

Automated computation of Feynman integrals with BootLoops

Matthew D. Schwartz*
Department of Physics, Harvard University

September 24, 2026

Abstract

Bootstrap methodology for scattering amplitude calculations has been developing for over two decades: symbols and canonical forms, differential equations, high-precision numerics, and so on. Progress has been steady but relatively fragmented: expertise is so specialized that neighboring groups rarely use each other’s software packages, and most papers are either about a new integral, a new mathematical method, or a new software tool. The current generation of large language models (LLMs) can play all three roles at once, as a strong physicist to set up the problem, a great mathematician to recognize the functions, and an effective programmer to drive the tools. The fragmented accomplishments of the community are naturally pooled into an LLM harness which we call BootLoops. BootLoops ports a wide variety of powerful independent software packages for mathematical physics from their native languages (Mathematica, C++, Julia) into a uniform open toolkit. The toolkit can also grow as LLMs improve the tools, or as additional methods are transcribed or invented. BootLoops’ basic function is to automate the semi-numerical bootstrap: reducing to master integrals, fixing the function space from the geometry of the maximal cut, collapsing the unknowns with analytic constraints, pinning the residual coefficients with integer-relation fits against high-precision numerics, and accepting a result only when it reproduces independent values its derivation never used, to thirty digits or more. A decision layer, adjudicated by the LLM, classifies the integral and picks which of the tools is most valuable. Then the LLM runs the tools, manages the memory and CPU needs, and assembles the final answer. This paper gives an overview of the method. Details of many worked examples are in a separate note, and a dedicated exposition of the elliptic and Calabi–Yau bootstrap is in a third.

Contents

1	Introduction	2
2	Feynman integrals, differential equations and function classes	4
2.1	Integral families and canonical differential equations	5
2.2	Function classes from the maximal cut	6
2.3	Numerical evaluation and integer-relation reconstruction	8
3	The semi-numerical bootstrap	9
3.1	Singular locus and function space	9
3.2	Analytic constraints	11
3.3	Numerical data, coefficient fits and acceptance	12

*Work done with Anthropic, but results and views are not endorsed by Anthropic.

3.4	The value-fittability criterion	13
3.5	Operation by the language model	14
4	Calibration on the outer-mass double box	15
4.1	Target and published solutions	15
4.2	Alphabet recovery, constraints and the fit	16
4.3	Exact integrability and sample selection	18
5	The nonplanar two-loop W-pair family	19
5.1	Family, kinematics and reduction	19
5.2	Alphabet and maximal-cut geometry	21
5.3	Block structure and order-by-order solution	21
5.4	Boundary data	23
5.5	Validation	24
6	Constraints by function class	25
6.1	Polylogarithmic families	26
6.2	Elliptic families and value-fittability	27
6.3	Identification and derived boundary constants	28
6.4	Calabi–Yau periods and summary	31
7	Implementation	33
7.1	Codes and ports	34
7.2	Tools written for this work	34
8	Conclusions and outlook	36
A	Details of the F_C computation	38
A.1	Sector labels and alphabet	38
A.2	Non-logarithmic blocks and boundary nodes	39
B	Numerical details	40
B.1	Canonical basis and validation	40
B.2	Evaluators	41

1 Introduction

Precise predictions for collider experiments require scattering amplitudes at second order in perturbation theory and beyond. At those orders an amplitude is a sum of Feynman integrals with two or more loops. Evaluating these integrals as functions of the momenta and masses is often the step that limits the precision of a prediction [123, 2]. At one loop closed results cover the general N -point function [40]. At two loops the polylogarithmic integrals are now computed systematically from their differential equations [72]. When internal masses and thresholds are present, however, the maximal cut of an integral may define an elliptic curve, a K3 surface or a Calabi–Yau variety. The answer is then built from iterated integrals whose kernels involve the periods of that geometry [16, 6, 18], and each such Feynman integral has, in practice, been a separate research problem.

The methods now standard for these computations were developed over more than two decades. Integration-by-parts (IBP) identities, organized by Laporta’s algorithm and solved over finite fields,

reduce any family of integrals to a finite set of master integrals [36, 83, 77, 76, 112, 85, 66]. In a canonical basis the differential equations for the master integrals take an ε -factorized form whose connection is a sum of logarithmic one-forms [71, 86, 9, 97, 87, 62, 104]. The arguments of those logarithms (the letters) determine the symbol of the answer. The symbol bootstrap then fixes polylogarithmic amplitudes from an ansatz in the letters, constrained by integrability together with branch-cut and Steinmann conditions [64, 46, 43, 31, 44]. The candidate letters are now computable directly from the graph, through principal Landau determinants and related methods of computational algebraic geometry [98, 56, 57, 38, 35]. Beyond polylogarithms, the function theory of elliptic and Calabi–Yau integrals has matured (elliptic polylogarithms and iterated integrals of modular forms at genus one, and canonical bases built from the periods for more general geometries), and ε -factorized differential equations exist for many such families [24, 25, 7, 103, 50, 59, 93]. Numerically, several methods (auxiliary mass flow, series solution of the differential equations, sector decomposition) give values of master integrals to essentially arbitrary precision [91, 73, 10, 20, 113]. Finally, analytic expressions are routinely reconstructed from numerical samples on a known function basis, with integer-relation algorithms for the transcendental constants [55, 84]. Such reconstructions are then tested at kinematic points not used to obtain them, as in the two-loop five-point program [1, 3, 4].

The reduction, differential-equation, function-theory and numerical methods above were developed by different communities, and the software reflects that history. The packages are written in Mathematica, C++, Julia and Python, among others, each with its own conventions for integral families and bases. The expertise behind each of these methods is specialized enough that neighboring groups rarely run each other’s packages (few of the reduction programs, for example, share an input format). Most papers in the area accordingly report a new integral, a new mathematical method, or a new software tool. Progress on the integrals has been steady but piecemeal, and each new elliptic or Calabi–Yau computation has been put together by hand from the subset of these tools that its authors use. The automation available in the polylogarithmic case (canonical forms and symbol calculus) does not carry over to such computations unchanged.

Two lines of work have begun to join the analytic and numerical methods into a bootstrap for individual Feynman integrals. The Landau bootstrap [68, 67] constrains an integral from its singular locus, using the letters and the local behavior near each singularity to cut down an ansatz. Analytic regression [12] fixes the rational coefficients of an integral on a known basis by lattice reduction of high-precision samples. Its examples are polylogarithmic, and its authors leave the elliptic case for future work. Separately, a bootstrap for elliptic Feynman integrals exists at the level of the symbol [99, 124]. The established route to an explicit form for a given integral is to solve its canonical differential equation order by order in ε , fixing the boundary constants from regularity conditions and numerical values [72]. To our knowledge, the Landau bootstrap and analytic regression have not yet been combined with that route into one procedure that leads from the graph to an explicit form across function classes, boundary constants included. In this paper we describe such a procedure and the system that runs it, which we call BootLoops. The procedure combines the Landau bootstrap with analytic regression and applies to polylogarithmic and elliptic function spaces. To these we add a criterion, which we call value-fittability, evaluated on the leading-singularity prefactor of each master integral and on the modular properties of the inhomogeneous term in its differential equation. It predicts before any sampling whether a constant-coefficient fit can succeed or the integral must instead be obtained by transporting its differential equation. We also present a new integral family computed in this way and a pooled set of the community’s codes, operated throughout by a large language model (LLM) and assembled so that no result depends on a proprietary computer-algebra kernel.

BootLoops automates the semi-numerical bootstrap as follows. We first reduce the family to

master integrals with IBP identities. We then compute the Landau singular locus of the graph from the principal A -determinant of its Lee–Pomeransky polynomial and from a loop-by-loop Baikov analysis [79, 57, 38]. From the geometry of the maximal cut we read off the function class, which together with the locus fixes a finite graded ansatz for every master integral. Next we impose the analytic constraints (integrability of the connection, the first-entry condition, graph symmetries, the Galois structure of algebraic letters, regularity away from the singular locus), which typically leave a few hundred rational unknowns. We fix those unknowns by integer-relation fits [55, 89] to values of the canonical basis computed by auxiliary mass flow [91]. For the master integrals that fail the fittability criterion we derive or compute boundary data instead and transport the differential equation. We accept a result only if it reproduces, to at least thirty digits, values at kinematic points not used in its derivation.

An LLM carries out each of these steps. It does a physicist’s work in setting up the problem (which includes confirming that propagators taken from the literature define the intended family). It identifies, as a mathematician would, the function class of each master integral and the ring of constants in which a boundary value can lie. As a programmer, it maintains and drives the codes. Concretely, the model runs the singular-locus and cut-geometry analysis and evaluates the fittability criterion on its output. It selects the programs for each master integral, runs them within the memory and processor budget of the machine, and assembles the closed form and its numerical evaluator. When a step lacks a suitable tool the model writes one or ports a published package from its proprietary computer-algebra system to open code (Julia, Python), and adds it to the toolkit.

We first demonstrate the method on the outer-mass double box of Caron-Huot and Henn, the main example of both the Landau bootstrap and analytic regression [32, 68, 12]. From the graph we recover its alphabet and weight-four function space, and from the fit its known symbol as a single basis vector. We depart from those papers at two steps, by imposing integrability as an exact linear system over the rationals and by setting the number of sample points from the conditioning of the lattice reduction. Our second example, which is new, is the nonplanar two-loop family F_C of W -pair production with exact top-quark-mass dependence (the planar families are known [69]). From the geometry of its maximal cuts we find that exactly one sector is elliptic. We obtain the 76 master integrals of the family in explicit form on a one-dimensional kinematic line, largely by transport of the differential equation from a single boundary point, and validate each integral at two independent points on the line. For each function class (polylogarithmic, elliptic or Calabi–Yau) we then determine which analytic constraints apply and which master integrals a coefficient fit can fix. We collect many further worked examples in a companion paper [108], and treat the elliptic and Calabi–Yau bootstrap in detail in another [107].

The paper is organized as follows. In section 2 we review canonical differential equations, the function classes fixed by the maximal cut, and integer-relation reconstruction. In section 3 we state the bootstrap as an algorithm, including the acceptance criterion, the value-fittability criterion and the part taken by the LLM. In section 4 we calibrate the method on the outer-mass double box and in section 5 we compute the F_C family. In section 6 we examine the constraints by function class, in section 7 we describe the codes, and in section 8 we conclude with an outlook.

2 Feynman integrals, differential equations and function classes

In this section we review established material and fix notation. We recall integral families and their canonical differential equations, then the singular locus and the function classes determined by the maximal cut, and finally the numerical and integer-relation methods we rely on.

2.1 Integral families and canonical differential equations

We consider a family of dimensionally regularized L -loop Feynman integrals [72, 123],

$$I_{\nu_1 \dots \nu_N}(\vec{x}; \varepsilon) = \int \prod_{j=1}^L \frac{d^D \ell_j}{i\pi^{D/2}} \frac{1}{D_1^{\nu_1} \dots D_N^{\nu_N}}, \quad D = 4 - 2\varepsilon. \quad (2.1)$$

Here the ℓ_j are loop momenta and $\vec{x}$ collects dimensionless ratios of the Mandelstam invariants and internal masses. The D_k with $k \leq K$ are inverse propagators with the $+i0$ prescription, and the remaining D_k are irreducible scalar products (which enter only with $\nu_k \leq 0$). IBP identities [36, 120], ordered by Laporta's algorithm [83] and solved over finite fields or symbolically [77, 76, 75, 112, 85], express every member of the family through n master integrals $\vec{I} = (I_1, \dots, I_n)$. Since derivatives with respect to x_i stay in the family, the reduction also yields a closed system of differential equations (DEs) [72, 2],

$$\partial_{x_i} \vec{I}(\varepsilon, \vec{x}) = A_i(\varepsilon, \vec{x}) \vec{I}(\varepsilon, \vec{x}), \quad (2.2)$$

where each A_i is an $n \times n$ matrix, rational in ε and $\vec{x}$.

A change of basis $\vec{I} = T(\varepsilon, \vec{x}) \vec{J}$ with T invertible acts on the connection matrices as a gauge transformation,

$$A_i \longrightarrow T^{-1} A_i T - T^{-1} \partial_{x_i} T. \quad (2.3)$$

Henn observed [71] that for large classes of families we can choose T such that the ε dependence factorizes and the connection becomes logarithmic,

$$d\vec{J} = \varepsilon \left(\sum_a \tilde{A}_a d \log W_a(\vec{x}) \right) \vec{J}, \quad (2.4)$$

where the $\tilde{A}_a$ are constant rational matrices and the letters W_a are algebraic functions of $\vec{x}$. We call $\{W_a\}$ the alphabet and eq. (2.4) the canonical or ε -form of the DE. A rational or algebraic T can be constructed by Lee's balance transformations [86], by the Magnus expansion [9] or by Meyer's algorithm. Public implementations are *Libra*, *Fuchsia* and *epsilon* for Lee's algorithm and *CANONICA* for Meyer's [97, 87, 62, 104]. A basis that satisfies eq. (2.4) has uniform transcendentality (UT) weight, in the sense that the ε^k coefficient of each J_m is a pure function of weight k whose differential contains only weight- $(k-1)$ functions times $d \log$ forms [43, 72]. Equivalently, $\vec{J}$ has uniform weight zero if we assign weight -1 to ε .

The matrices A_i for different x_i are not independent, because the mixed second derivatives of $\vec{I}$ commute. With $A = \sum_i A_i dx_i$, applying d to $d\vec{I} = A\vec{I}$ gives the flatness condition

$$dA = A \wedge A, \quad (2.5)$$

valid in any basis and for any class of functions, where the wedge product includes matrix multiplication. In particular, in the ε -form of eq. (2.4) the left side of eq. (2.5) vanishes identically and the condition reduces to $\sum_{a < b} [\tilde{A}_a, \tilde{A}_b] d \log W_a \wedge d \log W_b = 0$, which is empty in one variable. In several variables it relates the matrices $\tilde{A}_a$ whose letters depend on different x_i , and in our experience it is the strongest constraint on a multivariable ansatz [2, 43].

We write the solution of eq. (2.4) along a path γ from a boundary point $\vec{x}_0$ as a path-ordered exponential. The coefficients of its expansion $\vec{J} = \sum_k \varepsilon^k \vec{J}^{(k)}$ are given by Chen iterated integrals [34],

$$\vec{J}(\vec{x}) = \mathbb{P} \exp \left[\varepsilon \int_{\gamma} \sum_a \tilde{A}_a d \log W_a \right] \vec{J}(\vec{x}_0), \quad \vec{J}^{(k)}(\vec{x}) = \sum_a \tilde{A}_a \int_{\gamma} d \log W_a \vec{J}^{(k-1)} + \vec{J}^{(k)}(\vec{x}_0), \quad (2.6)$$

where the constants $\vec{J}^{(k)}(\vec{x}_0)$ are the boundary values. When the letters are rational functions of $\vec{x}$, or algebraic functions whose square roots can be rationalized by a change of variables, these iterated integrals are multiple polylogarithms (MPLs) [63, 46, 122].

The symbol of a weight- k pure function $F^{(k)}$ is a sum of tensor products of k letters from its iterated-integral representation [64, 46], taken in the same order as the integrations of eq. (2.6),

$$\mathcal{S}(F^{(k)}) = \sum_{a_1, \dots, a_k} c_{a_1 \dots a_k} W_{a_1} \otimes \dots \otimes W_{a_k}, \quad \sum_{a, b} c_{\dots a b \dots} d \log W_a \wedge d \log W_b = 0, \quad (2.7)$$

where the $c_{a_1 \dots a_k}$ are rational numbers. We impose the second relation on each adjacent pair of slots with the others held fixed. It is the integrability condition for a tensor to be the symbol of a function and coincides with eq. (2.5). The first entry W_{a_1} is further restricted to letters that vanish at physical thresholds, because a Feynman integral has discontinuities only in physical channels [43]. For a finite alphabet the admissible tensors at each weight span a finite-dimensional $\mathbb{Q}$ -vector space, whose enumeration is the starting point of the symbol bootstrap [43, 31, 44]. The symbol does not detect terms proportional to transcendental constants, which we must fix from boundary data.

2.2 Function classes from the maximal cut

The graph polynomials alone determine where an integral can be singular. In the Lee–Pomeransky representation the integrand is a power of the polynomial $\mathcal{G} = \mathcal{U} + \mathcal{F}$ built from the Symanzik polynomials [88]. If we promote the coefficients of its monomials to independent variables, the integral becomes an A -hypergeometric function in the sense of Gel’fand, Kapranov and Zelevinsky (GKZ) [61, 41, 78]. Its singularities lie on the zero set of the principal A -determinant E_A . This determinant factorizes over the faces τ of the Newton polytope of $\mathcal{G}$,

$$E_A = \pm \prod_{\tau} \Delta_{\tau}^{m_{\tau}}, \quad (2.8)$$

where Δ_{τ} is the discriminant of $\mathcal{G}$ restricted to the face τ and m_{τ} is a positive integer [79]. Applied to the second Symanzik polynomial, the principal A -determinant gives the codimension-one Landau variety of the graph Γ [79],

$$L_1(I_{\Gamma}) = V(E_{A_{\mathcal{F}}}(\mathcal{F})), \quad (2.9)$$

where V denotes the zero set. The faces correspond to the leading Landau singularity, to the normal, anomalous and pseudo-thresholds, and to the second-type singularities [82, 53, 79].

At physical kinematics the coefficients of $\mathcal{G}$ are not generic. The analog of E_A for such constrained coefficients is the principal Landau determinant [57, 56], which we compute with `PLD.jl` and `SOFIA.jl` [38, 98]. Equivalently, the singular locus is the set of kinematic points at which the Euler characteristic of the Lee–Pomeransky hypersurface complement drops [35, 70]. One can also construct it recursively in the loop order [30]. The irreducible components of this locus are the candidate letters W_a of eq. (2.4).

The Landau variety fixes the possible branch points of the integral, and we determine the class of functions with those branch points from the maximal cut. Cutting every propagator of a sector in a loop-by-loop Baikov parametrization [123, 60] leaves an integral of an algebraic function over the remaining variables. The leading singularity is this integral taken over a closed cycle. When the cut integrand has only poles, the leading singularities are algebraic in $\vec{x}$ and so is the transformation T of eq. (2.3). Each order in ε of the integral is then expected to be a combination of MPLs over the alphabet.

At fixed weight the MPLs over a given alphabet span a finite-dimensional space, which we constrain further by the following conditions. Products of MPLs obey the shuffle algebra, so the independent functions at each weight are counted by Lyndon words in the letters [46]. The Steinmann relations forbid consecutive discontinuities in partially overlapping channels [114] and in their extended form constrain every pair of adjacent symbol entries [31]. Cluster adjacency restricts which letters may be neighbors [44]. Finally, the last entry must be one of the letters that appear in the total differential [43].

Suppose instead that the maximal cut of a sector leaves the square root of a quartic (or cubic) polynomial in the last integration variable z ,

$$y^2 = P_4(z; \vec{x}), \quad \psi_i(\vec{x}) = \oint_{\gamma_i} \frac{dz}{y} \quad (i = 1, 2), \quad \tau = \frac{\psi_2}{\psi_1}, \quad q = e^{2\pi i \tau}. \quad (2.10)$$

The cut is then a family of elliptic curves over the kinematic space, with periods ψ_1, ψ_2 over two independent cycles $\gamma_{1,2}$ and modulus τ [16, 6, 5]. The leading singularity is now the period ψ_1 , a transcendental function of $\vec{x}$ annihilated by a second-order Picard–Fuchs operator. The transformation T to an ε -factorized basis accordingly involves ψ_1 and its derivative [7, 60, 126]. Each master integral carries its own period prefactor [50, 59, 33]. It is of the first kind (ψ_1), of the second kind (i.e. the quasi-period, which we denote η_1) or of the third kind, the last attached to a marked point of the curve.

