than a constraint.

Theorem 12.134 (12.134: the anchor equation)**.** *Let $\Delta\Phi$ denote the increment of the canonical height invariant along the Shimura family from a CM specialisation to the generic point, measured on the γ scale. With $\hat{f}_{\text{CM}}$ the conductor mass of the CM endpoint, the calibration offered in Section 12.20 is*

$$B := 3c\big(4 - \hat{f}_{\text{CM}}/L\big) + \Delta\Phi = \delta_\gamma, \qquad \kappa_0 = \frac{1}{3c}(\delta_\gamma - \Delta\Phi) + \frac{\hat{f}_{\text{CM}}}{L}.$$

The two displayed relations are algebraically equivalent, and $\gamma = 4 - 12\kappa = 3 - R$ holds identically on the seven proven anchor values $1/12, 1/8, 1/6, 5/24, 1/4, 1/3, 5/12$.

Proof. Solving the first relation for $\Delta\Phi$ and substituting into the second returns the first relation unchanged, so the two displayed formulas are mutual inverses and carry no independent

information. The mixed-scale form of the second formula is corrected in Appendix LLLL, where $\kappa_0 = (4-\gamma_0)/12$ is used instead. The dictionary identities are checked in exact rational arithmetic on the seven anchor values, twelve equalities in total, all exact. □

Theorem 12.135 (12.135: locality of the CM endpoint)**.** *The fifty-eight CM specialisation parameters of the X_7 and X_9 families are all non-rational, yet every corresponding field discriminant is a perfect fourth power. Consequently the CM normalisation constant c_{CM} depends only on the CM type and carries no tower parameter.*

Proof. The parameter files give an irreducible minimal polynomial for each of the fifty-eight points, so rationality fails at every point. The LMFDB receipts give discriminants of the form $(\prod p^{a_p})^4$ in fourteen of fourteen full records, and the fifty-eight labels reproduce this fourth-power shape through the exact resultant norms. A normalisation constant extracted from such data is a function of the reflex type alone. □

Theorem 12.136 (12.136: the calibration condition)**.** *$c_{\mathrm{CM}} = c$ holds exactly when $\hat{f}_{\mathrm{CM}} = 4L$, and $4 - 12\kappa_{\mathrm{CM}} = 1$ is precisely the $X(N)$ anchor $R = 2$.*

Proof. The condition $\hat{f}_{\mathrm{CM}}/L = 4$ is the calibration convention adopted in Section 12.20; under the dictionary $\gamma = 4 - 12\kappa$ it forces $\gamma_{\mathrm{CM}} = 1$ and therefore $R = 3 - \gamma = 2$, which is the $X(N)$ anchor. No closed formula for the absolute CM height is claimed here. □

Theorem 12.137 (12.137: three horizontal limits)**.** *On the $C_{7,t}$ family one has $\delta_\gamma \geq 0.3633592645$ at $t = 2$; the Qiu dilution limit gives $\delta_\gamma \to 0$; and the horizontal family $t \to \infty$ makes $L \to \infty$ while the S-unit window is finite. Hence a single-point calibration cannot be continued along the whole horizontal family.*

Proof. The first value is $48c/\log\Delta(2)$ with $\log\Delta(2) = 90\log 2 + 36\log 3 + 10\log 7$. The second follows from $L_n \to \infty$ in the dilution tower. The third is the Evertse finiteness statement for S-unit solutions of $t + (1-t) = 1$ already used in Appendices WWW and XXX. □

Theorem 12.138 (12.138: the two scale readings)**.** *The criteria $\kappa_A^{(0)} < 1/3 \iff \Delta\Phi > 4c/L$ and $\gamma_0 > 0 \iff \Delta\Phi > 48c/L$ are the same criterion written on the κ and γ scales respectively; their thresholds differ by the exact factor twelve.*

Proof. $\gamma = 4 - 12\kappa$ converts an increment $\Delta\kappa$ into $\Delta\gamma = -12\Delta\kappa$, and the normalisation band $cg/L = 4c/L$ on the κ scale becomes $12cg/L = 48c/L$ on the γ scale. □

Theorem 12.139 (12.139: a two-point calibration is necessary)**.** *Fixing the normalisation on a horizontal family requires two CM points, not one, in order to remove the field discriminant contribution $(2g-2)\log|D_K|$.*