In the variable τ we replace the logarithmic kernels of eq. (2.4) by modular ones,

$$d \log W_a \longrightarrow \omega_n = 2\pi i f_n(\tau) d\tau, \quad f_n \in \mathcal{M}_n(\Gamma), \quad (2.11)$$

where $\mathcal{M}_n(\Gamma)$ is the space of modular forms of weight n for the monodromy group $\Gamma \subset \text{SL}(2, \mathbb{Z})$ of the family [24, 25, 128]. The algebraic $d \log$ letters attached to the rational part of the singular locus remain in the connection alongside the ω_n , which in the language of the curve are Kronecker–Eisenstein forms [26, 28]. The iterated integrals of eq. (2.6) become iterated integrals of modular forms, or elliptic multiple polylogarithms (eMPLs), graded by modular weight in addition to length [24, 6, 123]. These functions are not labeled by words in a finite alphabet, and neither Lyndon-word counting nor the adjacency conditions carry over. However, eq. (2.5) and the first-entry condition still apply, with the cusps of Γ in the role of thresholds.

Integrating the differential of the second kind over the two cycles $\gamma_{1,2}$ of eq. (2.10) gives the periods of the second kind, η_1 and η_2 . The four periods satisfy the Legendre relation

$$\psi_1 \eta_2 - \psi_2 \eta_1 = 2\pi i \quad (2.12)$$

in the normalization of refs. [6, 25]. Thus the period matrix has constant nonzero determinant, and the quasi-period η_1 is linearly independent of its partner ψ_1 . The ratio η_1/ψ_1 is expressed through the Eisenstein series of weight two,

$$\frac{\eta_1}{\psi_1} \sim E_2(\tau) + (\text{Eichler integral}), \quad E_2\left(-\frac{1}{\tau}\right) = \tau^2 E_2(\tau) + \frac{12\tau}{2\pi i}, \quad (2.13)$$

where the second relation shows that E_2 is quasimodular. That is, E_2 satisfies the weight-two transformation law only up to the inhomogeneous term, which no combination of modular forms for Γ can cancel [128, 25]. The value-fittability criterion of section 3.4, by which we decide whether a coefficient fit can succeed, is based on this quasimodularity of η_1/ψ_1 .

Beyond genus one the maximal cut is a period of a higher-dimensional variety (a K3 surface for the three-loop banana graph and a Calabi–Yau variety of dimension $L - 1$ at L loops) [15, 103, 102,

18, 80]. The vector of periods $\vec{\omega}$ is annihilated by the Picard–Fuchs operator of the cut,

$$\mathcal{L} \vec{\omega}(x) = \left[\theta^r + \sum_{j=0}^{r-1} p_j(x) \theta^j \right] \vec{\omega}(x) = 0, \quad \theta = x \frac{d}{dx}, \quad (2.14)$$

where the p_j are rational functions and r is the number of independent periods. Near a point of maximal unipotent monodromy (MUM) the Frobenius basis of eq. (2.14) is organized by powers of $\log x$. Griffiths transversality imposes quadratic relations among the periods, packaged in the Yukawa coupling and the Y -invariants [49, 45, 121], and the ε -factorized connection is built from the periods and these invariants [49, 93]. In all these cases the transformation to ε -form thus contains the periods themselves [7, 60]. We use the following correspondence between cut geometry and function class throughout the paper. An algebraic maximal cut gives MPLs, an elliptic curve gives eMPLs and iterated integrals of modular forms, and a Calabi–Yau variety gives iterated integrals of its periods [107, 11, 100].

2.3 Numerical evaluation and integer-relation reconstruction

We obtain high-precision values of the master integrals at chosen points with the auxiliary-mass-flow method [91]. Every propagator of eq. (2.1) is shifted by an auxiliary mass η ,

$$\frac{1}{D_k} \rightarrow \frac{1}{D_k - \eta}, \quad \partial_\eta \vec{I}_{\text{aux}}(\varepsilon, \eta) = A_\eta(\varepsilon, \eta) \vec{I}_{\text{aux}}(\varepsilon, \eta), \quad \vec{I}(\varepsilon, \vec{x}_0) = \lim_{\eta \rightarrow i0^-} \vec{I}_{\text{aux}}(\varepsilon, \eta), \quad (2.15)$$

where $\vec{I}_{\text{aux}}$ are the master integrals of the auxiliary family at the fixed point $\vec{x}_0$ and the DE in η follows from their IBP reduction. As $\eta \rightarrow \infty$ the auxiliary integrals reduce to equal-mass vacuum integrals, which supply the boundary values. We recover the physical integrals by solving the η -equation numerically through its singular points down to $\eta = 0$ [91]. We transport these values to further points with series solutions of the kinematic DE along a path [73, 10]. Sector decomposition [20, 113] gives an independent numerical evaluation by direct integration over Feynman parameters.

Suppose we know a value v and candidate basis numbers $b_1, \dots, b_m$ (zeta values, periods, L -values or basis functions at the same point) to P digits. The PSLQ algorithm [55] and LLL lattice reduction [89] then search for integers $c_0, \dots, c_m$, not all zero, such that

$$c_0 v + c_1 b_1 + \dots + c_m b_m = 0 \quad (2.16)$$

holds to within 10^{-P} [118, 116, 119]. An accidental relation with coefficients bounded by H is expected to fail beyond roughly $m \log_{10} H$ digits. In practice a relation with small coefficients that persists when we raise the precision is strong evidence for an exact identity. For example, Laporta identified the constants of the four-loop electron anomalous magnetic moment in this way [84]. Analytic regression uses the same step to fix the rational coefficients of a polylogarithmic ansatz from values at several points [12].

We also reconstruct the connection matrices themselves from samples. Consider an entry of A_i in eq. (2.2) as a function of one variable y with the others fixed. We recover it from numerical or finite-field samples $A(y_1), \dots, A(y_M)$ by Thiele’s interpolation formula,

$$A(y) = a_0 + \frac{y - y_1}{a_1 + \frac{y - y_2}{a_2 + \frac{y - y_3}{a_3 + \dots}}}, \quad (2.17)$$

where the inverse differences a_j are computed from the samples and the fraction terminates when A is rational. We accept the reconstruction when it reproduces samples that were not used to build it, as in the finite-field interpolation of FireFly and Kira [76, 75, 77]. The same test at points not used in the reconstruction is standard when analytic expressions are assembled from numerical or finite-field evaluations, as in the two-loop five-point computations [1, 3, 4] and in computations with internal masses [69].

3 The semi-numerical bootstrap

The method takes as input an integral family (a graph with its propagators, internal masses and external kinematics), the orders in ε required, and a kinematic region. For each master integral it returns an explicit form over a given alphabet and an arbitrary-precision evaluator of that form, validated at kinematic points not used in the derivation. It has five steps (Figure 1), the fourth with two alternative routes. We first reduce the family to a canonical basis of master integrals [77, 76, 71]. We then compute the singular locus and the geometry of the maximal cut, which fix the function space. On that space we write a graded ansatz and impose the analytic constraints. Next we fix the surviving coefficients by integer-relation fits to high-precision values. When the criterion of section 3.4 shows in advance that no constant-coefficient fit can succeed, we instead derive or compute boundary data and transport the ε -form differential equation. Finally, we accept a result only if it satisfies eq. (3.6).

3.1 Singular locus and function space

After the reduction to master integrals we determine where in kinematic space the integral can be singular. The singular locus depends only on where the propagators can pinch the integration contour, and is therefore a property of the graph polynomials [88, 79, 98, 30]. We compute the Landau variety of eq. (2.9) with one or more of three codes and, when more than one is run, take the union of their outputs. The codes are `PLD.jl` [57, 56], which computes the principal Landau determinant, `SOFIA.jl` (our port of the SOFIA package [38]), which performs a loop-by-loop Baikov analysis, and an implementation of the graphical Landau analysis of ref. [68]. As a test of completeness we add any component on which the Euler characteristic of the Lee–Pomeransky critical-point variety drops and which these computations missed [35]. The principal-determinant construction relies on genericity conditions of the GKZ system that hold for most families but not all, and at physical kinematics the completeness of the locus is proven only at one loop [57]. Moreover, the locus degenerates at equal-mass and zero-mass points, and we recompute it there [57, 70]. The output of this analysis is the candidate alphabet, including second-type letters and letters that a threshold analysis by hand would be likely to miss. An example is the quadratic $L_{\text{bub}} = x^2 - 10x + 5$ of the two-loop ice-cream cone [107], with x the kinematic variable of that family.

For a polylogarithmic family the alphabet already determines the function space, since the candidate functions are the weight-graded words in the kernels $d \log a$ of its letters [64, 46]. Beyond polylogarithms the variety of eq. (2.10) left by the maximal cut determines the function class (section 2.2) [24, 25, 50]. From an elliptic cut we compute the j -invariant, the modular level, the cusps and Kodaira fibers, and whether the curve has complex multiplication. These data fix the integration kernels and the ring in which boundary constants are sought.

We next normalize each master integral and write the ansatz. The leading singularity $\text{LS}_i(x)$ of

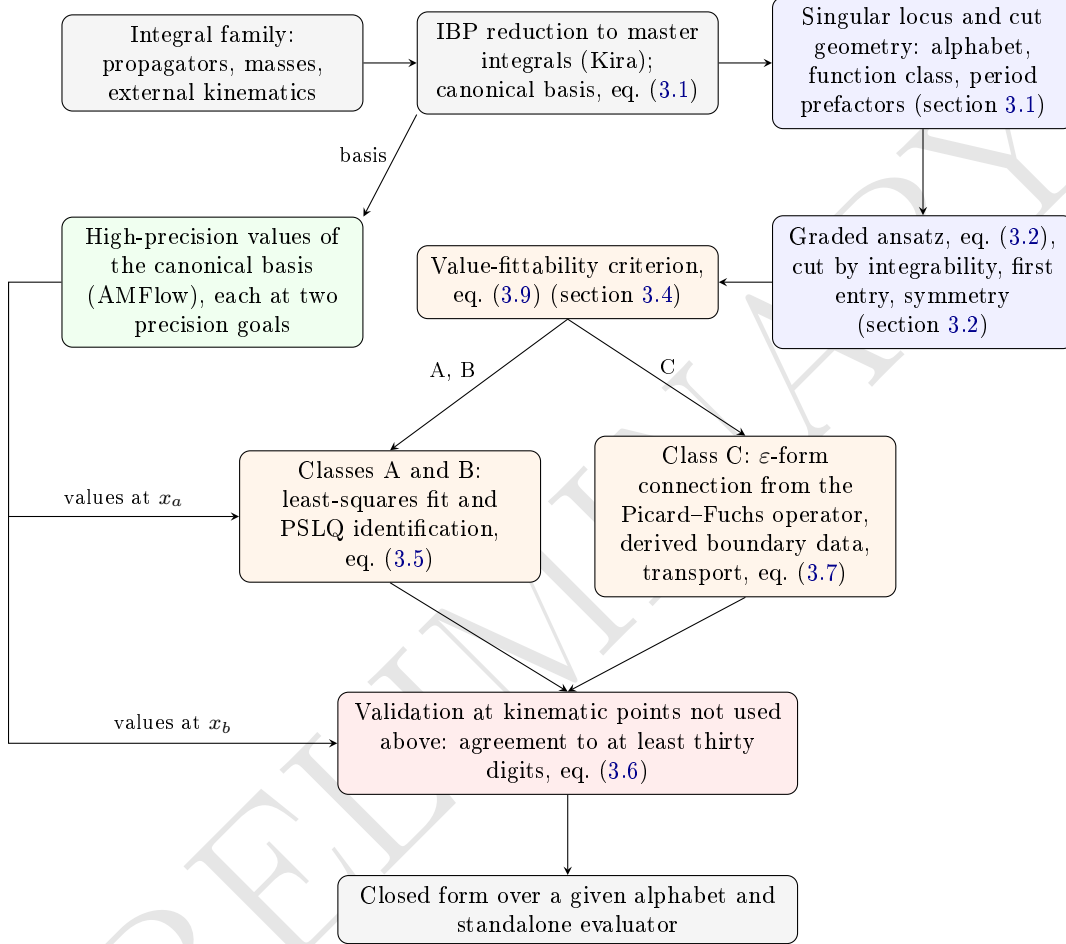


Figure 1: The semi-numerical bootstrap as one procedure with two routes. Reading from the top, reduction and the singular-locus and cut-geometry analysis fix a canonical basis and its function space, and the analytic constraints cut the graded ansatz. Each master integral is then computed, according to the value-fittability criterion of section 3.4, either by a coefficient fit against high-precision values (classes A and B) or by transport of the ε -form differential equation from derived or computed boundary data (class C). Both routes end at the same test, eq. (3.6), at kinematic points x_b disjoint from the fit points x_a .

the maximal cut supplies the prefactor that brings a master to uniform transcendental weight,

$$g_i(\varepsilon, x) = \varepsilon^{p_i} e^{2\varepsilon\gamma_E} \text{LS}_i(x) \sum_j T_{ij}(\varepsilon, x) I_j(x), \quad (3.1)$$

where the I_j are the masters returned by the reduction, T is here a rational IBP rotation, p_i an integer, and g is the canonical basis written $\vec{J}$ in eq. (2.4). The exponential is the usual normalization at two loops ($e^{L\varepsilon\gamma_E}$ at L loops). Beyond polylogarithms LS_i is replaced, master by master, by a period of the cut geometry of the first, second or third kind (ψ_1 , the quasi-period η_1 , or a puncture period) [50, 59]. At order ε^k the ansatz for a normalized master is

$$g_i^{(k)}(x) = \sum_{j=1}^{N_k} c_j \phi_j(x), \quad c_j \in \mathbb{Q} \text{ (or a stated ring of constants)}, \quad (3.2)$$

where the ϕ_j are the weight- k iterated integrals over the alphabet just determined. For an elliptic sector they are iterated integrals over the Kronecker–Eisenstein kernels of the cut curve and the surviving dlog letters, graded by modular weight [24, 25, 99]. The unknowns are the rational numbers c_j and a small number of boundary constants.

3.2 Analytic constraints

We impose on eq. (3.2) all constraints that follow from the differential equation, the analytic structure of the integral and the graph [68, 43, 31]. Written as conditions on the coefficients, with $c_{a_1 \dots a_k}$ the coefficient of the word $a_1 \otimes \dots \otimes a_k$ in the symbol of eq. (2.7), they are

$$\begin{aligned} (\text{integrability}) \quad & \sum_{a,b} c_{w a b w'} d \log a \wedge d \log b = 0 \quad \text{for all words } w, w', \\ (\text{first entry}) \quad & c_{a w} = 0 \quad \text{unless } \{a = 0\} \text{ is a physical threshold,} \\ (\text{regularity}) \quad & g_i^{(k)} \text{ is finite at every point off the Landau variety,} \\ (\text{symmetry}) \quad & g(\sigma \cdot x) = R(\sigma) g(x) \quad \text{for every automorphism } \sigma \text{ of the graph,} \\ (\text{parity}) \quad & c_j = 0 \quad \text{unless } \phi_j \text{ lies in the Galois sector fixed by } \text{LS}_i, \\ (\varepsilon\text{-stratification}) \quad & g_i^{(k)} \text{ has uniform weight } k, \end{aligned} \quad (3.3)$$

where $R(\sigma)$ is the constant matrix by which σ acts on the master integrals. Integrability is the flatness condition of eq. (2.5) read on the symbol [2, 46]. The first-entry condition admits in the first slot only letters that vanish in a physical channel, since the integral is real and analytic at pseudo-thresholds and at infinity. Regularity forbids singularities off the Landau variety of eq. (2.9) [68, 98]. When the alphabet contains square roots, LS_i fixes the parity of each master under the sign changes of the roots, and we remove the words of the other parity [68, 67]. The last condition holds on the canonical basis of eq. (3.1) order by order in ε and removes words of lower weight. In the elliptic case we use the same list with two changes: the grading is by modular weight as well as by length, and integrability becomes a set of relations among Eisenstein series of adjacent modular weight [24, 6, 25].

We write the effect of the constraints as the sequence of dimensions of the space of allowed coefficient vectors,

$$N_0 \longrightarrow N_{\text{int}} \longrightarrow N_{\text{fe}} \longrightarrow N_{\text{sym}}, \quad (3.4)$$

after integrability, first entry (with parity) and symmetry in turn. Other conditions, such as the extended Steinmann or last-entry relations, lengthen the sequence where they apply. For the outer-mass double box at weight four the sequence runs from 20736 to 161 (eq. (4.5)), and for the

maximal-cut ansatz of the top sector of the three-loop crossed light-by-light box from 2401 to 34 (Table 3). In the families treated here the remaining unknowns are of order 10^2 to 10^3 rational coefficients and a handful of boundary constants, few enough for an integer-relation fit. A fit to the unconstrained ansatz, by contrast, would only interpolate the sample [68, 12].

3.3 Numerical data, coefficient fits and acceptance

The numerical input is a set of values of the canonical basis g of eq. (3.1) at kinematic points x_a . We compute them by auxiliary mass flow with **AMFlow** [91], on reductions generated with Kira [77, 75]. We sample the canonical basis itself, because the rational prefactors of the raw reduction basis mix transcendental weights and make the linear system for the coefficients ill-conditioned. Equation (3.1) fixes only the ε^0 coefficient of g_i , and at higher orders we complete the UT basis by subtracting dimension-shifted subsector integrals through ordinary IBP identities [123, 77] (Appendix B). In practice we compute each value twice, at a precision goal G and at a second, higher goal ($G + 12$ in the double-box calibration, $G + 30$ for all other results) on a different ε -grid. We keep only the digits on which the two runs agree.

With K sample points the ansatz of eq. (3.2) becomes the linear system

$$\sum_{j=1}^N \phi_j(x_a) c_j = g^{\text{num}}(x_a), \quad a = 1, \dots, K, \quad (3.5)$$

where N is the dimension left by eq. (3.4) and $g^{\text{num}}(x_a)$ are the **AMFlow** values. When $K \geq N$ we solve eq. (3.5) by least squares and identify each c_j and each boundary constant with PSLQ, eq. (2.16) [55, 89], in the ring of constants fixed by the cut geometry. When the c_j are rationals of small height, lattice reduction applied jointly to the numbers $\phi_j(x_a)$ and $g^{\text{num}}(x_a)$ at all K points recovers them also for $K < N$ [12, 89], as in the calibration of section 4. We keep an identification only if it is stable when the working precision is varied. This is the analytic-regression step of ref. [12], here applied beyond polylogarithms as well.

To test a fit we evaluate it at a second set of kinematic points $\{x_b\}$, disjoint from the fit points $\{x_a\}$. Where possible we choose the x_b in a different kinematic region (fitting at large $|t|$ and predicting at small $|t|$, say). The acceptance criterion is

$$|g^{\text{fit}}(x_b) - g^{\text{num}}(x_b)| < 10^{-30} |g^{\text{num}}(x_b)| \quad \text{for every } x_b, \quad (3.6)$$

order by order in ε , with the fit stable under deletion of any one sample point. Comparison with independent numerical evaluations is standard practice [1, 3, 4, 69, 39], and eq. (3.6) merely makes the threshold definite and requires the points x_b to be disjoint from the fit points. As with the four-loop identifications of ref. [84], a lattice-reduction identification, checked to many digits at points not used in the fit, is a numerically overwhelming but not symbolically proven statement.

The second route to a master integral is transport of the ε -form differential equation. Given the connection of eq. (2.4), written as $\varepsilon d\tilde{A}$ with $\tilde{A} = \sum_a \tilde{A}_a \log W_a$, and the value of g at one point x_0 , the value at x_b is

$$g(x_b) = \mathbb{P} \exp\left(\varepsilon \int_{x_0}^{x_b} d\tilde{A}\right) g(x_0), \quad (3.7)$$

which we evaluate by matched series expansions along a path avoiding the singular locus [72, 73, 10]. We derive the boundary vector $g(x_0)$ where possible, from regularity at a cusp or from a parametric integral representation, and compute it with **AMFlow** otherwise. When no ε -form is available in the kinematic region considered, we apply the same series expansions directly to eq. (2.2), either

order by order in ε or at fixed rational values of ε . Several blocks of the differential equation of the W -pair family require this variant. For the class-C masters of eq. (3.10) below we obtain the result from eq. (3.7) and apply eq. (3.6) to the transported value. For a master obtained by the fit the transported value gives an independent test, since it does not use the `AMFlow` value at x_b .

We regard a master integral as computed, in the sense used throughout this paper and its companions [108, 107] (the `BootLoops` standard), when three requirements are met. The result must be an explicit form over a given alphabet, in which every constant is classical or is defined by a convergent integral or series. It must come with an evaluator that contains no reduction or auxiliary-mass-flow code and runs at arbitrary kinematics and working precision. Finally, the evaluator must satisfy eq. (3.6). These requirements, with the same thirty-digit threshold, apply to each result below whether it is obtained by a fit or by transport.

3.4 The value-fittability criterion

For a polylogarithmic master with a complete alphabet, eq. (3.5) has an exact solution within eq. (3.2). For an elliptic master such a solution may not exist, because the de Rham cohomology $H_{\text{dR}}^1(E)$ of the cut curve is two-dimensional while the maximal cut supplies only the period of the first kind, ψ_1 . The period of the second kind enters through the differential equation, and a constant-coefficient combination of graded functions may be unable to represent it. We therefore decide, before any value is computed, whether a fit can succeed.

Let us write a master integral at fixed order in ε as

$$I_{\text{phys}}(x) = P(x) \hat{I}(x), \quad (3.8)$$

where $\hat{I}$ is a constant-coefficient, weight-graded combination of elliptic polylogarithms or iterated Eisenstein integrals in a frame derived from the Picard–Fuchs operator, and P is the prefactor. By a frame we mean here a choice of basis for the periods together with a choice of modular variable. In the construction of refs. [50, 59] the period matrix factorizes as $W = W_{\text{ss}} W_{\text{u}}$ into a semisimple part of weight zero, which generalizes the algebraic leading singularity, and a unipotent part. A final rotation U_t removes a residual nilpotent piece. We call I_{phys} value-fittable if the fit of eq. (3.5) determines it without integrating its differential equation. This holds if and only if the canonical rotation can be replaced by one with x -independent coefficients, which requires

- (i) W_{ss} (equivalently P) is removable by an x -independent map, and
 - (ii) the inhomogeneous kernel K of the particular solution is a modular q -series on the sampled region, so that U_t is trivial.
- (3.9)

Condition (i) sorts prefactors into three classes,

$$P(x) = \begin{cases} R(x) & \text{(A) algebraic, read from the cut leading singularity and the IBP rotation,} \\ (\psi_1(x)/\pi)^k R(x) & \text{(B) a single global power of the holomorphic period,} \\ R_1(x) \psi_1(x) + R_2(x) \eta_1(x) & \text{(C) period mixing,} \end{cases} \quad (3.10)$$

with $R, R_1, R_2 \in \overline{\mathbb{Q}}(x)$ and with (i) failing only in class C. Condition (ii) is needed because an algebraic P removes only the homogeneous ψ_1 -dressing. The kernel

$$K = \frac{\psi_1^3}{W}, \quad (3.11)$$

with W here the Wronskian of the Picard–Fuchs operator (not the period matrix above), can still bring in the second-kind period, and is globally a modular q -series only when the monodromy group is a congruence subgroup. Thus the classifying datum is the pair (P, K) , as the equal-mass and unequal-mass kites of section 6.2 illustrate. In what follows we also assign to class C a master whose prefactor is of type A or B but whose kernel fails (ii).

The class-C obstruction can be seen in the frame in which ψ_1 is holomorphic. There the ratio of the two periods obeys eq. (2.13), in which E_2 is quasimodular, and the Legendre relation of eq. (2.12) implies that ψ_1 and η_1 are linearly independent [25, 6, 128]. A class-C prefactor therefore multiplies a graded word by a non-graded object with x -dependent coefficients, and no assignment of weights removes it. The criterion of eq. (3.9) rests on this obstruction argument and on the factorization of the period matrix described above. Its forward direction is proven given the canonical-differential-equation construction, its obstruction direction is argued at the level of a linear-independence statement over the ring of cusp constants and is not a transcendence theorem, it has one out-of-sample confirmation, and its sufficiency in the Calabi–Yau case is open. The quasimodular language is ours, and eq. (3.8) is obtained in refs. [50, 59] without it.

We compute a master in class A or B whose kernel satisfies (ii), such as the equal-mass kite top, by the fit of eq. (3.5). A master in class C is computed by transport. For such masters we build the ε -form connection from the Picard–Fuchs operator, which keeps the second-kind period inside the connection and the coefficients algebraic [7, 60, 50]. We then derive or compute the boundary vector $g(x_0)$, transport it with eq. (3.7), and compare the result with `AMFlow` values at points away from x_0 . The ice-cream-cone top integral, the unequal-mass kite top integral and the elliptic $gg \rightarrow Z\gamma$ masters are in class C [107, 108]. Table 4 lists for each of them how far a direct fit falls short of eq. (3.6), next to a control fit that satisfies it in the same framework.

In practice each step can fail in a recognizable way. If the singular locus is incomplete, a singular point that is not in the alphabet appears on the transport path of eq. (3.7). Computing the locus in two ways and adding the Euler-characteristic test of section 3.1 should prevent this failure. If the alphabet or the frame is wrong, or a word is missing from eq. (3.2), the residual of the fit stalls at the same level at the fit points and at the validation points. Neither more sample points nor higher precision lowers it (an example is the equal-mass kite fit of Table 4 in a different frame with one word omitted). In particular, for a class-C master the residual already stalls at a few digits, as eq. (3.9) predicts. Sampling the raw basis instead of eq. (3.1) has the same effect through ill-conditioning, and insufficient reference precision shows up as integer relations that change between the two precision goals. Apart from these failures, the number of master integrals can be a limitation, and for Calabi–Yau threefold families whose full master vector is too large for auxiliary mass flow on one machine we work on the maximal cut [107].

3.5 Operation by the language model

The steps above are all performed with deterministic codes, those of section 7, which a person could equally well run. In this work they are run by a large language model. The model together with those codes is the LLM harness that we call `BootLoops`. The model first classifies the integral by running the singular-locus and cut-geometry analysis of section 3.1. The output of that analysis (genus, j -invariant, modular level, complex multiplication or not) determines which codes are needed. The model then evaluates the criterion of eq. (3.9) from the prefactor class and the modularity of the kernel, and assigns each master to the fit or to transport. This step is the decision layer adjudicated by the LLM, although the criterion itself is definite and gives the same assignment to anyone who applies it.

The model also launches and supervises the reductions, the auxiliary-mass-flow runs and the

fits. It sets a memory cap and a core count for each job at launch and estimates completion times from the job's own log. It thereby keeps the memory and processor load of many concurrent runs within the machine's capacity. When a run fails the model reads the log and distinguishes a job that exceeded its resources from a defect in a code. The first is relaunched with corrected settings; the second is repaired, and the patches to `amflow-cpp` described in section 7 were made in this way. Where a step had no code the model wrote one, for example `numkin` for a reduction that would not finish symbolically. Finally, the model assembles the identified coefficients and constants into the closed form, writes the standalone evaluator, and runs the validation. The model's judgment plays no part in acceptance, which depends only on eq. (3.6). The steps draw in turn on algebraic geometry, computer algebra and perturbative quantum field theory. We expect the procedure to be practical for a single operator mainly because the model can work in all three areas.

4 Calibration on the outer-mass double box

The outer-mass double box of Caron-Huot and Henn [32] is the main worked example of the two methods this paper builds on, the Landau bootstrap and analytic regression. Its symbol was derived from the Landau singularities of the graph in ref. [68], and it is Example 4 of ref. [12], where the same symbol was recovered by a fit to high-precision numerical values. The integral is polylogarithmic of weight four and its alphabet contains three square roots. Once the leading singularity is removed, its symbol is a single element of the basis of weight-four symbols that satisfy the constraints imposed below. We use this integral to calibrate the procedure of section 3, since computing it requires each step of that procedure and its answer is known in closed form. We compare our alphabet and symbol with those of ref. [68] and do not use its maximal cuts as input. The one input we take from that paper is its Regge-limit argument, which excludes six of the odd-letter candidates at the alphabet step.

4.1 Target and published solutions

The integral is the planar two-loop double box with all six perimeter propagators of mass m , a massless central propagator and massless on-shell external legs (Figure 2), and it is finite in four dimensions. With $u = 4m^2/(-s)$ and $v = 4m^2/(-t)$ we normalize by the leading singularity,

$$g_{10} = -\frac{1}{8} s^2 t \beta_u \beta_{uv} I_5, \quad \beta_u = \sqrt{1+u}, \quad \beta_{uv} = \sqrt{1+u+v}, \quad (4.1)$$

where I_5 denotes the scalar integral of Figure 2 with unit propagator powers in four dimensions, and g_{10} is then a pure multiple polylogarithm of weight four in (u, v) [32, 64]. The three square roots β_u , $\beta_v = \sqrt{1+v}$ and β_{uv} rationalize simultaneously in the variables (w, z) used by both earlier computations [68, 12],

$$u = \frac{(1-w^2)(1-z^2)}{(w-z)^2}, \quad v = \frac{4wz}{(w-z)^2}, \quad \beta_u = \frac{1-wz}{w-z}, \quad \beta_v = \frac{w+z}{w-z}, \quad \beta_{uv} = \frac{1+wz}{w-z}, \quad (4.2)$$

where the Euclidean region $u, v > 0$ is $0 < z < w < 1$. The symbol of g_{10} is written in the twelve-letter alphabet of ref. [68, Tab. II],

$$\begin{aligned} L_1 &= u, & L_2 &= v, & L_3 &= 1+u, & L_4 &= 1+v, \\ L_5 &= u+v, & L_6 &= \frac{\beta_u - 1}{\beta_u + 1}, & L_7 &= \frac{\beta_v - 1}{\beta_v + 1}, & L_8 &= \frac{\beta_{uv} - 1}{\beta_{uv} + 1}, \\ L_9 &= \frac{\beta_{uv} - \beta_u}{\beta_{uv} + \beta_u}, & L_{10} &= \frac{\beta_{uv} - \beta_v}{\beta_{uv} + \beta_v}, & L_{11} &= 1+u+v, & L_{12} &= \frac{\beta_{uv} - \beta_u \beta_v}{\beta_{uv} + \beta_u \beta_v}, \end{aligned} \quad (4.3)$$




Figure 2: The outer-mass planar double box of Caron-Huot and Henn [32], the calibration integral of this section. The six perimeter propagators carry a common mass m (thick red), and the central propagator and the four external legs are massless (thin black). The channels are $s = (p_1 + p_2)^2$ and $t = (p_2 + p_3)^2$, with dimensionless ratios $u = 4m^2/(-s)$ and $v = 4m^2/(-t)$. After the leading singularity of eq. (4.1) is removed the integral is a pure weight-four polylogarithm in (u, v) , which we recover in this section as a single basis vector with coefficient 1.

where β_v enters only through L_7 , L_{10} and L_{12} , and each letter is defined up to a constant factor, which drops out of $d \log L_i$. Moreover, in the variables of eq. (4.2) every letter is a product of integer powers of the twelve multilinear polynomials

$$\{w, z, 1-w, 1+w, 1-z, 1+z, w-z, w+z, 1-wz, 1+wz, 1-w+z+wz, 1+w-z+wz\}, \quad (4.4)$$

and we exploit this factorization in the integrability step below. The maximal cut has no elliptic or Calabi–Yau component, so the only non-rational objects in the problem are the three square roots.

The two earlier computations determine the symbol of g_{10} in different ways. In ref. [68] the alphabet is obtained from the Landau singularities [82, 98], and integrability, first-entry and square-root parity conditions reduce the weight-four symbol space until the remaining freedom is fixed by the maximal cuts. In ref. [12] a conjectured symbol basis is fit by lattice reduction [89] to high-precision values of I_5 computed with **AMFlow** [91] at “a set of 200 kinematic points,” at about 80 CPU-minutes each, with the **AMFlow** precision set to 30 digits. Both computations lead to the same closed form, which is the target of our calibration.

4.2 Alphabet recovery, constraints and the fit

All computations in this section were performed with the programs of section 7. These include a Julia implementation of the ansatz construction of ref. [68] in exact arithmetic [58], **SOFIA.jl** with its **Effortless** alphabet predictor [38], and a lattice-reduction and PSLQ fitter [89, 55, 116]. The numerical values come from **amflow-cpp** [91] with Kira [77] and Fermat [90]. Following section 3.1, we determine the alphabet from the graph and its mass assignment with **SOFIA.jl**. From the singularity analysis we obtain the seven four-dimensional Landau loci, which give the six even letters. In addition, the analysis yields two cubic second-type loci, which we do not include in the alphabet.

Finding the odd letters required two repairs to our port of the alphabet predictor. First, the filter that removes dependent letters had to be applied within each square-root class separately; applied to all letters at once, it discarded odd letters whose rational prefactor is trivial. Second, the letter L_{12} , which contains all three roots, required a triple-product construction that the pairwise-root search of our **Effortless** port lacked. With these two repairs we obtain twelve odd-letter candidates, the

six odd letters of eq. (4.3) and six further letters containing $\sqrt{u}$ or $\sqrt{v}$. The Regge-limit argument of [68], the same argument that paper uses to arrive at twelve letters, excludes those six, and the surviving set is exactly the alphabet of eq. (4.3).

With the alphabet fixed, we impose three conditions on the space of weight-four words in the twelve letters [64, 46],

$$20736 \xrightarrow{\text{d}\mathcal{S}\wedge\text{d}\mathcal{S}=0} 6993 \xrightarrow{\text{parity } (1,0,1)} 861 \xrightarrow{\text{physical first entry}} 161, \quad (4.5)$$

where each number is the dimension of the surviving space. The first arrow imposes integrability of the symbol $\mathcal{S}$. The second keeps the Galois sector that is odd under $\beta_u \rightarrow -\beta_u$ and $\beta_{uv} \rightarrow -\beta_{uv}$ and even under $\beta_v \rightarrow -\beta_v$, as the prefactor in eq. (4.1) requires. The third admits only the seven first entries with physical branch cuts [43],

$$\{L_1/L_3, L_2/L_4, L_3L_8/(L_5L_{10}), L_4L_8/(L_5L_9), L_6, L_7, L_{12}\}, \quad (4.6)$$

namely the s - and t -channel thresholds and pseudo-thresholds. The three dimensions in eq. (4.5) agree with the weight-four counts of refs. [68, 12]. In ref. [68] the maximal cuts would now select one vector in the 161-dimensional space. We leave them aside and determine the vector by a fit.

We evaluate g_{10} with **AMFlow** [91, 77] on a grid of Euclidean points, selected as described in the next subsection, and fit the 161 rational coefficients by lattice reduction [89, 12]. We obtain

$$g_{10} = \text{LS} \cdot \sum_{a=1}^{161} c_a G_a, \quad c_a = \delta_{a,52}, \quad (4.7)$$

where $\{G_a\}$ is the basis of the 161-dimensional space and LS the leading-singularity prefactor of eq. (4.1). A single coefficient is thus nonzero and equal to one, and our fitted symbol coincides with the symbol of refs. [32, 68] term by term. At the 20 Euclidean points reserved from the fit (the 16 unused grid points and the 4 validation points below) the residuals relative to **AMFlow** values at the precision goals of the grid samples lie between 6×10^{-26} and 1.9×10^{-24} . We also continued the closed form by transport to two physical-region points, $(u, v) = (-\frac{1}{2}, 1)$ and $(-\frac{2}{3}, 2)$, where it agrees with complex **AMFlow** values to within 2×10^{-23} in both real and imaginary parts.

Four Euclidean validation points,

$$(u, v) \in \left\{ \left(\frac{16}{3}, \frac{10}{3}\right), \left(\frac{10}{3}, \frac{16}{3}\right), \left(\frac{23}{4}, \frac{9}{4}\right), \left(\frac{5}{2}, 4\right) \right\}, \quad (4.8)$$

are not used in any fit. We recomputed them with **AMFlow** at precision goals 40 and 52, higher than the goals 25 and 37 of the grid samples. The prediction of eq. (4.7) agrees with these values to at least 39 digits at all four points, which satisfies the acceptance criterion of eq. (3.6).¹ The agreement is limited by the accuracy of the goal-40 reference values. Two sample values of g_{10} are

$$g_{10}(u=\frac{16}{3}, v=\frac{10}{3}) = 0.043932114\dots, \quad g_{10}(u=v=4) = 0.053038629\dots, \quad (4.9)$$

the second of which is the value at the symmetric point quoted in ref. [32].

¹The evaluator supplied with this integral [19] repeats the comparison at eight points, and one of the eight validation points, $(u, v) = (4, 4)$, is compared against the published value rather than against an **AMFlow** run [32].

4.3 Exact integrability and sample selection

Two steps of our computation differ in implementation from refs. [68, 12]. One is the integrability condition, the leftmost arrow of eq. (4.5). In ref. [68] denominators are cleared so that each square root appears linearly, and the condition, expanded in the 110 independent monomials that remain, is then row-reduced weight by weight. In ref. [12] the wedge products are evaluated at random numerical points and the null space is found by lattice reduction. We instead apply the substitution of eq. (4.2) before building the symbol space, after which each letter is represented exactly by its vector of integer exponents over the twelve factors of eq. (4.4). The wedge products of the dlog forms of these factors span a space of dimension exactly 35. The condition $d\mathcal{S} \wedge d\mathcal{S} = 0$ is then a finite linear system over $\mathbb{Q}$, and we compute its null space in integer arithmetic. The full cascade of eq. (4.5), from 20736 words to the 161-dimensional space, runs in 20 seconds of wall time on an otherwise idle machine. Neither earlier paper quotes a timing for this step, so the 20 seconds is reported as our measurement, not as a speedup.

The other step that differs is the choice of numerical samples (Table 1). We computed 46 AMFlow samples in total (40 on a Euclidean grid, the 4 validation points of eq. (4.8) and 2 physical-region points) and used 24 of them in the fit. A single sample takes about 6 minutes of wall time on 12 threads, both runs of the staggered pair included. Each sample is computed twice: the grid samples at precision goal 25 to order ε^4 and at goal 37 to order ε^6 , and the validation points at goals 40 and 52 with the same two depths.

The number of fit points, 24, is set by the conditioning of the lattice reduction. When all 40 grid points are included, reduction of the lattice built from the sample values as in ref. [12] returns only the trivial all-zero relation and the correct relation is lost. To locate the usable range we vary the number K of points included in the fit. In a first series of fits we recover the correct relation for all $10 \leq K \lesssim 30$. In a second series we find it from $K = 5$ up to $K = 30$, and only the trivial all-zero relation from $K = 32$ on. We therefore cap the fit at 24 points and reserve the remainder for validation. In the scaling study of ref. [12] (Example 2 there), adding points “helps, but saturates,” with little further change from 50 points to 362, and for the present integral that work uses 200 points. For the fit of eq. (4.7), however, we find that points beyond the usable range of K degrade the reduction.