Proof. A single point determines one linear combination of the normalisation constant and the discriminant term. Two points with distinct field discriminants give two independent combinations, so the constant can be separated. This is the two-anchor refinement of Section 12.19. □

Remark 12.140. Four corrections are filed in Appendix LLLL. (i) The equation $B = \delta_\gamma$ is an identity once the normalisation part of $\Delta\Phi$ is recognised as δ_γ itself, so it cannot determine c. (ii) The CM anchor forces the test point onto the critical orbit $\gamma_0 = 0$, matching Theorem 12.130 exactly. (iii) The quoted ratio 36.14 between δ_γ and $4c/L$ is withdrawn: on a single L the ratio is twelve, and 36.14 mixes $L \approx 40.306$ with $L_{\max} = 121.392390$. (iv) The name Colmez height gap is withdrawn since the generic point has no CM type. All statements are relative to the internal DIC dictionary; no external theorem is reproved, no new No-Go statement, mathematical hypothesis or external source is introduced, and the referee verdict remains Major Revision.

12.21 : what can be computed about $\Delta\Phi$

We keep the notation of Appendices BB, DD, LL, SSS, ZZ, MMM, QQQ, WWW, JJJJ and KKKK. Let $L = \log \mathrm{cond}(A)$, $c = \frac{1}{2}\log(2\pi)$, $g = 4$, $\gamma = 3 - R = 4 - 12\kappa$. Section 12.20 called $\Delta\Phi$ the Colmez height gap of the Mumford-type fourfold and proposed an anchor equation that fixes the physical Faltings height normalisation. This appendix computes $\Delta\Phi$ as far as the existing inputs allow: its scale, its decomposition, its reachable window and the formal Hurwitz-zeta spectral increment. All arithmetic is exact rational or fifty-digit; no computer algebra system, no period matrix and no external theorem is used.

Theorem 12.141 (12.141: the scale dictionary)**.** *The three Faltings height normalisations differ by cg, so the normalisation band is $\delta_\kappa = cg/L = 4c/L$ on the κ scale and*

$$\delta_\gamma = 12\delta_\kappa = \frac{48c}{L} = \frac{44.10904959382429}{L}$$

on the γ scale. The ratio $\delta_\gamma/\delta_\kappa = 12$ is independent of L. On $C_{7,2}$ one has $L_{\max} = 90\log 2 + 36\log 3 + 10\log 7 = 121.392390\ldots$, $L_{\min} = 9\log 3 + 9\log 7 = 27.400702\ldots$, hence $\delta_\gamma \in [0.3633592644934122,\ 1.6097780885759230]$ and $4c/L \in [0.0302799387078,\ 0.134148174048]$.

Proof. $\gamma = 4 - 12\kappa$ is affine with slope -12, so an additive shift of κ multiplies by twelve on the γ scale. The numerical endpoints are direct fifty-digit evaluations of $48c/L$ and $4c/L$ at the two conductor bounds; the ratio was checked at five distinct values of L with residual below 10^{-40}. □

Theorem 12.142 (12.142: domain of definition)**.** *The generic Mumford-type fourfold has Mumford–Tate group of SL_2^3 type, $\mathrm{End}^0 = \mathbb{Q}$ and no CM type, so a quantity Φ_{nonCM} does not exist. The correct definition is the increment of the canonical height invariant along the Shimura family, on the γ scale:*

$$\Delta\Phi := \gamma_{\mathrm{gen}} - \gamma_{\mathrm{CM}} = 12(\kappa_{\mathrm{CM}} - \kappa_{\mathrm{gen}}).$$

The Colmez average formula supplies only the CM endpoint.

Proof. A CM type is a section of the coset space of a CM field with complex conjugation acting freely; a non-toral Mumford–Tate group admits no such section. The height invariant is defined on the whole family, hence its increment exists; the Colmez closed formula requires a CM type. □

Theorem 12.143 (12.143: the anchor equation is an identity)**.** *Decompose $\Delta\Phi = \delta_\gamma + \Delta\Phi_{\mathrm{arch}} + \Delta\Phi_{\mathrm{base}}$, where the first term is the normalisation part. Then the Section 12.20 relation $B = 3c(4 - \hat{f}_{\mathrm{CM}}/L) + \Delta\Phi = \delta_\gamma$ is equivalent to the zero-sum condition*

$$\Delta\Phi_{\mathrm{arch}} + \Delta\Phi_{\mathrm{base}} = 3c\big(\hat{f}_{\mathrm{CM}}/L - 4\big),$$

and at the calibration point $\hat{f}_{\mathrm{CM}}/L = 4$ the right hand side vanishes exactly. Hence the anchor equation is an identity, not a solvable constraint: it cannot determine the normalisation constant c.