	Ref. [12], Example 4	this work
numerical samples	200	46 (24 in the fit)
cost per sample	~ 80 CPU-min	~ 72 CPU-min (6 min wall, 12 threads)
accepted value per sample	requested precision (30) trusted as set	only digits on which two runs at different goal and ε -depth agree
fit precision	not stated for Example 4	pair-agreed digits per sample
number of fit points	fixed budget	measured window $10 \leq K \lesssim 30$; cap 24
outcome	exact symbol	same exact symbol, $c_{52} = 1$, $\kappa = 1$
deepest quoted agreement	30-digit AMFlow setting	minimum 39.79 digits (4 validation points)

Table 1: Numerical sampling for the same integral in Example 4 of ref. [12] (middle column, as reported there) and in this work (right column). Both recover the same exact weight-four symbol. The fit here uses 24 of 46 samples (the earlier computation used 200), each sample is accepted only where two differently configured **AMFlow** runs agree, and the number of fit points is set by a measured conditioning window. The four validation points are those of eq. (4.8), and κ is the overall factor between the fitted symbol and that of ref. [68]. Per-sample costs were measured on different machines.

5 The nonplanar two-loop W -pair family

He et al. [69], who “focus primarily on planar topologies”, computed the two-loop master integrals for $q\bar{q} \rightarrow W^+W^-$ with exact top-quark mass dependence for the three planar top-quark branches $\mathcal{T}_1$, $\mathcal{T}_2$ and $\mathcal{T}_3$. Here we compute the nonplanar branch $\mathcal{T}_4$ (their family F_C), drawn in Figure 3. The family has 76 master integrals, of which 23 are new and 53 are shared with the planar branches. We obtain all of them in explicit form on a one-dimensional kinematic line, from ε^{-3} through ε^1 . Each of the 79 integrals of our working basis (the 76 masters and three auxiliary integrals) reproduces independent **AMFlow** values at two points not used in the construction to at least 39 digits.

The computation of F_C is larger than the calibration of section 4, since the differential equation on the line couples 79 components in 36 diagonal blocks. One sector has an elliptic maximal cut, and by the criterion of eq. (3.9) its integrals cannot be determined by a coefficient fit. Most of the result is obtained by transporting the differential equation along the line from one boundary point, subject to the acceptance criterion of eq. (3.6). We also give closed forms for the blocks whose solutions are polylogarithmic and for five of the six radical blocks defined below. The full expressions are given in the corresponding entry of ref. [108], and Appendix A lists the alphabet.

5.1 Family, kinematics and reduction

We use the conventions of ref. [69]. All external momenta are incoming, with $p_1^2 = p_2^2 = 0$ for the quarks and $p_3^2 = p_4^2 = m_W^2$, and the nine propagators of the family are

$$\begin{aligned}
D_1 &= \ell_1^2, & D_2 &= (\ell_1 + p_1)^2, & D_3 &= (\ell_1 - p_2)^2, \\
D_4 &= \ell_2^2, & D_5 &= (\ell_2 + p_1)^2, & D_6 &= (\ell_2 - p_2)^2, \\
D_7 &= (\ell_1 - \ell_2)^2, & D_8 &= (\ell_1 - \ell_2 - p_4)^2 - m_t^2, & D_9 &= (\ell_2 - p_2 - p_3)^2 - m_t^2,
\end{aligned} \tag{5.1}$$

where D_4 and D_5 appear only as numerators (irreducible scalar products) and the top sector is $[1, 1, 1, 0, 0, 1, 1, 1, 1]$. The dimensionless variables are

$$x = -\frac{s}{m_t^2}, \quad y = -\frac{t}{m_t^2}, \quad z = -\frac{m_W^2}{m_t^2}, \quad s + t + u = 2m_W^2, \tag{5.2}$$




Figure 3: The nonplanar top-quark branch $\mathcal{T}_4$ of two-loop $q\bar{q} \rightarrow W^+W^-$ with exact top-quark mass (family F_C of ref. [69], their fig. 1d), with top sector $[1, 1, 1, 0, 0, 1, 1, 1, 1]$ in the propagators of eq. (5.1) and all momenta incoming. Thin black lines are the massless propagators D_1, D_2, D_3, D_6, D_7 , the two violet lines carry the top-quark mass (D_8, D_9), and the W bosons (blue, wavy) attach with $p_3^2 = p_4^2 = m_W^2$. The lines D_7 and D_9 cross without a vertex, which is the nonplanar feature.

and we work on the line $(x, z) = (17/2, 5/3)$ with y symbolic, taking boundary data at $y_0 = 1/3$. The line lies in a continuation region ($s < 0$ and $m_W^2 < 0$, with u above the $t\bar{t}$ threshold) on which 36 of the 79 basis integrals introduced below are complex.

Because the dependent momenta p_3 and p_4 appear in the massive propagators D_8 and D_9 , the list in eq. (5.1) defines the family only together with the momentum convention. Our first definition of the family combined the all-incoming propagator strings of ref. [69] (where $p_4 = -p_1 - p_2 - p_3$) with an outgoing- W convention ($p_4 = p_1 + p_2 - p_3$). Such a mixture defines a different integral family, with 98 master integrals instead of 76 and a different maximal-cut geometry. Its IBP reduction, numerical values and differential equation are consistent with one another, and the error cannot be detected from within that family. We found the mismatch by imposing momentum conservation at each vertex of Figure 3. The corrected family coincides with that of ref. [69]. Its graph matches their fig. 1d edge by edge, including the $D_7 \times D_9$ crossing, and a unique affine map between the loop momenta takes our second Symanzik polynomial onto theirs and induces the identity permutation of the propagators. In particular, our AMFlow [91] values agree with their ancillary MPL results, evaluated with GiNaC [13, 122], to 47–49 digits on four masters at their benchmark point.

IBP reduction with Kira [77, 94] onto a Laporta basis [83] gives 76 master integrals. Of these, 53 can be mapped onto masters of the planar branches, and the other 23, in sectors 480–487 of eq. (A.1) (those containing all four of D_6, D_7, D_8 and D_9), are new. Our working basis contains three further sector-487 integrals with numerators, so the system on the line has 79 components,

$$\partial_y \vec{f}(d, y) = A(d, y) \vec{f}(d, y), \quad A_{ij} \in \mathbb{Q}(d, y), \quad i, j = 1, \dots, 79. \quad (5.3)$$

We obtain $A(d, y)$, where $d = 4 - 2\varepsilon$ is the spacetime dimension, from numerical reductions. Running Kira at 91 numerical points (d, y) , we reconstruct each of the 1240 nonzero entries as a rational function, by Thiele’s continued fraction in y (eq. (2.17)) at 24 values of d followed by rational reconstruction in d . The same reconstruction is performed by FireFly [76, 75] inside a reduction, whereas here we apply it to the already reduced system. As a check, the reconstructed matrix reproduces direct Kira output exactly at three further values of d kept out of the reconstruction, and again at one more value of d combined with two values of y ($d = 890/223$, $y = 29/211$ and $y = 17/139$). In the rest of this section we use the resulting matrix $A(d, y)$ as exact input; the reconstruction procedure is described in section 7.2.

5.2 Alphabet and maximal-cut geometry

We obtain the alphabet from the graph as in section 3.1, here with the graphical Landau analysis of ref. [68] over the full face lattice and with the principal Landau determinant [98, 57, 38], and we check their output with the Euler-characteristic test of ref. [35]. The alphabet has 36 letters (32 rational and 4 of square-root type), listed in eq. (A.2) and eq. (A.3). Four of these letters, among them the pseudo-threshold $x + y - z + 1$, were found only by the Euler-characteristic test, which is why we take the union of its output with that of the two Landau computations. On the line, the polylogarithmic expressions constructed below involve 15 letters in y (13 rational points and one square-root pair).

Exactly one sector has an elliptic maximal cut. In sector 481 = [1, 0, 0, 0, 0, 1, 1, 1, 1], which has six master integrals, the loop-by-loop Baikov representation [123] leaves two integrations, over a Baikov variable z_4 and an irreducible scalar product z_5 , and the integrand on the maximal cut is $1/\sqrt{g_1 G_2}$ with

$$Y^2 = g_1(z_4; z_5) G_2(z_4; z_5), \quad g_1 = -\frac{1}{4} [4z z_4 + (y - z_5 + 1)^2], \quad G_2 = a z_4^2 + b z_4 + c, \quad (5.4)$$

where $a = \frac{1}{16} x(x-4z)$, $b = \frac{1}{8} x(xy-x-2yz+yz_5-2y+2z^2+zz_5+2z)$ and $c = \frac{1}{16}(xy+x+yz_5-zz_5)^2$. Since $g_1 G_2$ is cubic in z_4 , eq. (5.4) is an elliptic curve for each value of z_5 , and it is not the curve of the planar branch $\mathcal{T}_3$ [69]. At a reference fiber over y_0 its j -invariant is $j \simeq 23625.44$, and the curve has no complex multiplication. As a family over z_5 it is a rational elliptic surface, with singular fibers of Kodaira types $2 \times I_1 + 5 \times I_2$ (seven in all, Euler number 12) and Mordell–Weil rank 3. This fiber configuration, together with the positive Mordell–Weil rank, excludes an identification with the universal curve over $\Gamma_1(5)$, and the curve is not modular for any congruence subgroup. Its Picard–Fuchs operator in z_5 has order two and its Wronskian has a closed form.

The curve of sector 481 degenerates as $z \rightarrow 0$ but not at the point $(y, z) = (0, 1)$ at which the boundary values of the planar branches are fixed in ref. [69]. That point thus provides no simple boundary condition for the sector. Moreover, since the monodromy group is not a congruence subgroup, no ring of cusp constants is known in which boundary constants could be sought. By the criterion of section 3.4 the sector is then in class C, and we obtain its integrals by transport of the differential equation [107]. Indeed, two conditions that might have fixed its six constants per order without boundary data are too weak on the line. The sector-481 block of $A(d, y)$ has no pole at $y = 0$ in any entry, so regularity there does not constrain the constants. Matching the solution at y_0 to the ε -expansion of the period integrals of the cut over z_5 gives a linear system whose rank is at most two (of six) per order and zero at the pole orders. The six constants are therefore part of the boundary data described below.

5.3 Block structure and order-by-order solution

On the line the connection of eq. (5.3) is block lower triangular (Figure 4),

$$A(d, y) = \begin{pmatrix} A_1 & 0 & \cdots & 0 \\ * & A_2 & & \vdots \\ \vdots & & \ddots & 0 \\ * & \cdots & * & A_{36} \end{pmatrix}, \quad A_k \simeq \begin{cases} \varepsilon \sum_a C_{k,a} \partial_y \log \alpha_a(y) & 29 \text{ blocks (dlog)}, \\ \text{radical} & 6 \text{ blocks}, \\ \text{elliptic, eq. (5.4)} & 1 \text{ block (sector 481)}, \end{cases} \quad (5.5)$$

where in the first case the block is equivalent, by a rational change of basis, to the ε -factorized dlog form shown [71], with constant matrices $C_{k,a}$ and letters α_a from the line alphabet of section 5.2. Of the 29 dlog blocks, 28 ε -factorize in the Laporta basis and sector 487 does so once the basis

element $[1, 1, 1, -2, 0, 1, 1, 1, 1]$ is replaced by $[1, 1, 1, -1, -1, 1, 1, 1, 1]$. We call the six blocks radical because their homogeneous solutions involve square roots of polynomials in y . The lower triangle couples each block to its subsectors only.

The 40 basis integrals whose subsector tree lies inside dlog blocks are pure combinations of multiple polylogarithms [63, 122]. We solve their blocks of eq. (5.5) order by order in ε and write the result as words $G(a_1, \dots, a_w; y)$ over the 15 letters in y introduced above, with the constants fixed at y_0 . The words reach depth five for the 34 integrals coupled through subsectors and depth six for a tower of 12 among those. The remaining six decouple from the others and are known exactly.

In five of the six radical blocks (sectors 417, 421, 449, 453 and 482, fifteen basis integrals) the reduction of the ε -dependent block toward ε -form by Moser-type balances [86, 104, 62] generates square roots that obstruct a rational ε -form. We have no ε -factorized form for these blocks and solve them order by order in ε instead. Writing each component as a Laurent series $f = \sum_K f_K \varepsilon^K$ and expanding $A = \sum_{k \geq 0} A^{(k)} \varepsilon^k$, eq. (5.3) becomes

$$f'_K(y) = A^{(0)}(y) f_K(y) + \sum_{k \geq 1} A^{(k)}(y) f_{K-k}(y), \quad (5.6)$$

so every order obeys the same homogeneous equation with a source built from lower orders, and

$$f_K(y) = \Phi(y) \left[c_K + \int_{y_0}^y \Phi^{-1}(y') \sum_{k \geq 1} A^{(k)}(y') f_{K-k}(y') dy' \right], \quad \Phi' = A^{(0)} \Phi, \quad (5.7)$$

with Φ a fundamental matrix of the ε^0 block and c_K the boundary constants at y_0 . Since the coefficients $A^{(k)}$ are rational in y , the square roots enter only through Φ , and their radicands reduce to two quadratic polynomials in y with complex roots. For these five blocks we evaluate the integrals of eq. (5.7) in closed form, through order ε^2 for sectors 417, 449 and 482 and through ε^1 for sectors 421 and 453. In each case we confirm in exact arithmetic that the closed form satisfies eq. (5.6) identically. The ε^2 order of the last two sectors, which lies outside the range ε^{-3} through ε^1 of the result, is not solved. The two blocks that receive sector 481 as a source, the sixth radical block (sector 483) and the block of sector 487, are not themselves elliptic. The homogeneous solutions of sector 483 are three rational columns and a pair with nested square roots over four further letters in y , and those of sector 487 are rational at $d = 4$.

The sector-481 block is also solved order by order in ε , on a rank-six basis of homogeneous solutions. At ε^0 the block is reducible to first order on this line, meaning that its operator factorizes into first-order factors with exponents in $\frac{1}{2}\mathbb{Z}$ at the letters of the sector. The ε^0 solutions are thus d'Alembertian, that is, iterated integrals of algebraic functions [108]. The full ε -dependent block is regular singular at every singular point, and at each of three rational values of d that we examined, a rational gauge transformation brings it to Fuchsian form. Its local exponents lie in $\mathbb{Z} + \mathbb{Z}\varepsilon$ at every letter except the two quadratic letters

$$Q_1 = 225y^2 + 5370y + 19801, \quad Q_2 = 900y^2 - 12180y - 7751, \quad (5.8)$$

where the exponent is $-3/2$ independently of d . Hence the block has no ε -factorized form with rational letters on the line, although one over the function field extended by $\sqrt{Q_1}$ and $\sqrt{Q_2}$ is not excluded [7, 60].

The rank-six basis of homogeneous solutions for sector 481 consists of two integrals of the cut over independent cycles (Lefschetz thimbles), two exact rational solutions, and two completion columns fixed by exact initial conditions at y_0 . Searching over rational, square-root and logarithmic trial




Figure 4: Block structure of the connection $A(d, y)$ of eq. (5.3) on the line, with 29 dlog blocks (blue), 6 radical blocks (orange, labeled by sector) and 1 elliptic block (sector 481, violet); the shaded lower triangle is the one-way coupling through subsector sources. Blocks are grouped by class for display and their sizes are not to scale. Sectors 483 and 487 receive sector 481 only as a source (the violet entries in their block rows), so the six integrals of sector 481 enter the equations for the other 73 basis integrals only through the lower triangle.

functions built from the letters in y , we find that the completion columns are of none of these types. The curve has rational two-torsion, so its j -invariant factors through the Hauptmodul of $X_0(2)$ (denoted t in the next equation, not to be confused with the Mandelstam invariant of eq. (5.2)),

$$j = \frac{(t + 256)^3}{t^2}, \quad (5.9)$$

but this parametrization does not by itself lead to an Eichler-integral representation of the solutions [16, 7]. We solve the block instead by the variation of parameters of eq. (5.7) on the thimble basis.

5.4 Boundary data

Transport of eq. (5.3) requires the values of all 79 components at the boundary point $y_0 = 1/3$. Fixed- d auxiliary-mass-flow runs [91], one per value of d , take more than four hours each for this family. Instead we use a single AMFlow run at y_0 , which takes 61 minutes and returns the Laurent expansion in ε through order 22, and we generate all boundary data from it. From the Laurent series we form the values of $\vec{f}$ at 24 nodes $\varepsilon = 1/q$, with q prime on an approximately geometric grid from 1013 to 10037, where the truncated tail is below 10^{-58} . We then transport each fixed- ε vector along the line by repeated Taylor expansion of eq. (5.3), in the manner of refs. [73, 10] but on the exact connection, at expansion order 60 with 500-digit arithmetic. Each such transport takes about 120 steps, whose size is set by the distance to the nearest complex singularity and by a bound on the tail. Transporting out and back again reproduces the starting vector to within 9.9×10^{-126} .

To compare with an ε -expanded reference value at the endpoint we must re-extract the Laurent coefficients from the 24 fixed- ε vectors. The accuracy of this extraction depends on where the nodes are placed in ε . With a narrow band of nodes the extracted coefficients are good to only 25 digits, because the Vandermonde system of the extraction is ill conditioned, while with a wide band they are contaminated by the first neglected Laurent coefficient. We therefore choose the node layout using a transport from y_0 out and back again, in which only the extraction error remains. For the geometric grid above we find 39-digit coefficients in this out-and-back test, the same accuracy as at the validation points.

We then replace the numerical boundary values from the single **AMFlow** run, as far as we can, by quantities with an independent definition. Of the 79 components of the boundary vector, 23 are limits of the MPL expressions in the ancillary files of ref. [69] for the shared masters, 11 have closed forms and 5 have explicit integral representations, both derived in this work. The rest include the radical blocks and sector 481, and for the result as a whole 40 of its 79 boundary constants are defined by a standalone written-integral transport rather than by closed forms. The definition is the auxiliary-mass-flow representation itself [91]. With an auxiliary mass η in the propagators, the basis obeys a differential equation in η whose boundary values at $\eta \rightarrow \infty$ are single-scale vacuum integrals (products of Γ functions). The physical value is recovered in the limit $\eta \rightarrow i0^-$, as in eq. (2.15). We set

$$B_i \equiv \lim_{\eta \rightarrow i0^-} I_{\text{aux},i}(\varepsilon, \eta) \Big|_{y=y_0}, \quad \partial_\eta \vec{I}_{\text{aux}}(\varepsilon, \eta) = A_\eta(\varepsilon, \eta) \vec{I}_{\text{aux}}(\varepsilon, \eta), \quad \vec{I}_{\text{aux}}(\varepsilon, \eta) \xrightarrow{\eta \rightarrow \infty} \vec{I}_{\text{vac}}(\varepsilon, \eta), \quad (5.10)$$

write out the rational matrix A_η and the vacuum values $\vec{I}_{\text{vac}}$ once, and evaluate B_i to any requested precision by transporting eq. (5.10) from $\eta = \infty$ to $\eta \rightarrow i0^-$ with a solver that uses neither **AMFlow** nor any reduction code at run time. Evaluated this way, the boundary constants are no longer limited by the precision of the original **AMFlow** run. Two evaluations of eq. (5.10) at different working precisions agree to 157 digits at the worst of the 24 nodes, whereas the Laurent coefficients of the original run are truncated after 49–51 significant figures.