Proof. Symbolic simplification of $B - \delta_\gamma$ gives $\Delta\Phi_{\mathrm{arch}} + \Delta\Phi_{\mathrm{base}} + 12c - 3c\hat{f}_{\mathrm{CM}}/L$; setting this to zero and substituting $\hat{f}_{\mathrm{CM}} = 4L$ leaves $\Delta\Phi_{\mathrm{arch}} + \Delta\Phi_{\mathrm{base}} = 0$. The two unknowns and one identity leave one degree of freedom, so c is unconstrained. □

Theorem 12.144 (12.144: the anchor forces the critical orbit)**.** *With $\Delta\Phi_{\mathrm{eff}} := \Delta\Phi_{\mathrm{arch}} + \Delta\Phi_{\mathrm{base}}$ one has $\gamma_0 > 0$ if and only if $\Delta\Phi_{\mathrm{eff}} > 0$. The CM anchor of Theorem 12.143 forces $\Delta\Phi_{\mathrm{eff}} = 0$, hence $\gamma_0 = 0$: calibrating the normalisation on a CM point places the test point exactly on the orbit whose sign is not normalisation invariant, in agreement with Theorem 12.130.*

Proof. $\gamma_0 = \Delta\Phi - \delta_\gamma = \Delta\Phi_{\mathrm{eff}}$ by the decomposition of Theorem 12.143. Theorem 12.128 states that $\gamma_0 = 0$ produces the triple $(-\delta_\gamma, 0, +\delta_\gamma)$, which always contains a zero and two opposite signs. □

Theorem 12.145 (12.145: the reachable and criterion windows)**.** *From* $\gamma_{\mathrm{gen}} \in [-1, 3]$ *and* $\gamma_{\mathrm{CM}} \in [5/2, 4]$ *one gets* $\Delta\Phi \in [-5, 1/2]$ *in exact rational arithmetic. Intersecting with* $\Delta\Phi > \delta_\gamma$ *gives* $0.3633592644934122 < \kappa_0 \le 1/2$*, an interval of length* 0.1366407355065878 *at* $L_{\max}$*,* $(0, 1/2]$ *of length* $1/2$ *in the tower limit, and the empty set at* $L_{\min}$ *since* $\delta_\gamma = 1.6097780886 > 1/2$*.*

Proof. The endpoints are differences of the interval endpoints; the lower bound $\gamma_{\mathrm{CM}} \ge 5/2$ is the CM–Parseval bound of Appendix ZZ, and $\gamma_{\mathrm{gen}} \ge -1$ is the saturated model of Appendix MMM. The intersections are exact rational subtractions. □

Theorem 12.146 (12.146: formal Hurwitz spectral increment)**.** *Put* $F(x) = \zeta(-1, x) + \zeta'(-1, x)$ *and evaluate the formal increment*

$$\Delta\Phi_{\mathrm{eff}}^{\mathrm{AB}} = \tfrac{1}{2} \sum_{\mathrm{gen}} F(x_i) - \tfrac{1}{2} \sum_{\mathrm{CM}} F(x_i)$$

on sixty three-point branch weight spectra, with $N \in \{7, 9, 11, 13, 15, 21\}$*, five face spectra and two generic spectra. The values lie in* $[-0.0771019408,\ 0.4358946950]$*, with fifty-five positive and five negative; three exceed* $\delta_{\min}$ *and none exceeds* $\delta_{\max}$*. Consequently the formal spectra cannot decide* $\gamma_0 > 0$*.*

Proof. The sixty evaluations are fifty-digit numerical substitutions; the classification of signs and the two threshold comparisons are exact inequalities against the endpoints of Theorem 12.141. The substitution is formal because the extension of the canonical height formula to a non-CM point is not proved. □

Remark 12.147. Three ledger corrections follow. (i) The Section 12.20 ratio $\delta_n = 36.14\,\Delta\Phi^*$ is withdrawn: on a single L the κ-to-γ conversion gives exactly twelve, and 36.14 implicitly uses $L \approx 40.306$ for δ_n while $\Delta\Phi^*$ uses $L_{\max} = 121.392390$; as a check $12L_{\max}/L_{\min} = 53.1631884764$. (ii) The mixed-scale expression for κ_0 in Theorem 12.134 is replaced by $\kappa_0 = (4 - \gamma_0)/12$. (iii) Section 12.20's single-sided divergence is refined to the window table of Theorem 12.145. The remaining input is the sign of $\Delta\Phi_{\mathrm{eff}}$, equivalently the Mumford–Tate corank gap inequality; condition (II) stays open. No new No-Go statement, mathematical hypothesis, conditional result or external source is introduced, the fourteen No-Go statements and all constants are unchanged, *abc* remains O, and the referee verdict remains Major Revision. This appendix is not an independent verification of Theorem A.

12.22 : two-point calibration does not shrink the band

We keep the notation of Appendices BB, ZZ, UUU, JJJJ, KKKK and LLLL. Let $g = 4$, $c = \frac{1}{2}\log(2\pi)$, $\gamma = 3 - R = 4 - 12\kappa$, $L = \log \mathrm{cond}(A)$ and $\delta_\gamma = 48c/L$. showed that the sign of γ_A is a three-value judgement on the three Faltings height normalisations, and reduced item (II) to the sign of $\Delta\Phi_{\mathrm{eff}}$. This appendix tests the last purely data-level proposal, namely to fix the physical normalisation by calibrating on two CM specialisation points. All statements below are exact rational identities or fifty-digit evaluations; no computer algebra system, no period matrix and no external theorem is used.

Theorem 12.148 (12.148: the two-point system has a unique solution)**.** *Parametrise the three normalisations by one global additive shift* $s \in \{-c, 0, +c\}$*, so that the observed height of a point* i *is* $H_i^{(s)} = h_i + sg$*. Two CM points with Colmez values* h_i, h_j *give two observations and one unknown* s*; the normal equation has determinant* $2g^2 = 32$*, hence* s *is uniquely determined by*

$$s^* = \frac{(H_i - h_i) + (H_j - h_j)}{2g}.$$

Five exact rational data sets reproduce the prescribed shift exactly.

Proof. The least-squares functional $\sum_i (H_i - h_i - sg)^2$ has second derivative $2\sum_i g^2 = 2g^2 > 0$ and a unique minimiser s^*. Substituting $H_i = h_i + sg$ gives $s^* = s$ identically. The five evaluations are exact rational substitutions with zero residual. □

Theorem 12.149 (12.149: joint cancellation, per-point retention)**.** *For a two-point joint calibration define the single global bookkeeping band*

$$\bar{\delta} = \frac{48c}{\bar{L}}, \qquad \bar{L} = \frac{L_i + L_j}{2},$$

so that $\delta_i = \delta_j = \bar{\delta}$ and $\delta_i - \delta_j = 0$ exactly. Point i carries the triple $(\gamma_{0,i} - \bar{\delta}, \gamma_{0,i}, \gamma_{0,i} + \bar{\delta})$ with the same $\varepsilon \in \{-1, 0, +1\}$ for both points. Then the joint difference $\gamma_i - \gamma_j = \gamma_{0,i} - \gamma_{0,j}$ cancels the band for all three values, while each individual point retains three distinct values whenever $\bar{\delta} > 0$. Consequently calibration determines a joint quantity but cannot determine any single point value.

Proof. Both points share the same $\varepsilon\bar{\delta}$, so subtraction removes it. Since $\bar{\delta} > 0$, the three values of each point are pairwise distinct: the triple is not a singleton. Five exact rational pairs and all three ε were checked. This is the precise content of the statement "the band cancels in the difference but $\pm\bar{\delta}$ is retained at each point". □

Theorem 12.150 (12.150: domain of the anchor-forces-critical claim)**.** *The Section 12.21 anchor equation is $\Delta\Phi_{\text{eff}} = 3c(\hat{f}_{\text{CM}}/L - 4)$. With one conductor L and the calibration value $\hat{f}_{\text{CM}}/L = 4$ it is a rank-one system forcing $\Delta\Phi_{\text{eff}} = 0$, hence $\gamma_0 = 0$ (12.144). With two distinct conductors $L_i \neq L_j$ the right hand sides are two independent constraints on one global unknown $\Delta\Phi_{\text{eff}}$, so the system has rank two and is generically inconsistent; the anchor no longer forces $\gamma_0 = 0$. 12.144 is therefore valid on the one-parameter family only.*