5.5 Validation

We transport the assembled solution from $y_0 = 1/3$ to $y = 1/2$ and to $y = 3/5$ and compare it there, order by order in ε , with independent **AMFlow** values at those two points [91], neither of which enters the construction. The comparison in Table 2 starts at order ε^{-4} , where all 79 coefficients vanish identically, so six orders are compared (all 79 basis integrals at ε^0 and above and, at the negative orders, those integrals whose expansions reach that deep). The window of the deliverable is ε^{-3} through ε^1 ; ε^2 was compared and descope. The smallest agreement over the six orders and both points is 39 digits.

The lower part of Table 2 shows that the rational coefficients of the symbolic forms are exact. When the polylogarithmic words, the radical-block solutions and the sector-481 block are evaluated and transported, their agreement with the reference values increases with the working precision. With one wrong rational coefficient it would instead saturate at the size of that coefficient's error. The polylogarithmic words also agree, order by order in ε , with the values obtained by transporting the Laurent expansion at y_0 , and, for the shared masters, with the planar-branch expressions of ref. [69] restricted to the line. For the 36 complex basis integrals we also compared the imaginary parts, which agree to at least 42 digits at both points.

The whole result lives on the one-dimensional line $(x, z) = (17/2, 5/3)$; the extension to full three-variable kinematics is not attempted here. The depth in ε is set by the reference values, which lose about 7.1 digits per order relative to their working precision, so orders beyond ε^2 are

Laurent order	integrals compared	minimum agreement (digits)	
		$y = 1/2$	$y = 3/5$
ε^{-4}	0 (identically zero)	–	–
ε^{-3}	11	59.0	59.4
ε^{-2}	61	53.8	53.8
ε^{-1}	73	49.7	49.8
ε^0	79	44.1	44.1
ε^1	79	41.6	42.1
ε^2	79	39.2	39.1
assembled symbolic forms, worst case over the six orders and both points			
evaluated at 40-digit working precision	79		37
evaluated at 60-digit working precision	79		57
sector 481 on its 90-digit evaluation chain	6		78–90

Table 2: Every basis integral of the F_C family on the line agrees with independent **AMFlow** values at two kinematic points not used in its construction to at least 39 digits, order by order in ε . Upper part: for each Laurent order, the number of the 79 basis integrals whose expansion is nonzero at that order and the smallest number of agreeing digits among them at $y = 1/2$ and at $y = 3/5$ (to one decimal); the ε^2 row lies outside the window ε^{-3} through ε^1 of the result. Lower part: the same comparison for the assembled symbolic forms at two working precisions, and for the six integrals of the elliptic sector 481 on their own evaluation chain. In every row the re-extracted Laurent coefficients reproduce the transported values at the four verification nodes of the ε grid (Appendix A) to at least 52 digits.

transported but not compared. Per-integral closed forms for the η -transport boundary constants remain open.

On the integral’s page of the BootLoops site (bootloops.ai/diagrams/) [19] we provide a standalone script, built on **mpmath** [118], that evaluates the family at any y on the line to a requested precision. It contains the exact connection of eq. (5.3), the symbolic forms and the boundary data, and uses them to transport the solution along the line to the chosen point.² The extension off the line is likely to require an ε -factorized form of the full system in (x, y, z) with the algebraic letters of the radical sectors, boundary data at a point where the curve of eq. (5.4) degenerates, and a representation of the sector-481 solutions as elliptic multiple polylogarithms on that curve [25] in place of transported values.

6 Constraints by function class

The constraints of section 3.2 (integrability, first entry, regularity away from the Landau locus, graph symmetry) hold for every Feynman integral, because they are statements about the differential equation and the graph. The geometry of the maximal cut determines whether the polylogarithm-specific constraints have an analog, whether a coefficient fit can fix the surviving unknowns (section 3.4), and in which ring the boundary constants lie. In each class we use an established formalism, namely the symbol constraints of the planar amplitude bootstrap [2, 43, 31, 44], the elliptic symbol calculus and the symbol-level elliptic bootstrap [24, 99], and Calabi–Yau period geometry [49, 93, 45]. Table 5

²By default it starts from stored values at the 24 boundary nodes instead of re-running eq. (5.10). Two recomputations of the stored values at different working precisions agree to 124 digits, which caps the default working precision near 109 digits. The η -transport chain that defines the nodes is not distributed with the script (Appendix B).

lists, for each of the three classes, which constraints apply and what replaces the polylogarithmic ones that do not.

6.1 Polylogarithmic families

For iterated integrals of $d \log$ forms the alphabet determines the function space, and the symbol-level constraints reviewed in section 2.2 apply [64, 46, 114, 31, 44]. We illustrate them on two families, the first of which is the three-loop crossed light-by-light box with a massive fermion loop [108]. On the maximal cut its alphabet is the seven letters

$$\mathcal{A}_{\text{cut}} = \{s, t, u, s - 4m^2, t - 4m^2, \sqrt{s(s - 4m^2)}, \sqrt{t(t - 4m^2)}\}, \quad (6.1)$$

where $u = -s - t$, m is the loop mass, and the two square roots rationalize independently. In Table 3 we give the dimension of the weight- n symbol ansatz over $\mathcal{A}_{\text{cut}}$ after each constraint (computed by Gaussian elimination over finite fields). Integrability already removes about half of the words at weight four and more at weight five, and the first-entry condition a further $\sim 36\%$. At weight four the Steinmann relation on the second entry, the last-entry condition and the $s \leftrightarrow t$ symmetry of the graph then leave 34 functions. We did not impose extended Steinmann (full adjacency), which is conjectural for a general Feynman integral [31, 44]. Single-variable blocks are much smaller from the start and, since integrability is automatic for them, only the first-entry condition reduces the count further. For example, the t -only harmonic-polylogarithm tower over $\{0, 1, -1\}$ goes from 3^n words to 1, 2, 6, 18, 54, 162, 486 through weight six, with 197 irreducible Lyndon generators in total. The remaining unknowns, of order 10^2 in number, are well within reach of the coefficient fit, provided the alphabet is complete and the function is polylogarithmic.

However, the full light-by-light family is not polylogarithmic [108]. The counts in Table 3 therefore describe only the polylogarithmic ansatz over eq. (6.1). We find that the alphabet of its top-sector connection is the eight rational letters $\{s, t, s + t, s + t + 4, s + t + 16, s - 4, t - 4, 2s + t\}$ (with $m^2 = 1$). Moreover, a four-propagator subsector is the equal-mass three-loop banana, whose maximal cut is a K3 surface [103, 15]. Thus the top master is a mixed polylogarithmic and K3 function, which we evaluated by transporting the differential equation in m^2 and validated at six independent points [108]. We obtained the u -only block, sampled in t on the slice $s = -3$, in closed form through weight two, in agreement with AMFlow [91] to at least 113 digits at 19 points. At weight three a fit in the raw reduction basis was limited by poor conditioning to about five digits, and this failure motivated the normalization of section 3.3.

The second family is the two-loop double box with a massive central propagator (the topology labeled C in ref. [106], nineteen masters, known in multiple polylogarithms through weight four). Its alphabet is seven rational letters with no square root [108]. Reading the alphabet off the maximal cuts loop by loop, we miss the rational letter $m^2(s + t) - st$ and obtain two square roots that appear in no master integral of the family. We found that a value fit over that incorrect alphabet is ill-conditioned, with a residual that does not decrease as the working precision is raised. Adding a candidate eighth letter $s^2 - t$ lowers the residual without fixing the function, and in fact $s^2 - t$ is absent from the family's differential equation. Over the correct seven letters, with integrability and the physical first entries imposed, we obtain weights ≤ 2 as an eight-term formula with rational coefficients. It reproduces the reference values to 108 digits at eight kinematic points, four of which were not used in any fit. We obtain weights three and four as two-variable word lists of 39 and $123 + 29$ terms, with constants in $\mathbb{Q}[\zeta_2, \zeta_3, \zeta_4]$ and logarithms, identified by integer-relation fits [55, 89]. These lists satisfy eq. (3.6) by a similar margin at twenty further points.

The constants in this class are multiple zeta values (MZVs), graded by weight, with the first irreducible depth-two constant $\zeta(5, 3)$ at weight eight [22]. Their conjectural dimensions obey $d_k =$

weight n	words 7^n	after integrability	after first entry
1	7	7	5
2	49	39	25
3	343	213	135
4	2401	1166	740
5	16807	6388	4056
weight 4, continued: Steinmann (second entry) $740 \rightarrow 283$; last entry $\rightarrow 68$; $s \leftrightarrow t$ symmetry $\rightarrow 34$			

Table 3: Dimension of the symbol ansatz for the three-loop crossed light-by-light box after each constraint, over the seven maximal-cut letters of eq. (6.1). Integrability removes about half of the weight-four words and first entry a further $\sim 36\%$; the three remaining polylogarithmic constraints then reduce 740 unknowns to 34. Extended Steinmann (full adjacency), conjectural for a general Feynman integral, would reach 69 before the symmetry and was not imposed. Dimensions are ranks over finite fields (two primes through weight four, one at weight five).

$d_{k-2} + d_{k-3}$ [127] and a motivic basis is known [27]. Since the algebraic independence of the odd zeta values is conjectural [127, 8], so is the linear independence of the basis of constants used in the integer-relation fit in this ring. A hidden relation among the basis constants would make a printed decomposition non-unique, but it could not invalidate an accepted result, because eq. (3.6) tests the function itself against independent numerical values.

6.2 Elliptic families and value-fittability

When the maximal cut is an elliptic curve, the dlog symbol is replaced by iterated integrals of the Kronecker–Eisenstein kernels of section 2.2, graded by modular weight [24, 25, 7]. The flatness condition (2.5) then becomes a set of relations among Eisenstein series of adjacent weight. For the equal-mass two-loop sunrise the maximal cut in $t = p^2/m^2$ is the universal family of elliptic curves over the modular curve of $\Gamma_1(6)$,

$$\mathcal{E}_t \longrightarrow X_1(6) = \overline{\mathbb{H}/\Gamma_1(6)}, \quad t \in \{0, 1, 9, \infty\} \text{ singular, with Kodaira fibers } I_2, I_3, I_1, I_6, \quad (6.2)$$

and the integral is an iterated integral of modular forms for $\Gamma_1(6)$ [16, 6, 115]. The grading by modular weight and the first-entry condition restrict the ansatz to kernels of low weight. Because the cusp-form spaces $S_2(\Gamma_1(6))$ and $S_3(\Gamma_1(6))$ are empty, the kernels of modular weight up to three are Eisenstein series, whose cusp data are fixed by the Landau singularities [115, 128]. Hence for the sunrise and the equal-mass kite the geometry determines the kernels through weight three without numerical input [5].

However, in this class the constraints, fully imposed, may still fail to fix the function. The obstruction is the second-kind structure of eq. (2.13), and whether it is present in a given master follows from the criterion of eq. (3.9) (see Table 4 for our tests). The equal-mass kite top integral at order ε^0 is class A. Its prefactor and modular variable are

$$P(t) = \frac{1}{4t}, \quad \tau = \frac{\psi_2}{\psi_1}, \quad q = e^{i\pi\tau}, \quad (6.3)$$

where q is the geometric nome of the curve, used in place of the half-shifted nome of the maximal cut. With the depth-two word $\bar{\text{Eb}}_{0,2}$ of ref. [25] in the basis, a value fit against **AMFlow** data gives five rational coefficients. The result satisfies eq. (3.6) at both validation points, with the same agreement there as at the fit points [108]. The same fit in the half-shifted frame and without that

word fails (Table 4). A missing basis element of this kind can be repaired by enlarging the basis, unlike the period obstruction of class C.

The equal-mass ice-cream-cone top integral of ref. [107] is class C. We used 36 sample points at 120 digits, enough that the number of samples does not limit the fit, and put the exact period-weighted weight-three words of the bubble letter L_{bub} in the basis. Even so, the agreement at the validation points remains far below the threshold of eq. (3.6), in both the scalar canonical frame and the matrix ε -factorized frame. For a sub-topology of this integral (the scoop topology of ref. [107], class A/B), by contrast, a fit of the same kind satisfies eq. (3.6). When we add words to the top-integral basis the agreement saturates, and the three second-kind words, the last of them η_1 itself, do not raise it further. The saturation follows from the form of the variation-of-parameters solution, whose particular part contains, schematically, the Eichler-type integral

$$\int \psi_1^2 \frac{\eta_1}{\psi_1} (\dots) f, \quad f \text{ a modular form for the monodromy group}, \quad (6.4)$$

which is not weight-graded and is frame-invariant, and therefore cannot be removed by a local change of basis [25, 50, 59]. A weight-three boundary constant of the top integral (denoted $T(-1)$ as in ref. [107]) must then be supplied and the master transported through its differential equation.

We tested the kernel condition (ii) of eq. (3.9), which requires the inhomogeneous kernel $K = \psi_1^3/W$ to be a modular q -series, on the equal-mass and unequal-mass kite families. On the equal-mass kite, with monodromy group $\Gamma_1(6)$, we use only the geometry (the exact Wronskian, no `AMFlow` input) and reproduce K with two words and a rational coefficient, while the coefficient of the second-kind period vanishes numerically. On the unequal-mass kite, whose curve is not of congruence type, the agreement in the same test rises only slowly as the basis is enlarged from two to thirteen words. For each basis size the agreement is the same at the fit points as at the validation points, as it must be for any target outside the span of the basis, and the second-kind period enters with a nonzero coefficient (Table 4). The unequal-mass kite thus provides a test of the criterion on a case not used to formulate it. By the prefactor its top integral would be class A, whereas by the kernel it is class C. A value fit against `AMFlow` data does fail, with non-rational coefficients, while the same basis fits to high precision a reference function manufactured to lie in its span [108]. This contrast confirms that condition (ii) of eq. (3.9) is necessary.

Fittability is local in the kinematic variable, whereas modularity of the elliptic surface is a global property. Thus a non-modular surface does not by itself exclude a fit. An example is the threshold sunrise of ref. [107], whose surface is globally non-modular. Indeed, the closed form obtained there by a value fit reproduces the reference value at a Euclidean validation point away from the singular fibers, to the precision in the last line of Table 4.

6.3 Identification and derived boundary constants

On the new elliptic integrals of this work the coefficient fit by itself fixed no integral: the results came from identification, derived boundary constants, and transport, carried out on the function space the geometry had fixed. Where a fit did satisfy eq. (3.6), for the equal-mass kite and the sunrise masters, the curve was one whose word basis was already in the literature [5, 6, 24].

One mechanism is identification of the function, as for the two-loop $gg \rightarrow Z\gamma$ masters with a massive loop [108]. There the cusp widths of the maximal-cut curve fix the monodromy group to $\Gamma_0(4)$. Since the relevant cusp-form spaces are empty, the homogeneous block is Eisenstein-only and is a specific cusp-regularized Eichler integral, whose only freedom is two constants fixed by cusp data [128]. The masters then follow by transport in s from the cusp, where all 25 boundary conditions are analytic, so that no constant is fitted numerically [108].

integral	class	test	fit agreement (digits)	control (digits)
equal-mass kite top, ε^0	A	value fit, frame of eq. (6.3), word $\text{Eb}_{0,2}$ included; validation at $t = -3, -10$	70, 71	same fit, half-shifted nome, word omitted: 9
equal-mass ice-cream-cone top [107]	C	value fit, 36 AMFlow points at 120 digits (four of them at ~ 160); scalar frame / matrix ε -factorization	15 / 11; word-by-word build-up saturates at 15.7 from 21 words onward	scoop sub-topology (A/B), same framework: 89–91
equal-mass kite, kernel K ($\Gamma_1(6)$)	A	geometry-only kernel test, exact Wronskian, two words	116 (coefficient -6 ; second-kind $\sim 10^{-108}$)	
unequal-mass kite, kernel K (non-congruence)	C	same test, two to thirteen words	$1 \rightarrow 9$	
unequal-mass kite top	C	value fit against AMFlow	9–12	same basis, manufactured reference: 72
$gg \rightarrow Z\gamma$ elliptic masters [108]	C	value fit against AMFlow	143	positive control: 78
threshold sunrise [107] (globally non-modular surface)	†	closed-form value fit; validation at $t = -3$	149	

Table 4: Coefficient-fit agreement at independent points compared with control fits, by criterion class. Where the criterion of eq. (3.9) allows a constant-coefficient fit (class A or B) the agreement is 70 to 149 digits; where it does not (class C) it is between 1 and 15 digits, while a control fit in the same framework gives 72 to 91. The kernel tests use no numerical input beyond the curve. “Digits” is the agreement with values not used in the fit; the acceptance threshold of eq. (3.6) is 30. †: fittability is a local property of (P, K) at the sample, which this line illustrates (see text).

Alternatively, a boundary constant can be derived from an analytic condition that over-determines it. For the ice-cream-cone family the condition is the absence of a forbidden branch at the infinity cusp. We derived the constants both by cusp matching and from a parametric integral representation. The bubble-subsector and ε^1 constants lie in a ring of polylogarithmic constants built on the golden ratio (the golden dilogarithm ring of ref. [107]), whereas the ε^0 top constant $T(-1)$ matches no named ring at the heights searched. The resulting expressions for the three ice-cream-cone integrals satisfy eq. (3.6) [107, 108]. For a further kite family treated in ref. [108] the condition is boundedness at the vacuum point $p^2 = 0$, a regular singular point at which the residue eigenvalues allow only one bounded solution. For the unequal-mass kite it is regularity at $s = 0$, where the residue matrix is semisimple. Together with the two-loop vacuum values this condition gives all the constants in terms of γ_E , ζ_2 , $\ln 2$ and Catalan’s constant. Transporting from that boundary, we reproduce all 42 components of the AMFlow boundary vector at $s = -2$, and the top integral at independent points to 125–131 digits, with no AMFlow input [108].

In this class the ring of constants depends on the arithmetic of the curve. Critical L -values of modular forms are periods [54, 111]. For a curve with complex multiplication (CM) the Chowla–Selberg formula places the period in the Γ -function ring of the CM field [110], and the period is then

algebraically independent of π [37, 101]. In particular, the constants $\text{Cl}_2(\pi/3)$ (conductor three), Catalan’s constant as $L(\chi_{-4}, 2)$ (conductor four) and the level-15 value $L(f_3, 2)$ below are instances. The ice-cream-cone top constant $T(-1)$ lies off the CM and modular locus, where we know of no dictionary between periods and known constants.

Arithmetic route. For a class-C integral the boundary constant is a period outside the weight-graded span and must be supplied to the transport. In the elliptic case we can obtain it from the Picard–Fuchs operator, because the periods of the curve are governed by the same arithmetic as its point counts over finite fields, which can be read off the operator [81]. Let $\varpi_0(t) = \sum_{n \geq 0} c_n t^n$ be the holomorphic solution of the second-order operator at its MUM point. For a prime p of good reduction the truncated series is the Hasse invariant of the fiber,

$$a_p(t) \equiv [\varpi_0]_p(t) \pmod{p}, \quad [\varpi_0]_p(t) = \sum_{n=0}^{p-1} c_n t^n, \quad (6.5)$$

where $a_p(t) = p+1 - \#\mathcal{E}_t(\mathbb{F}_p)$ is the Frobenius trace, up to a sign fixed by the model of the curve [81]. On the sunrise operator eq. (6.5) agrees with direct point counts for all 10 primes and 4 fibers we tried.