Proof. Exact rational evaluation gives the single-L right hand side 0 and the two-L right hand side pair $(0, 9/4)$ for $\hat{f}/L \in \{4, 9/2\}$; two distinct values cannot be met by one scalar. The rank count is one equation against two equations. □

Theorem 12.151 (12.151: the CM endpoint dictionary)**.** *For the 8T13 CM specialisations the Colmez average has $\mu(\Phi) = 4A$, $d_\Phi = \sqrt{\text{disc}(K)}$ and $m_\Phi = \frac{1}{2}\log d_\Phi$. The fifty-eight per-point receipts of Appendix BBBB all give $\text{disc}(K) = n^4$, hence $d_\Phi = n^2 \in \mathbb{Z}$ and $m_\Phi = 2\log n$ for 58/58 rows. The first four rows have strictly increasing n, giving the three adjacent conductor pairs used below.*

Proof. Parsing the receipt table yields fifty-eight rows of the shape $\text{disc}(K) = n^4$; the identity $n^4 = (n^2)^2$ is exact, so $d_\Phi = n^2$ is an integer and $m_\Phi = \frac{1}{2}\log n^4 = 2\log n$. The first four values $1673, 1757, 2653, 3437$ are strictly increasing. The earlier data-level qualifier "57/58 exact, 1/58 conditional" refers to the square status of the auxiliary integer q, not to $\text{disc}(K) = n^4$; at the level of d_Φ all fifty-eight rows are exact. □

Theorem 12.152 (12.152: the band does not shrink)**.** *With $L_{\min} = 9\log 3 + 9\log 7 = 27.4007019395\ldots$ and $L_{\max} = 90\log 2 + 36\log 3 + 10\log 7 = 121.3923901330\ldots$ one has*

$$\bar{\delta} \in [0.3633592644934122,\ 1.6097780885759230], \qquad \Delta\gamma_A \geq 0.3633592644934122.$$

The band is independent of the calibration of Theorem 12.148 and of numerical precision. The Section 12.21 windows are reproduced exactly: $(0.3633592644934122, 1/2]$ of length 0.1366407355065878 at $L_{\max}$, $(0, 1/2]$ in the tower limit, and empty at $L_{\min}$.

Proof. $\bar{\delta} = 48c/\bar{L}$ with $48c = 44.1090495938242916\ldots$; the endpoints are fifty-digit evaluations. Any two entries of the triple of Theorem 12.149 differ by at least $\bar{\delta}$, so the normalisation uncertainty of a single point is bounded below by the lower endpoint. The window lengths are exact subtractions from 1/2. □

Theorem 12.153 (12.153: the sign matrix after calibration)**.** *For every proved band* $\gamma_0 \in \{2, 5/2, 8/3, 3\}$ *and every* $\bar{\delta} \le 1.6097780886 < 2$*, all three normalisations give* $\gamma_A > 0$*; joint cancellation makes this verdict independent of the two-point calibration. The critical orbit* $\gamma_0 = 0$ *still contains a zero in every normalisation. A counterexample requires* $\kappa_0 > \frac{1}{3} + \bar{\delta}/12 > 0.3636132720$*. Hence item (II) reduces to the single inequality*

$$\Delta\Phi_{\text{eff}} > 0,$$

the Mumford–Tate corank gap, which no two-point calibration can supply.

Proof. Positivity of $\gamma_0 - \bar{\delta}$ follows from $\bar{\delta} < 2 \le \gamma_0$; the triple of Theorem 12.149 then has three positive entries. For $\gamma_0 = 0$ the triple is $(-\bar{\delta}, 0, +\bar{\delta})$. The threshold is $\kappa_0 = (4 - \gamma_0)/12$ with $\gamma_0 < -\bar{\delta}$. The final reduction is Theorem 12.142 together with $\gamma_0 = \Delta\Phi_{\text{eff}}$. □

Remark 12.154. Three ledger dispositions follow. (i) The proposal that a CM anchor is necessary and that two-point calibration upgrades the sign to a point value is withdrawn: Theorem 12.149 shows the point band is retained. (ii) 12.144 is restricted to the one-parameter family by Theorem 12.150; with two distinct conductors the anchor does not force the critical orbit. (iii) The CM endpoint is exact for 58/58 receipts at the level of d_Φ, and the sign matrix is upgraded to calibration-independent positivity on all proved bands, while the dangerous orbit $\gamma_0 = 0$ remains sign-unstable. Item (II) stays open, with the unique remaining input being the new analytic inequality $\Delta\Phi_{\text{eff}} > 0$. No new No-Go statement, mathematical hypothesis, conditional result or external source is introduced; the fourteen No-Go statements and all constants are unchanged, *abc* remains O, and the referee verdict remains Major Revision. This appendix is not an independent verification of Theorem A.

References

[1] C. Zeng, *The Height-Conductor Ratio on Curve Towers: Closed Forms for the Fermat Tower, the One-Third Barrier, and the Attainable Set*, preprint 2610.00544v2 (2026).

[2] C. Zeng, *Asymptotic closed forms for the Faltings height of Fermat Jacobians and the spectrum of logarithmic derivatives of L-functions*, preprint 2610.00712v2 (2026).

[3] C. Qiu, *A Note on Jacobians with Prescribed Factors*, arXiv:2606.04344v1 (2026).

[4] A. Obus, *The (local) lifting problem for curves*, arXiv:0911.1103v3.

[5] C. Soulé, preprint arXiv:2611.05217v1 (Theorems A/B).

[6] A. Brumer, K. Kramer, *The conductor of an abelian variety*, Compositio Math. **92** (1994), 227–248.

[7] P. Colmez, *Périodes des variétés abéliennes à multiplication complexe*, Ann. of Math. **138** (1993), 625–683. MR1239797.

[8] R. Coleman, W. McCallum, *Stable reduction of Fermat curves and Jacobi sum Hecke characters*, J. Reine Angew. Math. **385** (1988), 41–101.

[9] A. Dolce, G. Mercuri, *Intersection matrices for the minimal regular model of $X_0(N)$ and applications to the Arakelov canonical sheaf*, J. London Math. Soc. (2) **110** (2024), no. 2, e12964.

[10] J.-H. Evertse, H. Schlickewei, W. Schmidt, *Linear equations in variables which lie in a multiplicative group*, Ann. of Math. **155** (2002), 807–836.

[11] G. Faltings, *Endlichkeitssätze für abelsche Varietäten über Zahlkörpern*, Invent. Math. **73** (1983), 349–366.

[12] G. Faltings, semistable degeneration asymptotics for the Faltings height (2021), cited at text level.

[13] F. Fité, V. Florit, X. Guitart, *On the product of Shimura curves over totally real number fields*, in preparation; cited at text level.

[14] J. Jorgenson, J. Kramer, *Bounds on the Faltings delta function through covers*, Ann. of Math. (2) **170** (2009), no. 1, 1–43.

[15] C. Khare, J.-P. Wintenberger, *Serre's modularity conjecture (I)*, Invent. Math. **178** (2009), 485–504.

[16] Q. Liu, *Conducteur et discriminant minimal de courbes de genre* 2, Compositio Math. **94** (1994), 51–79; *Inequality for conductor and differentials of a curve over a local field* (with T. Saito), J. Algebraic Geom. **9** (2000), 409–424.

[17] T. Matsusaka, *On a generating curve of an abelian variety*, Nat. Sci. Rep. Ochanomizu Univ. **3** (1952), 1–4.

[18] H. Mayer, *Self-intersection of the relative dualizing sheaf on modular curves $X_1(N)$*, J. Théor. Nombres Bordeaux **26** (2014), 111–161.

[19] L. Moret-Bailly, *Métriques permises*, Astérisque **183** (1990), 29–87.

[20] A. Obus, P. Srinivasan, *Conductor-discriminant inequality for hyperelliptic curves in characteristic* 2, Compositio Math. **158** (2022), 2435–2476.

[21] A. N. Parshin, *Finiteness theorems and hyperbolic manifolds*, Sémin. Théor. Nombres Paris 1986–87, Progr. Math. **75**, Birkhäuser, Boston, 1988, 299–312.

[22] M. Raynaud, *Spécialisation du foncteur de Picard*, Publ. Math. IHÉS **38** (1970), 27–76 (1990 text-level reference).

[23] K. Ribet, *Abelian varieties over* $\mathbf{Q}$ *and modular forms*, in: Modular curves and abelian varieties, Birkhäuser (2004), 241–261.