Hecke multiplicativity extends $\{a_p\}$ to $\{a_n\}$. The approximate functional equation of the weight-three, level-15 newform f_3 , which has CM by $\mathbb{Q}(\sqrt{-15})$, then gives us $L(f_3, 2) = 0.88045982\dots$ to 48 digits. With a PSLQ search [55] we identify this value with the Chowla–Selberg period

$$L(f_3, 2) = \frac{\Gamma(\frac{1}{15})\Gamma(\frac{2}{15})\Gamma(\frac{4}{15})\Gamma(\frac{8}{15})}{\pi \cdot 2^3 \cdot 3^{3/2} \cdot 5}, \quad (6.6)$$

a relation that persists at the 55-digit working precision, which exceeds the precision to which the L -value itself has converged. This is the constant of Broadhurst and of Bloch, Kerr and Vanhove [23, 15], in terms of which the equal-mass sunrise special value is $(12\pi/\sqrt{15}) L(f_3, 2) = 8.5702804\dots$. The target is classical and we claim none of it. The new element is the route, in which the recipe of ref. [81] (operator, a_p , Hecke eigenform, L -value, Γ -product) is applied to a Feynman integral. This route fixes a boundary constant from the p -adic Frobenius data of the Picard–Fuchs operator, without evaluating any Feynman integral and without solving for the master vector. It does not yet extend to the half-integer K3 branches, for two reasons. The modular object needed there is the cusp series of a half-integral-weight Eichler integral, which we have not identified [96]. In addition, the numerical connection from the MUM point to threshold crosses three singular fibers ($t = 4, 16, 36$).

Eichler-sector route. By the criterion, the second-kind component of the equal-mass sunrise block, with its $\psi_1^2 \eta_1$ prefactor, is class C and would be obtained by transport, but we can fix that component directly. The quantity to be fixed is

$$\Pi_2(t) = \psi_1(t) \eta_1(t) = \frac{\pi^2}{3} E_2(\tau_F(t)), \quad (6.7)$$

where η_1 is the quasi-period paired with ψ_1 and τ_F the modular parameter of the fiber. By the Legendre relation Π_2 does not depend on the frame.

A linear fit of $E_2(\tau_F)$ in the three-dimensional weight-two Eisenstein space $M_2(\Gamma_0(6))$, supplemented by the squared period and a constant, fails on a five-point Euclidean grid, with non-rational coefficients. The agreement at the validation points is at most 2 digits, and the residual at the

fit points is equally large, as expected for a target outside the span of the basis. Instead we use Ramanujan’s identity

$$DE_2 = \frac{E_2^2 - E_4}{12} \equiv g_4, \quad D = q \frac{d}{dq}, \quad (6.8)$$

which makes $E_2 = c_0 + \alpha \int^q g_4 d \log q'$ the depth-one Eichler integral of the known weight-four form g_4 , with a degree-zero period polynomial and therefore a single pair of constants (c_0, α) . Fitting at $t = -1$ and identifying the two constants with PSLQ, we obtain $c_0 = 1$ and $\alpha = 1$, which recovers the known Ramanujan values. With these fixed, we reproduce Π_2 at five further points ($t = -\frac{1}{2}, -2, -5, -7, -13$) to 99–160 digits, limited by the working precision. In the frame used to compute Π_2 the top sunrise master has, in the holomorphic-period sector, the form

$$J_{111}^{\varepsilon^0}(t) = -\frac{\psi_1(t)}{\pi} (B_0 + I(1, f_3)), \quad (6.9)$$

where B_0 is a constant and $I(1, f_3)$ an iterated integral of modular forms for $\Gamma_1(6)$ built on f_3 [6, 108] (this f_3 , a modular form for $\Gamma_1(6)$, is not the level-15 newform of eq. (6.6)). Equation (6.9) agrees with AMFlow values to more than 99 digits at each of the six Euclidean points where ≥ 260 -digit reference values exist.

Although the ingredients of this Eichler-integral construction of Π_2 are standard [25, 115, 128], it lets us fix directly, to the accuracy of eq. (3.6), a component that the criterion places in class C and that the graded value fit of the same object fails to determine. The first limitation of the construction is that the period polynomial is of degree zero, so the test constrains only (c_0, α) . At degree one the same construction, applied as a modular-forms exercise to the smallest non-CM weight-three newform (LMFDB label 12.3.d.a [117]), fixes its period constants $L(f, 1)$ and $L(f, 2)$ to 71 digits. That newform belongs to no Feynman integral we can reach, however, since every accessible maximal cut has Frobenius traces in the rational integers [107]. Second, $J_{111}^{\varepsilon^1}$, a weight-4 second-kind object, is not delivered: the weight-4 Eichler word basis is incomplete and only 6 AMFlow grid points exist.

6.4 Calabi–Yau periods and summary

When the maximal cut is a Calabi–Yau variety the function space is governed by its period vector [49, 93, 45]. Integrability (2.5) is unchanged, and the canonical-form constructions for these geometries are built on it [49, 93]. First entry and regularity carry over, because the singular locus is still computable from the graph, with the principal Landau determinant and GKZ data giving the banana loci at all loops [80, 18, 79]. Graph symmetry and the acceptance criterion apply as before. The constraint route has not been applied to a genuine Calabi–Yau integral; every Calabi–Yau result in this work was obtained by transport of the differential equation.

In the Calabi–Yau case a further constraint, with no elliptic analog, comes from the structure of the period vector itself. We write z for the complex-structure modulus (the variable x of eq. (2.14)) and reserve t for the flat coordinate, as is customary for Calabi–Yau periods. At the MUM point $z = 0$ the Frobenius basis $\varpi_0, \dots, \varpi_n$ is organized by powers of $\log z$, and in dimension three the mirror map and Yukawa coupling are

$$t(z) = \frac{\varpi_1(z)}{\varpi_0(z)}, \quad q = e^t, \quad Y(q) \propto \theta_q^2 \frac{\varpi_2}{\varpi_0}, \quad \theta_q = q \frac{d}{dq}, \quad (6.10)$$

in the normalization of refs. [49, 45, 29]. The monodromy preserves a pairing on the periods (symplectic in odd dimension, orthogonal for a K3 surface). For the three-loop equal-mass banana,

Constraint	Polylog	Elliptic	Calabi–Yau
<i>General (geometry-independent)</i>			
Integrability $dA=A\wedge A$	applies	applies	applies
First entry (physical channels)	applies	applies (cusps)	applies (MUM/loci)
Regularity (Landau locus only)	applies	applies	applies
Graph symmetry	applies	applies	applies
Validation at independent points (≥ 30 digits)	applies	applies	applies
<i>Polylog-specific (iterated $d \log$)</i>			
Symbol over finite alphabet	yes	$\rightarrow$ elliptic symbol (Kron.–Eis. kernels)	$\rightarrow$ period vector
Shuffle / Lyndon count	yes	no (not word-spanned)	no
Steinmann / ext. Steinmann	yes	$\rightarrow$ low-wt. kernels Eisenstein ($S_2=S_3=0$)	(Hodge/self-duality)
Cluster adjacency	yes	no	no
Last entry	yes	survives in spirit (L_{bub} -type)	loci-restricted
<i>Geometry-new</i>			
Modular-weight grading	–	yes	–
Y -invariant / self-duality	–	–	yes
PSLQ constant ring	MZV ring	cuspid ring (congruence)	<i>none</i> ($\rightarrow$ transport)
<i>Sufficiency (value-fittability)</i>			
Constraints suffice?	yes (alphabet complete)	A/B yes; C no (E_2 /Eichler)	class-B argued; A open $\rightarrow$ transport

Table 5: The constraints by function class. Reading across shows how one constraint changes as the geometry of the maximal cut changes; reading down shows which constraints are available in a given class. The general constraints (top block) hold everywhere. The polylogarithm-specific ones (middle block) either have no analog (“no”) or are replaced by a geometric counterpart ($\rightarrow$). The third block lists the constraints new to each geometry and the ring of constants available to the integer-relation fit. The last line is the sufficiency statement of the criterion of eq. (3.9), with the evidence of Table 4.

whose maximal cut is a one-parameter K3 surface, the Picard–Fuchs operator is the symmetric square of that of the sunrise curve of eq. (6.2) [15, 103], so that

$$(\varpi_0, \varpi_1, \varpi_2) \propto (\psi_1^2, \psi_1\psi_2, \psi_2^2), \quad (6.11)$$

with $\psi_{1,2}$ the periods of the sunrise curve. The Y -invariant relations built from eq. (6.10) constrain the connection beyond flatness and take the place of the modular-weight grading [49] (the semisimple–unipotent split of the period matrix is that of ref. [45]).

For Calabi–Yau cuts, however, no ring of constants is available for the integer-relation fit. At a MUM point the normalization involves the Y -invariants, which are local analytic series and not algebraic numbers. Moreover, there is no enumerable dictionary of Calabi–Yau period constants for a PSLQ search. Whether the normalizer is algebraic (class A) or period-valued is an open question, and we can argue only for the class-B form (a single power of the holomorphic period). Transport, in turn, is limited in precision by the radius of convergence of the series at the MUM point and by the distance to the nearest singular point.

The three-loop equal-mass banana is the simplest Calabi–Yau case, because its K3 is the symmetric square (6.11) of the sunrise curve. The integral then has an ε -factorized form in meromorphic modular forms [103, 48], and our evaluation of it is a calibration on a known function space. Both its closed-form words and a four-master transport satisfy eq. (3.6) at four independent AMFlow points, the former with agreement to at least 110 digits [108]. At the mass configuration

$$\sqrt{M_4} = \sqrt{M_1} + \sqrt{M_2} + \sqrt{M_3} \quad (6.12)$$

the unequal-mass K3 acquires non-unipotent monodromy of order two [107]. In general the unequal-mass three-loop banana requires the full period machinery [51, 105], and transport appears to be the only route.

The four-loop banana with mass partition $\{4, 1\}$ has a Calabi–Yau variety of dimension three on its cut, with an order-six Picard–Fuchs operator [102]. Evaluating the full 28-master vector at the required ε -order by auxiliary mass flow exceeds the memory of a single large-memory machine [107]. On the maximal cut the system reduces to the six top-sector masters, which ref. [107] treats by transport of the periods. Finally, a non-CM weight-three period does not arise in any of the banana integrals we have examined [107]. The three-equal-mass banana K3 is Hilbert modular over $\mathbb{Q}(\sqrt{5})$ [47], but its cut factorizes into two copies of the equal-mass sunrise cut with CM by $\mathbb{Q}(\sqrt{-15})$. Its weight-three period is thus again a CM Γ -period [48]. By Livné’s theorem every rank-two K3 over $\mathbb{Q}$ carries a CM newform, and a non-CM example would therefore need transcendental rank three or more, where no L -value dictionary is available [92, 109].

Most of Table 5 is established material, credited above where each object was used [2, 71, 43, 24, 99, 49]. The comparison itself is new, namely the statement of which constraints are geometry-independent, which are specific to iterated dlog forms, and what replaces the latter in the elliptic and Calabi–Yau cases. The last line of the table, the sufficiency statement supplied by the value-fittability criterion, is also new, with the status stated in section 3.4.

7 Implementation

Each step of the method is carried out with a public code or with a program written for this work, and no result depends on a proprietary computer-algebra kernel. To this end we ported several public packages out of their native languages. With these ports, we and the language model that operates the codes (section 3.5) can modify and test all of them in one software environment. Documentation and tests for each tool are collected in the online registry [19].

7.1 Codes and ports

Table 6 lists the codes by the step of the method they serve. Codes that run without a proprietary kernel are used as released, in particular Kira with its FireFly back end [77, 76, 75]. We drive Kira through its job files and parse its output with our own scripts. Packages whose published version requires Mathematica are replaced by ports. Thus we run the auxiliary-mass-flow method [91] with the C++ program `amflow-cpp` on FLINT. For the singular locus we use `SOFIA.jl`, our port to Julia (on Nemo and FLINT [58]) of the package of Correia, Giroux and Mizera [38], alongside the native Julia package `PLD.jl` [57]. Finally, we rewrote research code published in Mathematica, using Julia with Nemo, Hecke and Arb [58, 42, 74] for exact and ball arithmetic and Python with `mpmath` [118] for high-precision numerics.

The Julia package `LandauBootstrap`, which implements the ansatz, constraint and fit steps, is our rewrite of the Mathematica code of refs. [68, 12]. At most a single Mathematica workstation license is retained, for occasional cross-checks on which no result reported here depends. The transports and basis normalizations reported in this paper were carried out with our own Julia and Python programs; `DiffExp`, `SeaSyde` [73, 10] and the canonical-form codes [97, 86, 9] are the public references for these steps.

The program `amflow-cpp` (repository `chang18/amflow-cpp` of the “AMFlow.cpp contributors”, MIT license) is a public C++ reimplementation of the auxiliary-mass-flow method of Liu and Ma [91], to whom its notice assigns the algorithmic credit. We adopted this port after running its test suite and comparing its output with closed-form values of known integrals. On our side the algorithm was not reimplemented; our code is driver and basis-builder glue.³ Our driver writes the family specification, supplies a complete quadratic propagator basis, and parses the ε -expansion of the output. Three of our patches affect the values quoted in this paper. At precision goals of 100 digits and above, the command-line output was truncated at 30 digits, and we removed the truncation. The boundary matching at large auxiliary mass could fail when several resonant orders occur in one sector, and we restored the behavior of the original Mathematica implementation. A request whose ε -order was set too low could return a zero finite part produced by the truncation alone, and such a request now stops with an error.

For integer relations we use the PSLQ routine of `mpmath` [118, 55] and, at higher precision, `fpLLL` [116], the `linddep` function of PARI/GP [119] and the LLL reduction in Nemo [89, 58], all as released. PARI/GP is also used for the Weierstrass models, conductors and L -function values of the elliptic curves that appear as maximal cuts, and Oscar and Hecke [42] for Gröbner bases and number-field arithmetic. Modular forms and L -functions are referred to by their labels in the LMFDB [117].

7.2 Tools written for this work

We wrote two programs for steps at which no public code was available. The first, `numkin`, completes an IBP reduction whose multivariate finite-field reconstruction stalls on a few high-degree entries after the system of identities has been generated. We substitute a rational value for one kinematic invariant directly in the generated system files and run one inexpensive reduction per value with Kira and FireFly [77, 75] (with d kept symbolic). The same substitution applies to an auxiliary-mass reduction at fixed rational $\varepsilon = P/Q$, where

$$d \longrightarrow \frac{4Q - 2P}{Q} \quad (7.1)$$

³Our working copies of `amflow-cpp` and of Kira are public as `mattschwartz-sketch/amflow-cpp` and `mattschwartz-sketch/kira`. With our patches the `amflow-cpp` test suite passes in full (572 of 572).

Step	Code (version)	Language as run here	Role in this work	Ref.
IBP reduction	Kira 3.1 with FireFly 2.0.3 and Fermat 7.9a	C++, as released	reduction at numerical kinematic points, differential-equation matrices, input to AMFlow	[94, 77, 76, 75, 90]
Singular locus	PLD.jl / Landau.jl	Julia, as released	principal Landau determinant of the graph	[57, 98]
	SOFIA $\rightarrow$ SOFIA.jl	Mathematica $\rightarrow$ Julia (Nemo, FLINT)	loop-by-loop Baikov singular locus, candidate letters including square roots	[38, 58]
Canonical form	CANONICA, Lee's algorithm, Magnus exponential	Mathematica, not run	baselines for the normalization by leading singularities and the ε -factorization check	[71, 97, 86, 9]
Ansatz, constraints, fits	LandauBootstrap	Mathematica $\rightarrow$ Julia (Nemo)	alphabet layer, weight-graded words, constraints over $\mathbb{Q}$, coefficient fits	[68, 12]
Numerical values	AMFlow algorithm via amflow-cpp (FLINT 3.5.0)	Mathematica $\rightarrow$ C++, public port, adopted and patched	values of the canonical basis at the requested precision goal	[91]
	pySecDec	Python and C++, as released	sector-decomposition cross-check of the closed-form box integral used to test amflow-cpp, sharing no code with Kira or AMFlow	[20]
	DiffExp, SeaSyde	Mathematica, not run	public references for series transport of the differential equation	[73, 10]
Integer relations	mpmath PSLQ, fpLLL, PARI/GP linden, Nemo LLL	Python, C++, C, Julia, as released	identification of fitted coefficients and of boundary constants	[118, 55, 116, 119, 89]
Exact arithmetic, curves	PARI/GP, Oscar, Nemo, Hecke, Arb	C and Julia, as released	Weierstrass models, conductors, Kodaira fibers, L -values, Gröbner bases, ball arithmetic	[119, 42, 58, 74]
Elliptic and modular numerics	Eichler.jl	Julia on Arb, this work	periods, quasi-periods, Eichler integrals and L -values as enclosures, q -series transport	[74, 19]
Reference data	LMFDB	online	labels of modular forms and L -functions	[117]

Table 6: Codes used at each step of the method, with the language in which we run each one and its role in this work. No code on which a result depends requires a Mathematica kernel. Arrows mark ports out of the native language, and “not run” marks packages that we cite as the public references for steps we carried out with our own programs. Version numbers, where given, are those of the runs reported here.

leaves the auxiliary mass as the only symbolic variable. We then reconstruct the dependence on the substituted variable exactly, entry by entry, by the Thiele continued fraction of eq. (2.17), and accept an entry only if it reproduces further samples that did not enter the interpolation. The procedure is the univariate rational reconstruction that FireFly performs internally [76, 75], here applied from outside to the system that has already been generated, which avoids repeating the generation step for each value. The reduction of section 5.1 was obtained in this way, with the difference that d was sampled as well and reconstructed in a second step. The same procedure should help whenever a reconstruction stalls on a small subset of entries.

Our first application of `numkin` was to the differential equation of the three-loop crossed light-by-light box of ref. [108] restricted to the curve $1/s + 1/t = -3/8$, where one block was missing 442 entries and one further row another 17. We reduced the system at 87 rational values of s on the curve, interpolated from 85 of them and reserved $s = -77/10$ and $s = -53/9$, at which all 459 reconstructed entries agree exactly with the direct reduction. Boundary data transported with the completed block agree with an independent `AMFlow` evaluation to 49 digits [108]. This substitution method is not economical when many evaluations at one kinematic point are needed, and block-triangular reduction [65, 66] is then the better route.

The second program, `Eichler.jl`, is a library for elliptic and modular numerics, written in Julia on `Arb` [74]. Each routine returns a ball, i.e. an enclosure whose radius includes a proved bound on the truncation error. The library computes the periods and quasi-periods of a cut curve, the periods of the third kind attached to its punctures, cusp-regularized Eichler integrals, and L -values. Boundary constants are identified by PSLQ against a basis of constants distributed with the library. The basis covers Dirichlet L -values of quadratic characters, the Eisenstein cusp constants, the conductor-four ring that contains Catalan’s constant $L(\chi_{-4}, 2)$, and the conductor-15 Chowla–Selberg ring. The variation-of-parameters integral shared by the elliptic transports of section 6.2 is evaluated as a single q -series, at 0.9 ms per point, whereas the nested quadrature it replaces took on the order of fifteen minutes. Roughly one hundred further utilities, among them sampling-grid generators, symmetry finders and drivers for reconstruction and transport, are documented with their tests in the online registry [19].

8 Conclusions and outlook

In this paper we have described a semi-numerical bootstrap for multi-loop Feynman integrals, together with `BootLoops`, the LLM harness that operates it. The Landau singular locus of the graph and the geometry of its maximal cuts fix the alphabet and the function class of each master integral. Integrability and first-entry conditions then reduce a graded ansatz on that class to a few hundred rational unknowns. We fix those unknowns by integer-relation fits to high-precision auxiliary-mass-flow values. For master integrals that fail the value-fittability criterion, we instead derive or compute boundary data and transport the ε -factorized differential equation. A result is accepted only when it reproduces independent numerical values not used in its derivation. The LLM classifies each integral, evaluates the fittability criterion, runs the codes and assembles the result.

We calibrated the procedure on the outer-mass double box, recovering the alphabet and symbol of refs. [68, 12] from the graph. In that calculation we imposed integrability by exact linear algebra over $\mathbb{Q}$ and set the number of numerical samples by the measured conditioning of the integer-relation fit. We then applied the procedure to the nonplanar two-loop W -pair family F_C with exact top-quark-mass dependence, of which only the master integrals shared with the planar branches had been computed [69]. From the cut geometry we found that exactly one of its sectors is elliptic. On a one-dimensional kinematic line we obtained all 76 master integrals in explicit form through

order ε^1 , most of them by transporting the differential equation from one boundary point. Each integral reproduces independent values at two further points on the line. Across the polylogarithmic, elliptic and Calabi–Yau families of section 6, the geometry-independent constraints continue to hold. Beyond the polylogarithmic class, however, the remaining unknowns of some master integrals can no longer be fixed by a coefficient fit.

We expect the value-fittability criterion to be useful beyond the integrals computed here. It is evaluated from the leading-singularity prefactor and the modularity of the inhomogeneous kernel, without numerical input, and it excludes a fit when the periods of the maximal cut mix under the connection. An elliptic boundary constant, moreover, can be fixed without solving the master integral. We did this once from the p -adic arithmetic of the Picard–Fuchs operator [81] and once from a Ramanujan identity among Eisenstein series, in both cases reproducing a classical value. The toolkit, documented in the online registry of ref. [19], uses no Mathematica kernel at any step of the calculation, and we extend it as new function classes require new bases and evaluators.

The most immediate open problem is to make the value-fittability criterion a theorem. A proof should follow from identifying its fittable and period-mixing classes with the semisimple and unipotent parts of the period matrix of ref. [50]. For the singular locus, one would like a proof that the principal Landau determinant is complete beyond one loop, or a characterization of the components it misses where the locus degenerates at special kinematics [57, 56]. With recursive Landau analysis and the geometric master count of refs. [30, 70], such a result would make the determination of the number of master integrals, the cut geometry and the alphabet automatic, even for alphabets with higher roots [21].

For Calabi–Yau geometries the periods and ε -factorized bases have been obtained [17, 49, 93], but we know of no basis of the corresponding iterated integrals whose evaluation is fast and precise enough for integer-relation fitting. The natural first target for such a fit is the four-loop banana with mass partition $\{4, 1\}$ on its maximal cut (a far smaller differential system than that of the full family). At genus two, both a function space complete enough for an ansatz [14] and Siegel modular kernels for the ε -form [52] remain to be built. A genus-drop test [95, 125] should precede either construction, since for many apparently genus-two integrals the periods of the maximal cut reduce to those of elliptic curves. Finally, both the fit and the transport presuppose a Fuchsian differential system and fail at irregular singular points. It would be interesting to see how many of the integrals obtained here only by transport can eventually be fixed from their singularities and a modest number of numerical values, with no differential equation solved along the way.

A Details of the F_C computation

We collect here, for the family F_C of section 5, the sector labels, the alphabet and the grid of fixed- ε boundary nodes. In addition, Table 7 gives the sectors and homogeneous solutions of the seven diagonal blocks of eq. (5.5) that are not in dlog form on the line $(x, z) = (17/2, 5/3)$.

A.1 Sector labels and alphabet

We number the sectors as in Kira [77]. The sector $[\theta_1, \dots, \theta_9]$, in which the propagator D_i of eq. (5.1) is present when $\theta_i = 1$, has the number

$$S = \sum_{i=1}^9 \theta_i 2^{i-1}, \quad (\text{A.1})$$

so that the top sector $[1, 1, 1, 0, 0, 1, 1, 1, 1]$ is $S = 487$ and the sector $[1, 0, 0, 0, 0, 1, 1, 1, 1]$ with the elliptic maximal cut is $S = 481$.

The alphabet of F_C , obtained from the Landau analysis of section 3.1 [68, 57, 35], has 36 letters in the variables (x, y, z) of eq. (5.2), each defined up to a constant factor. Its 32 rational letters are the polynomials

$$\begin{aligned} W_1 &= x, & W_2 &= y, & W_3 &= z, & W_4 &= x + 4, & W_5 &= y + 1, & W_6 &= y - z, \\ W_7 &= z + 1, & W_8 &= x - 4z, & W_9 &= x + y - z, & W_{10} &= y - z - 1, \\ W_{11} &= xz + x + 1, & W_{12} &= xz + x + y - z, & W_{13} &= x - z^2 - 2z - 1, \\ W_{14} &= x + yz + y - z^2 - z, & W_{15} &= xy + y^2 - 2yz + z^2, \\ W_{16} &= xy^2 + 2xy + x + 4y^2 - 8yz + 4z^2, \\ W_{17} &= x + y - 2z, & W_{18} &= x - yz - y, & W_{19} &= x + 2y - 2z, \\ W_{20} &= x + y - 2z - 1, & W_{21} &= 4xy - 4y^2 + 4y - z, \\ W_{22} &= x^2 + 4x - 4z^2 - 8z - 4, \\ W_{23} &= xy + x + y^2 - 2yz + z^2, \\ W_{24} &= xz + 2x + yz + y - 2z^2 - 2z, \\ W_{25} &= 4x^2 + 4x - 4y^2 + 4y - z^2 - 2z, \\ W_{26} &= 8xy + 32x + 4y^2 + 20y - 15z^2 - 38z, \\ W_{27} &= 16x^2 + 8x - 4y^2 + 16yz - 4y - 19z^2 + 4z, \\ W_{28} &= 4y^2z - 16yz^2 - 4yz - 8y + 15z^3 + 4z^2 + 12z, \\ W_{29} &= x + y - z + 1, & W_{30} &= x - 4z - 4, \\ W_{31} &= x^3 + 2x^2y - 4x^2z + 2x^2 + xy^2 - 4xyz \\ &\quad + 6xy + 4xz^2 - 4xz + x + 4y^2 - 8yz + 4z^2, \\ W_{32} &= xz + yz + y - z^2 - z. \end{aligned} \quad (\text{A.2})$$

The remaining four letters $r_1, \dots, r_4$ are of square-root type, with

$$\begin{aligned} r_1^2 &= x(x + 4), & r_2^2 &= x(x - 4z), \\ r_3^2 &= x(xy^2 + 2xy + x + 4y^2 - 8yz + 4z^2), \\ r_4^2 &= -xy(xy + 4y - 4z - 4). \end{aligned} \quad (\text{A.3})$$

The letters W_1 to W_{16} and the four square roots already occur in the alphabets of the planar branches $\mathcal{T}_1$ and $\mathcal{T}_2$ of ref. [69], whereas W_{17} to W_{28} do not. The last four, W_{29} to W_{32} , were found only by the Euler-characteristic test [35], among them the pseudo-threshold $W_{29} = x + y - z + 1$. On the line $(x, z) = (17/2, 5/3)$ the letters free of y , including r_1 and r_2 , are constant. The polylogarithmic expressions of section 5.3 then involve 15 letters in y .

A.2 Non-logarithmic blocks and boundary nodes

We transport the boundary vector from $y_0 = 1/3$ at fixed $\varepsilon = 1/q$ (section 5.4), solving the differential equation (5.3), whose coefficients are then rational functions of y alone, once for each of the 24 primes

$$q \in \{1013, 1117, 1237, 1367, 1511, 1663, 1847, 2039, 2243, 2503, 2741, 3037, 3347, 3697, 4091, 4513, 4987, 5519, 6089, 6733, 7433, 8219, 9091, 10037\}, \quad (\text{A.4})$$

in roughly geometric progression. At a validation point y_b we recover the Laurent coefficients $f_{i,K}(y_b)$ of the 79 components from the transported values through the square linear system

$$f_i(y_b) \Big|_{\varepsilon=1/q} = \sum_{K=-4}^{15} f_{i,K}(y_b) q^{-K}, \quad q \in \mathcal{Q}_{\text{fit}}, \quad i = 1, \dots, 79, \quad (\text{A.5})$$

where $\mathcal{Q}_{\text{fit}}$ contains 20 of the 24 primes and the remainder of order ε^{16} is neglected. We solve eq. (A.5) at 420-digit working precision. The four primes outside $\mathcal{Q}_{\text{fit}}$, interleaved with the others along the grid, are the verification nodes of Table 2. At these nodes the mismatch between the two sides of eq. (A.5) measures how well the truncated expansion reproduces the transported values, whereas the accuracy of the individual coefficients $f_{i,K}$ is estimated by the out-and-back transport test of section 5.4. The choice of grid and of truncation order is a compromise between the size of the neglected remainder and the conditioning of the matrix (q^{-K}). We select both using that out-and-back test. Finally, for twenty of the boundary constants with analytic expressions, the recomputation from the vacuum values of eq. (5.10) agrees with an independent auxiliary-mass-flow evaluation at fixed d [91] to 93–94 digits, the precision of that evaluation.

S	propagators	block type	homogeneous solutions on the line
417	D_1, D_6, D_8, D_9	radical	closed form, orders ε^{-2} to ε^2
421	D_1, D_3, D_6, D_8, D_9	radical	closed form, orders through ε^1
449	D_1, D_7, D_8, D_9	radical	closed form, orders ε^{-2} to ε^2
453	D_1, D_3, D_7, D_8, D_9	radical	closed form, orders through ε^1
482	D_2, D_6, D_7, D_8, D_9	radical	closed form, orders ε^{-2} to ε^2
483	$D_1, D_2, D_6, D_7, D_8, D_9$	radical	three rational solutions and a pair with nested square roots over four further letters in y
481	D_1, D_6, D_7, D_8, D_9	elliptic, eq. (5.4)	rank six: two integrals of the cut over Lefschetz thimbles, two rational solutions, two completion columns

Table 7: The seven sectors of the family F_C whose diagonal blocks in the connection $A(d, y)$ of eq. (5.3) are not in d log form on the line of section 5, and how each block is solved (section 5.3). S is the sector number of eq. (A.1) and the propagators are those of eq. (5.1). The block type follows eq. (5.5), where “radical” means that the homogeneous solutions involve square roots of polynomials in y . The last column gives the homogeneous solutions on which the order-by-order solution (5.7) is built, with the Laurent orders covered for the five blocks solved in closed form. All seven sectors contain both top-quark propagators D_8 and D_9 . The top sector $S = 487$ is in d log form after the change of one basis integral described in section 5.3 and does not appear.

B Numerical details

For the fits of section 3.3 we sample a basis of uniform transcendental weight at each order in ε . We describe its construction, the validation conditions and the precision of the evaluator scripts.

B.1 Canonical basis and validation

Equation (3.1) fixes only the ε^0 coefficient of g_i ; at higher orders we subtract dimension-shifted subsector integrals,

$$\tilde{g}_i = g_i - \sum_j c_{ij} \mathbf{D}^- I_j, \quad \mathbf{D}^- I_j = \sum_k R_{jk}(\varepsilon; s, t) I_k^{\text{raw}}, \quad (\text{B.1})$$

where j labels subsector masters, $\mathbf{D}^-$ lowers the dimension by two, and the c_{ij} are fixed by requiring that $\tilde{g}_i$ not mix into its subsectors. The second relation is Tarasov’s expression of a $(D-2)$ -dimensional integral through D -dimensional integrals with raised propagator powers [123], which we reduce with the same IBP system [77]. For each master that needs the subtraction we check that w_0 and w_1 vanish in the expansion $\tilde{g}_i = \sum_k \varepsilon^k w_k$. Without the subtraction a rational w_1 of order one remains, which a uniform-weight integral cannot have. On sectors with coupled differential equations no finite set of dotted integrals removes the mixing, and we complete the basis by ε -factorization of the connection [86, 9] where it succeeds.

All fits reported here satisfy four validation conditions.

1. Fit and validation points are disjoint, and where possible lie in separate kinematic regions, so the fit must extrapolate.
2. Equation (3.6) holds at every validation point.
3. Each reference value is computed with AMFlow [91] at two precision goals on different ε -grids, and only the digits common to both runs are kept.

4. The fit is stable under deletion of any fit point, and no numerical solution of the differential equation enters the prefactor or basis (an elliptic nome, for example, is computed in closed form from complete elliptic integrals).

B.2 Evaluators

Each result is accompanied by a standalone Python script, built on `mpmath` [118], that evaluates it anywhere in its kinematic domain to a requested precision. The script performs no IBP reduction or auxiliary-mass-flow computation, and contains the reference values only for comparison. Its series are truncated at orders that scale with the requested precision, and for several such precisions we checked the output against these reference values. Where a script instead solves a proved remainder bound for the truncation order at run time, the requested precision is attained by construction. Hence the digit counts quoted in this paper are measured agreements at independent points, not certified enclosures, except where an evaluator is stated to carry a proved error bound. For the W -pair family of section 5, part of the boundary vector is defined only through the η transport of section 5.4, which is not distributed with the script. The script therefore starts from stored boundary values, whose accuracy caps its precision. Scripts and data are available on each integral's page of the BootLoops site (bootloops.ai/diagrams/) [19].

References

- [1] S. Abreu, F. Febres Cordero, H. Ita, B. Page, and V. Sotnikov. Leading-color two-loop qcd corrections for three-jet production at hadron colliders. 2021. arXiv:2102.13609, JHEP **07** (2021) 095; see also arXiv:1809.06240.
- [2] Samuel Abreu, Ruth Britto, and Claude Duhr. The SAGEX review on scattering amplitudes, Chapter 3: Mathematical structures in Feynman integrals. 2022. arXiv:2203.13014, J. Phys. A **55** (2022) 443004.
- [3] Samuel Abreu, Dmitry Chicherin, Vasily Sotnikov, and Simone Zoia. Two-loop five-point two-mass planar integrals and double lagrangian insertions in a wilson loop. 2024. arXiv:2408.05201.
- [4] Samuel Abreu, Harald Ita, Ben Page, and Wladimir Tschernow. Two-loop hexa-box integrals for non-planar five-point one-mass processes. *JHEP*, 03:182, 2022.
- [5] Luise Adams, Christian Bogner, Armin Schweitzer, and Stefan Weinzierl. The kite integral to all orders in terms of elliptic polylogarithms. 2016. arXiv:1607.01571, MITP/16-069.
- [6] Luise Adams and Stefan Weinzierl. Feynman integrals and iterated integrals of modular forms. 2017. arXiv:1704.08895.
- [7] Luise Adams and Stefan Weinzierl. The ε -form of the differential equations for Feynman integrals in the elliptic case. 2018. arXiv:1802.05020, Phys. Lett. B **781** (2018) 270–278.
- [8] Roger Apéry. Irrationalité de $\zeta(2)$ et $\zeta(3)$. *Astérisque*, 61:11–13, 1979.
- [9] Mario Argeri, Stefano Di Vita, Pierpaolo Mastrolia, Edoardo Mirabella, Johannes Schlenk, Ulrich Schubert, and Lorenzo Tancredi. Magnus and Dyson series for master integrals. 2014. arXiv:1401.2979, JHEP **03** (2014) 082.
- [10] Tommaso Armadillo, Roberto Bonciani, Simone Devoto, Narayan Rana, and Alessandro Vicini. Evaluation of Feynman integrals with arbitrary complex masses via series expansions. *Comput. Phys. Commun.*, 282:108545, 2023. arXiv:2205.03345.
- [11] Piotr Bargiela, Hjalte Frellesvig, Robin Marzucca, Roger Morales, Florian Seefeld, Matthias Wilhelm, and Tong-Zhi Yang. The spectrum of Feynman-integral geometries at two loops. 2026. arXiv:2512.13794, JHEP **05** (2026) 057.
- [12] Oscar Barrera, Aurélien Dersy, Rabia Husain, Matthew D. Schwartz, and Xiaoyuan Zhang. Analytic regression of feynman integrals from high-precision numerical sampling. 2025. arXiv:2507.17815, JHEP **01** (2026) 014.
- [13] Christian W. Bauer, Alexander Frink, and Richard Kreckel. Introduction to the GiNaC framework for symbolic computation within the C++ programming language. 2002. arXiv:cs/0004015, J. Symb. Comput. **33** (2002) 1.
- [14] Konstantin Baune, Johannes Broedel, Egor Im, Artyom Lisitsyn, and Federico Zerbini. Schottky–Kronecker forms and hyperelliptic polylogarithms. 2024. arXiv:2406.10051, J. Phys. A **57** (2024) 445202.

- [15] Spencer Bloch, Matt Kerr, and Pierre Vanhove. A Feynman integral via higher normal functions. *Compos. Math.*, 151(12):2329–2375, 2015. arXiv:1406.2664.
- [16] Spencer Bloch and Pierre Vanhove. The elliptic dilogarithm for the sunset graph. 2015. arXiv:1309.5865, *J. Number Theory* **148** (2015) 328–364.
- [17] Kilian Bönisch, Claude Duhr, Fabian Fischbach, Albrecht Klemm, and Christoph Nega. Feynman integrals in dimensional regularization and extensions of Calabi–Yau motives. 2021. arXiv:2108.05310, *JHEP* **09** (2022) 156.
- [18] Kilian Bönisch, Fabian Fischbach, Albrecht Klemm, Christoph Nega, and Reza Safari. Analytic structure of all loop banana integrals. 2020. arXiv:2008.10574, *JHEP* **05** (2021) 066.
- [19] The BootLoops tool registry. Online registry of the BootLoops tools at bootloops.ai, with per-tool documentation, validation tables, and change history, 2026.
- [20] S. Borowka, G. Heinrich, S. Jahn, S. P. Jones, M. Kerner, J. Schlenk, and T. Zirke. pySecDec: a toolbox for the numerical evaluation of multi-scale integrals. 2018. arXiv:1703.09692, *Comput. Phys. Commun.* **222** (2018) 313.
- [21] Jacob L. Bourjaily, Cristian Vergu, and Matt von Hippel. Landau singularities and higher-order polynomial roots. 2022. arXiv:2208.12765, *Phys. Rev. D* **108** (2023) 085021.
- [22] D. J. Broadhurst and D. Kreimer. Association of multiple zeta values with positive knots via Feynman diagrams up to 9 loops. *Phys. Lett. B*, 393:403–412, 1997.
- [23] David Broadhurst. Feynman integrals, L-series and Kloosterman moments. *Commun. Num. Theor. Phys.*, 10:527–569, 2016. arXiv:1604.03057.
- [24] Johannes Broedel, Claude Duhr, Falko Dulat, Brenda Penante, and Lorenzo Tancredi. Elliptic symbol calculus: from elliptic polylogarithms to iterated integrals of Eisenstein series. 2018. arXiv:1803.10256, *JHEP* **08** (2018) 014.
- [25] Johannes Broedel, Claude Duhr, Falko Dulat, Brenda Penante, and Lorenzo Tancredi. Elliptic Feynman integrals and pure functions. 2019. arXiv:1809.10698, *JHEP* **01** (2019) 023.
- [26] Johannes Broedel, Claude Duhr, Falko Dulat, and Lorenzo Tancredi. Elliptic polylogarithms and iterated integrals on elliptic curves. Part I: general formalism. 2017. arXiv:1712.07089, *JHEP* **05** (2018) 093; see also 1712.07095 (part II: sunrise).
- [27] Francis Brown. Mixed Tate motives over $\mathbb{Z}$. *Ann. of Math.*, 175(2):949–976, 2012.
- [28] Francis Brown and Andrey Levin. Multiple elliptic polylogarithms. 2011. arXiv:1110.6917.
- [29] Philip Candelas, Xenia C. de la Ossa, Paul S. Green, and Linda Parkes. A pair of Calabi–Yau manifolds as an exactly soluble superconformal theory. 1991. *Nucl. Phys. B* **359** (1991) 21–74.
- [30] Simon Caron-Huot, Miguel Correia, and Mathieu Giroux. Recursive Landau analysis. 2024. arXiv:2406.05241, *Phys. Rev. Lett.* **135** (2025) 131603.
- [31] Simon Caron-Huot, Lance J. Dixon, Falko Dulat, Matt von Hippel, Andrew J. McLeod, and Georgios Papathanasiou. The Cosmic Galois Group and extended Steinmann relations for planar $\mathcal{N} = 4$ SYM amplitudes. 2019. arXiv:1906.07116, *JHEP* **09** (2019) 061.