[24] D. Rohrlich, *Jacobi sums and explicit reciprocity laws*, Compositio Math. **60** (1986), 97–114. MR867959.

[25] T. Saito, *Conductor, discriminant, and the Noether formula of arithmetic surfaces*, Duke Math. J. **57** (1988), 151–173. MR0952229.

[26] C. Stewart, K. Yu, *On the abc conjecture*, Math. Ann. **291** (1991), 225–230; *On the abc conjecture, II*, Math. Ann. **345** (2009), 731–771.

[27] H. Pasten, *The abc conjecture and the square-freeness of consecutive integers*, cited at text level.

[28] J. Tsimerman, *The existence of an abelian variety over* $\overline{\mathbf{Q}}$ *isogenous to no Jacobian*, Ann. of Math. **176** (2012), 637–650.

[29] J. Wilms, *New explicit formulas for the Faltings delta invariant*, Invent. Math. **209** (2017), 481–539.

[30] X. Yuan, *Arithmetic bigness*, Annals of Math. **203** (2026); *one-sided lower bound for the arithmetic self-intersection*, Theorem 1.4 (2026), cited at text level.

[31] X. Yuan, S.-W. Zhang, *The arithmetic Hodge index theorem for adelic line bundles*, Math. Ann. **367** (2017), 1123–1171.

[32] S. Zhang, *Admissible pairing on a curve*, Invent. Math. **112** (1993), 171–193; *Small points and adelic metrics*, J. Algebraic Geom. **4** (1995), 281–300.

[33] E. Bombieri, J. Vaaler, *On Siegel's lemma*, Invent. Math. **73** (1983), 11–32.

[34] T. Borek, *Successive minima of line bundles on arithmetic varieties*, Compositio Math. **142** (2006), 1123–1150.

[35] T. Fujita, *On the structure of polarized varieties with* Δ*-genera zero*, J. Fac. Sci. Univ. Tokyo **22** (1975), 103–115.

[36] L. Stoppino, *Slope inequalities for fibred surfaces via GIT*, Osaka J. Math. **45** (2008), 1027–1041.

[37] S. Mukai, *Curves and symmetric spaces*, Proc. Japan Acad. **68** (1992), 7–10.

[38] D. Mumford, *Stability of projective varieties*, Enseign. Math. **23** (1977), 39–110.

[39] R. Coleman, F. Oort, *Moduli of curves and Jacobians with complex multiplication*, in preparation; B. Moonen, F. Oort, *The Torelli locus and special subvarieties*, cited at text level.

[40] K. Lu, K. Zuo, *On the slope conjecture of Barja and Stoppino for fibred surfaces*, cited at text level; K. Lu, arXiv:1211.0451.

[41] R. Pries, *Wildly ramified covers with large genus*, Arizona Winter School notes, cited at text level.

[42] A. N. Parshin, *The Bogomolov–Miyaoka–Yau inequality for the arithmetic canonical class*, 1988, cited at text level.

[43] U. Kühn, *On the arithmetic self-intersection number of the dualizing sheaf on arithmetic surfaces*, cited at text level; C. Curilla, U. Kühn, preprint, cited at text level.

[44] R. de Jong, semistable degeneration asymptotics (2019), cited at text level.

[45] A. Moriwaki, *Relative Bogomolov's inequality and the cone of divisors*, cited at text level.

[46] J. Jang, *Examples of varieties with non-effective Chern classes*, characteristic p counterexamples, cited at text level.

[47] I. Biswas, K. Paranjape, *The Hodge conjecture for general Prym varieties*, J. Algebraic Geom. **11** (2002), 33–42.

[48] Y. Miyaoka, *On the Chern numbers of surfaces of general type*, Invent. Math. **42** (1977), 225–237; S.-T. Yau, *Calabi's conjecture and some new results in algebraic geometry*, Proc. Nat. Acad. Sci. **74** (1977), 1798–1799.

[49] P. Colmez, *Périodes des variétés abéliennes à multiplication complexe* (averaged form), cited at text level.

[50] A. Granville, T. Tucker, *It's as easy as abc*, Notices Amer. Math. Soc. **49** (2002), 1224–1231.

[51] Preprint arXiv:2406.19785, alignment source of Section 12, cited at text level.

[52] Preprint arXiv:2510.00093v2, degree-eight CM-type tables (58 fields), cited at text level.