- [32] Simon Caron-Huot and Johannes M. Henn. Iterative structure of finite loop integrals. 2014. arXiv:1404.2922, JHEP **06** (2014) 114.
- [33] Ekta Chaubey and Vasily Sotnikov. Elliptic leading singularities and canonical integrands. 2025. arXiv:2504.20897.
- [34] Kuo-Tsai Chen. Iterated path integrals. *Bull. Amer. Math. Soc.*, 83(5):831–879, 1977.
- [35] Vsevolod Chestnov, Giulio Crisanti, and Mathieu Giroux. Landau’s leviathans. 2026. arXiv:2606.29612.
- [36] K. G. Chetyrkin and F. V. Tkachov. Integration by parts: the algorithm to calculate β -functions in 4 loops. 1981. Nucl. Phys. B **192** (1981) 159.
- [37] Gregory V. Chudnovsky. *Contributions to the Theory of Transcendental Numbers*, volume 19 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 1984.
- [38] M. Correia, M. Giroux, and S. Mizera. SOFIA: Singularities of Feynman integrals automatized. 2025. arXiv:2503.16601.
- [39] Joshua Davies, Dominik Grau, Kay Schönwald, Matthias Steinhauser, Daniel Stremmer, and Marco Vitti. Two-loop QCD corrections to ZH and off-shell Z boson pair production in gluon fusion. 2025. arXiv:2509.07072.
- [40] A. I. Davydychev. Some exact results for N -point massive Feynman integrals. *J. Math. Phys.*, 32:1052–1060, 1991.
- [41] Leonardo de la Cruz. Feynman integrals as A-hypergeometric functions. 2019. arXiv:1907.00507, JHEP **12** (2019) 123.
- [42] Wolfram Decker, Christian Eder, Claus Fieker, Max Horn, and Michael Joswig, editors. *The Computer Algebra System OSCAR: Algorithms and Examples*. Algorithms and Computation in Mathematics. Springer, 2025.
- [43] Lance J. Dixon, James M. Drummond, and Johannes M. Henn. Bootstrapping the three-loop hexagon. 2011. arXiv:1108.4461, JHEP **11** (2011) 023.
- [44] James Drummond, Jack Foster, and Ömer Gürdoğan. Cluster adjacency properties of scattering amplitudes in $\mathcal{N} = 4$ supersymmetric Yang–Mills theory. 2018. arXiv:1710.10953, Phys. Rev. Lett. **120** (2018) 161601.
- [45] J. Dürcker, A. Klemm, and J. F. Piribauer. Calabi–Yau period geometry and restricted moduli in Type II compactifications. 2025. arXiv:2503.00272.
- [46] Claude Duhr. Mathematical aspects of scattering amplitudes. 2015. arXiv:1411.7538, TASI 2014 lectures, in: Journeys Through the Precision Frontier (WSP, 2015) 419–476.
- [47] Claude Duhr and Sara Maggio. Feynman integrals, elliptic integrals and two-parameter K3 surfaces. 2025. arXiv:2502.15326.
- [48] Claude Duhr and Sara Maggio. Three-loop banana integrals with three equal masses. 2025. arXiv:2511.19245.

- [49] Claude Duhr, Sara Maggio, Christoph Nega, Benjamin Sauer, Lorenzo Tancredi, and Fabian J. Wagner. Aspects of canonical differential equations for Calabi–Yau geometries and beyond. 2025. arXiv:2503.20655, JHEP **06** (2025) 128.
- [50] Claude Duhr, Sara Maggio, Christoph Nega, Benjamin Sauer, Lorenzo Tancredi, and Fabian J. Wagner. Canonical differential equations beyond polylogs. 2026. arXiv:2602.04371, RadCor 2025 proceedings.
- [51] Claude Duhr, Sara Maggio, Franziska Porkert, Cathrin Semper, and Sven F. Stawinski. Three-loop banana integrals with four unequal masses. 2025. arXiv:2507.23061, JHEP **12** (2025) 034.
- [52] Claude Duhr, Franziska Porkert, and Sven F. Stawinski. Canonical differential equations beyond genus one. 2024. arXiv:2412.02300, JHEP **02** (2025) 014.
- [53] R. J. Eden, P. V. Landshoff, D. I. Olive, and J. C. Polkinghorne. *The Analytic S-Matrix*. Cambridge University Press, 1966.
- [54] Martin Eichler. Eine Verallgemeinerung der Abelschen Integrale. 1957. Math. Z. **67** (1957) 267–298.
- [55] Helaman Ferguson, David Bailey, and Steve Arno. Analysis of PSLQ, an integer relation finding algorithm. 1999. Math. Comp. **68** (1999) 351.
- [56] Claudia Fevola, Sebastian Mizera, and Simon Telen. Landau singularities revisited: Computational algebraic geometry for feynman integrals. 2023. arXiv:2311.14669, Phys. Rev. Lett. **132** (2024) 101601; see also the companion arXiv:2311.16219.
- [57] Claudia Fevola, Sebastian Mizera, and Simon Telen. Principal Landau determinants. 2023. arXiv:2311.16219, Comput. Phys. Commun. **303** (2024) 109278; see also 2311.14669.
- [58] Claus Fieker, William Hart, Tommy Hofmann, and Fredrik Johansson. Nemo/Hecke: computer algebra and number theory packages for the Julia programming language. 2017. arXiv:1705.06134, Proc. ISSAC 2017, 157–164.
- [59] Felix Forner, Cesare Carlo Mella, Christoph Nega, Lorenzo Tancredi, and Fabian J. Wagner. Integrand analysis, leading singularities and canonical bases beyond polylogarithms. 2026. arXiv:2604.25270.
- [60] Hjalte Frellesvig. On epsilon factorized differential equations for elliptic feynman integrals. 2021. arXiv:2110.07968.
- [61] I. M. Gel’fand, M. M. Kapranov, and A. V. Zelevinsky. Generalized Euler integrals and A -hypergeometric functions. *Adv. Math.*, 84:255–271, 1990.
- [62] Oleksandr Gituliar and Vitaly Magerya. Fuchsia: a tool for reducing differential equations for Feynman master integrals to epsilon form. 2017. arXiv:1701.04269, Comput. Phys. Commun. **219** (2017) 329.
- [63] Alexander B. Goncharov. Multiple polylogarithms and mixed Tate motives. 2001. arXiv:math/0103059.

- [64] Alexander B. Goncharov, Marcus Spradlin, C. Vergu, and Anastasia Volovich. Classical polylogarithms for amplitudes and Wilson loops. 2010. arXiv:1006.5703, Phys. Rev. Lett. **105** (2010) 151605.
- [65] Xin Guan, Xiao Liu, and Yan-Qing Ma. Complete reduction of integrals in two-loop five-light-parton scattering amplitudes. 2020. arXiv:1912.09294, Chin. Phys. C **44** (2020) 093106.
- [66] Xin Guan, Xiao Liu, Yan-Qing Ma, and Wen-Hao Wu. Blade: a package for block-triangular form improved Feynman integrals decomposition. 2025. arXiv:2405.14621, Comput. Phys. Commun. **310** (2025) 109538.
- [67] Holmfridur S. Hannesdottir, Andrew J. McLeod, Matthew D. Schwartz, and Cristian Vergu. Constraints on sequential discontinuities from the geometry of on-shell spaces. *JHEP*, 07:236, 2023.
- [68] Holmfridur S. Hannesdottir, Andrew J. McLeod, Matthew D. Schwartz, and Cristian Vergu. Applications of the Landau bootstrap. 2024. arXiv:2410.02424, Phys. Rev. D **111** (2025) 085003.
- [69] Wen-Jie He, Ren-You Zhang, Liang Han, Yi Jiang, Zhe Li, Xiao-Feng Wang, Shu-Xiang Li, Pan-Feng Li, and Qing-hai Wang. Two-loop planar master integrals for nnlo qcd corrections to w-pair production in quark-antiquark annihilation. *JHEP*, 12:136, 2024. JHEP 12 (2024) 136, arXiv:2409.08879.
- [70] Martin Helmer and Felix Tellander. Geometric singularities of Feynman integrals. 2025. arXiv:2506.05042.
- [71] Johannes M. Henn. Multiloop integrals in dimensional regularization made simple. 2013. arXiv:1304.1806, Phys. Rev. Lett. **110** (2013) 251601.
- [72] Johannes M. Henn. Lectures on differential equations for Feynman integrals. *J. Phys. A*, 48:153001, 2015.
- [73] Martijn Hidding. DiffExp, a Mathematica package for computing Feynman integrals in terms of one-dimensional series expansions. *Comput. Phys. Commun.*, 269:108125, 2021. arXiv:2006.05510.
- [74] Fredrik Johansson. Arb: efficient arbitrary-precision midpoint-radius interval arithmetic. *IEEE Transactions on Computers*, 66:1281–1292, 2017.
- [75] Jonas Klappert, Sven Yannick Klein, and Fabian Lange. Interpolation of dense and sparse rational functions and other improvements in FireFly. 2021. arXiv:2004.01463, Comput. Phys. Commun. **264** (2021) 107968.
- [76] Jonas Klappert and Fabian Lange. Reconstructing rational functions with FireFly. 2020. arXiv:1904.00009, Comput. Phys. Commun. **247** (2020) 106951.
- [77] Jonas Klappert, Fabian Lange, Philipp Maierhöfer, and Johann Usovitsch. Integral reduction with Kira 2.0 and finite-field methods. 2021. arXiv:2008.06494, Comput. Phys. Commun. **266** (2021) 108024.
- [78] René Pascal Klausen. Hypergeometric series representations of Feynman integrals by GKZ hypergeometric systems. 2019. arXiv:1910.08651, JHEP **04** (2020) 121.

- [79] René Pascal Klausen. Kinematic singularities of Feynman integrals and principal A-determinants. 2021. arXiv:2109.07584, JHEP **02** (2022) 004.
- [80] Albrecht Klemm, Christoph Nega, and Reza Safari. The l -loop banana amplitude from GKZ systems and relative Calabi–Yau periods. 2019. arXiv:1912.06201, JHEP **04** (2020) 088.
- [81] Pyry Kuusela, Michael Lathwood, Miroslava Mosso Rojas, and Michael Stepniczka. Solutions of Calabi-Yau differential operators as truncated p-adic series and efficient computation of zeta functions. 2026. arXiv:2604.01191.
- [82] L. D. Landau. On analytic properties of vertex parts in quantum field theory. 1959. Nucl. Phys. **13** (1959) 181.
- [83] S. Laporta. High-precision calculation of multiloop Feynman integrals by difference equations. 2000. arXiv:hep-ph/0102033, Int. J. Mod. Phys. A **15** (2000) 5087.
- [84] S. Laporta. High-precision calculation of the 4-loop contribution to the electron $g-2$ in QED. 2017. arXiv:1704.06996; see also 1805.00227.
- [85] Roman N. Lee. LiteRed 1.4: a powerful tool for reduction of multiloop integrals. 2014. arXiv:1310.1145, J. Phys. Conf. Ser. **523** (2014) 012059.
- [86] Roman N. Lee. Reducing differential equations for multiloop master integrals. *JHEP*, 04:108, 2015. arXiv:1411.0911.
- [87] Roman N. Lee. Libra: a package for transformation of differential systems for multiloop integrals. 2021. arXiv:2012.00279, Comput. Phys. Commun. **267** (2021) 108058.
- [88] Roman N. Lee and Andrei A. Pomeransky. Critical points and number of master integrals. 2013. arXiv:1308.6676, JHEP **11** (2013) 165.
- [89] A. K. Lenstra, H. W. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. 1982. Math. Ann. **261** (1982) 515.
- [90] Robert H. Lewis. Fermat: a computer algebra system for polynomial and matrix computation. <https://home.bway.net/lewis/>, 2026.
- [91] Xiao Liu and Yan-Qing Ma. AMFlow: a Mathematica package for Feynman integrals computation via auxiliary mass flow. 2022. arXiv:2201.11669, Comput. Phys. Commun. **283** (2023) 108565, doi:10.1016/j.cpc.2022.108565; package amflow; amflow-cpp github.com/chang18/amflow-cpp.
- [92] Ron Livné. Motivic orthogonal two-dimensional representations of $\mathrm{Gal}(\bar{\mathbb{Q}}/\mathbb{Q})$. *Israel J. Math.*, 92:149–156, 1995.
- [93] Sara Maggio and Yoann Soehnle. On canonical differential equations for Calabi–Yau multi-scale Feynman integrals. 2025. arXiv:2504.17757, JHEP **10** (2025) 202.
- [94] Philipp Maierhöfer, Johann Usovitsch, and Peter Uwer. Kira—a Feynman integral reduction program. 2018. arXiv:1705.05610, Comput. Phys. Commun. **230** (2018) 99.
- [95] Robin Marzucca, Andrew J. McLeod, Ben Page, Sebastian Pögel, and Stefan Weinzierl. Genus drop in hyperelliptic Feynman integrals. 2023. arXiv:2307.11497, Phys. Rev. D **109** (2024) L031901.

- [96] Eleanor McSpirt and Larry Rolin. Quantum modular forms and resurgence. 2025. arXiv:2505.00799.
- [97] Christoph Meyer. Algorithmic transformation of multi-loop master integrals to a canonical basis with CANONICA. 2018. arXiv:1705.06252, Comput. Phys. Commun. **222** (2018) 295.
- [98] Sebastian Mizera and Simon Telen. Landau discriminants. 2021. arXiv:2109.08036, JHEP **08** (2022) 200.
- [99] R. Morales, A. Spiering, M. Wilhelm, Q. Yang, and C. Zhang. Bootstrapping elliptic feynman integrals using Schubert analysis. 2022. arXiv:2212.09762, Phys. Rev. Lett. **131** (2023) 041601.
- [100] Roger Morales. Special functions and geometries in scattering amplitudes: From particle physics to gravity. 2025. arXiv:2506.11911 (thesis-monograph).
- [101] Yu. V. Nesterenko. Modular functions and transcendence questions. *Sb. Math.*, 187(9):1319–1348, 1996. Russian original: Mat. Sb. **187**:9 (1996) 65–96.
- [102] Sebastian Pögel, Xing Wang, and Stefan Weinzierl. Taming Calabi–Yau Feynman integrals: the four-loop equal-mass banana integral. 2022. arXiv:2211.04292, Phys. Rev. Lett. **130** (2023) 101601.
- [103] Sebastian Pögel, Xing Wang, and Stefan Weinzierl. The three-loop equal-mass banana integral in ε -factorised form with meromorphic modular forms. 2022. arXiv:2207.12893, JHEP **09** (2022) 062.
- [104] Mario Prausa. epsilon: a tool to find a canonical basis of master integrals. 2017. arXiv:1701.00725, Comput. Phys. Commun. **219** (2017) 361.
- [105] Sebastian Pögel, Toni Teschke, Xing Wang, and Stefan Weinzierl. The unequal-mass three-loop banana integral. 2025. arXiv:2507.23594.
- [106] Niklas Schwanemann and Stefan Weinzierl. Electroweak double-box integrals for Möller scattering. *SciPost Phys.*, 18:172, 2025.
- [107] Matthew D. Schwartz. Bootstrapping elliptic and Calabi–Yau Feynman integrals. Companion paper (the threshold-geometry results of this program); in preparation, 2026.
- [108] Matthew D. Schwartz. Frontier Feynman integrals from a hybrid bootstrap. Companion paper (the computation portfolio of this program); in preparation, 2026.
- [109] Matthew D. Schwartz. The orthorhombic Watson integral is a rank-five K3 period. Companion note; in preparation, 2026.
- [110] Atle Selberg and Sarvadaman Chowla. On Epstein’s zeta-function. *J. Reine Angew. Math.*, 227:86–110, 1967.
- [111] Goro Shimura. Sur les intégrales attachées aux formes automorphes. 1959. J. Math. Soc. Japan **11** (1959) 291–311.
- [112] A. V. Smirnov and F. S. Chukharev. FIRE6: Feynman Integral REduction with modular arithmetic. 2020. arXiv:1901.07808, Comput. Phys. Commun. **247** (2020) 106877.

- [113] A. V. Smirnov, N. D. Shapurov, and L. I. Vysotsky. FIESTA5: numerical high-performance Feynman integral evaluation. 2022. arXiv:2110.11660, Comput. Phys. Commun. **277** (2022) 108386.
- [114] Othmar Steinmann. Über den Zusammenhang zwischen den Wightmanfunktionen und den retardierten Kommutatoren. *Helv. Phys. Acta*, 33:257–303, 1960.
- [115] Matija Tapušković. The cosmic Galois group, the sunrise Feynman integral, and the relative completion of $\gamma_1(6)$. 2023. arXiv:2303.17534.
- [116] The FPLLL development team. fplll, a lattice reduction library. <https://github.com/fplll/fplll>, 2026.
- [117] The LMFDB Collaboration. The L-functions and modular forms database. <https://www.lmfdb.org>, 2026.
- [118] The mpmath development team. mpmath: a Python library for arbitrary-precision floating-point arithmetic. <https://mpmath.org>, 2026.
- [119] The PARI Group. PARI/GP. Univ. Bordeaux, <https://pari.math.u-bordeaux.fr/>, 2026.
- [120] F. V. Tkachov. A theorem on analytical calculability of 4-loop renormalization group functions. 1981. Phys. Lett. B **100** (1981) 65.
- [121] Duco van Straten. Calabi–yau operators. 2017. arXiv:1704.00164.
- [122] Jens Vollinga and Stefan Weinzierl. Numerical evaluation of multiple polylogarithms. 2005. arXiv:hep-ph/0410259, Comput. Phys. Commun. **167** (2005) 177.
- [123] Stefan Weinzierl. *Feynman Integrals: A Comprehensive Treatment for Students and Researchers*. UNITEXT for Physics. Springer, 2022. arXiv:2201.03593, Springer (2022), DOI 10.1007/978-3-030-99558-4.
- [124] Matthias Wilhelm and Chi Zhang. Symboly for elliptic multiple polylogarithms and the symbol prime. 2022. arXiv:2206.08378, JHEP **01** (2023) 089.
- [125] Feiyu Yang, Jianyu Gong, and Yang Zhang. Genus drop involving non-hyperelliptic curves in Feynman integrals. 2026. arXiv:2605.07729.
- [126] Li Lin Yang and Yiyang Zhang. From dlog to dE: Canonical elliptic integrands and modular symbol letters with pure eMPLs. 2025. arXiv:2512.19370.
- [127] Don Zagier. Values of zeta functions and their applications. In *First European Congress of Mathematics (Paris, 1992), Vol. II*, volume 120 of *Progress in Mathematics*, pages 497–512. Birkhäuser, 1994.
- [128] Don Zagier. Elliptic modular forms and their applications. In Kristian Ranestad, editor, *The 1-2-3 of Modular Forms*, Universitext, pages 1–103. Springer, 2008.