


DRCC-Width-Zeta (DWZ)

*A Width-Controlled Euler-Maclaurin Reconstruction Method
for the Riemann Zeta Function in the Critical Strip*

Motivated by the DRCC/CPR Reconstruction-Width Framework

$t = 0$



The Zeta Spiral: $\zeta(0.5 + it)$ for $t \in [0, 60]$ on the Critical Line

Color gradient burgundy $\rightarrow$ gold with increasing t

Reza Hesamiy

DWZ-V2.1 • Corrected Revision, 04 August 2026

DRCC-WIDTH-ZETA Method (DWZ)

A Width-Controlled Euler-Maclaurin Reconstruction Method

for the Riemann Zeta Function in the Critical Strip

Motivated by the DRCC/CPR Reconstruction-Width Framework

Reza Hesamiy

DWZ-V2.1 – Corrected Revision, August 4, 2026

Abstract

This work is a reproducible case study in operationalizing an abstract reconstruction-width framework – DRCC, *Dimensional Reduction via Controlled Combinatorics* [7] – on a single, well-understood numerical problem: computing the Riemann zeta function $\zeta(s)$ in the critical strip $0 < \text{Re}(s) < 1$. The question we ask is not “how do we compute $\zeta(s)$ efficiently” (the Riemann-Siegel formula already answers that), but “what does it concretely take to turn an abstract stability condition, $0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq d_{\text{frag}}(X) < \infty$, into a checkable number on a specific, nontrivial example, end to end, with every approximation and limitation made explicit.” We show that the naive partial sum $\sum_{n=1}^W n^{-s}$ diverges in the critical strip as the width W grows, and that a first-order Euler-Maclaurin correction – a classical, eighteenth-century technique of analysis, not a contribution of this work – resolves this and yields a genuinely convergent width-zeta function (empirically fitted scaling exponent close to $-3/2$ on the critical line $\text{Re}(s) = \frac{1}{2}$, consistent with the first omitted Euler-Maclaurin term under the stated local simple-zero assumptions, Section 2.8.2; elsewhere in the strip the order depends on $\text{Re}(s)$), which we call the *DWZ method* (DRCC-Width-Zeta method; methodically: *Euler-Maclaurin width reconstruction method*). We deliberately use the simplest, first member of the correction family at which convergence occurs at all, since methodological traceability – not competitive speed – is the goal of this case study.

Throughout this work, “width” W refers exclusively to the numerical truncation depth W_{trunc} of the partial sum; an identification with the structural CPR-Width of Hesamiy [7] is not claimed (Section 2.5). From this we derive a closed-form, empirically fitted estimate $\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon) \approx (C(t)/\varepsilon)^{2/3}$ of the originally abstract DRCC condition – a concrete instantiation of the kind this case study is meant to demonstrate. The method is numerically checked against `mpmath` reference values for 100 nontrivial zeros (absolute deviation 1.5×10^{-6} – 1.3×10^{-5}): 98 are locally refined from known reference starting values and 2 are found by blind scans without reference knowledge, with 29 of the 100 (#13–#41) additionally hardened via the hybrid method, and a resulting pattern hypothesis is tested statistically without finding a significant relationship.

As expected for a minimal, order-1 correction chosen for traceability rather than speed, a non-workload-matched cumulative timing comparison against the established Riemann-Siegel formula (via `mpmath`) on all 100 zeros – descriptive rather than a strict like-for-like algorithm benchmark, since DWZ refines from an already-known starting value while the reference routine performs a full search – confirms a clear efficiency disadvantage: Riemann-Siegel is about 30 times faster and simultaneously more accurate – a result reported without embellishment, and consistent with the case-study framing above rather than in tension with it. Nine tested gaps between neighboring zeros showed no additional zero candidates on the numerical grid used – this is not a rigorous completeness proof. The results constitute a validity demonstration of the operationalization, not a contribution to the Riemann Hypothesis itself, whose statement ($\zeta(\rho) = 0$, $0 < \text{Re}(\rho) < 1 \Rightarrow \text{Re}(\rho) = \frac{1}{2}$) is not assumed here: the work restricts its numerical experiments to already-known zeros located on the critical line, but does not assume that *all* nontrivial zeros lie there. All figures in this document are drawn

directly from the underlying numerical data using `pgfplots` – no external image files are required.

Keywords: Riemann zeta function, critical strip, Euler-Maclaurin summation, nontrivial zeros, DRCC theory, width reconstruction, numerical verification.

Contents

Core Equation Register	3
1 Introduction	3
1.1 Starting Point	3
1.2 DRCC Origin and Scope	4
1.3 Contribution of This Work	5
1.4 DRCC Contribution, Equation Traceability, and Counterfactual Role	5
2 Methods	7
2.1 The Basic Problem of the Naive Width Definition	7
2.2 Variant A: Euler-Maclaurin Correction (the DWZ Method)	9
2.3 Variant B: Eta-Based Alternative	10
2.4 Variant C: Second-Order Euler-Maclaurin Correction	10
2.5 Order Degeneracy and Model Dependence	10
2.6 The DRCC Retraction Pair: What It Does and Does Not Establish	19
2.7 Preliminary RP-Tensor Scoring Framework for Reconstruction Variants	20
2.8 Operationalizing $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$	21
2.8.1 Conditional Numerical Operationalization	23
2.8.2 Analytical Origin of the Observed $W^{-3/2}$ Rate	24
2.8.3 General DRCC Operationalization Principle	26
2.9 Verification Strategies	28
2.10 Detailed Worked Example: Determination of Zero #1	30
3 Results	33
3.1 Comparison of the Repair Variants	33
3.2 $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$	33
3.3 Blind and Hybrid Zero Search	37
3.4 Hardening and Pattern Hypothesis	38
3.5 Numerical Reproduction and Reference Check for 100 Zeros	39
3.6 Gap Verification	42
4 Discussion	42
4.1 Scope and Limitations	42
4.2 Detailed Discussion	44
4.3 Why DRCC, Given That an Analytic Reconstruction Formula Already Converges	44
5 DRCC/CPR-Width Application Assessment of the DWZ Method	48
5.1 Purpose of the Application Assessment	48
5.2 Mapping and Resource Separation, in Brief	48
5.3 Boundary-Condition Assessment	49
5.4 Net Assessment	49
6 Conclusion and Outlook	49
Appendix A: List of Symbols	50

Core Equation Register

Eq.	Equation	Section
(1)	$\zeta(s) = \sum_{n=1}^{\infty} n^{-s}$ (Dirichlet series)	§1.1
(2)	Euler product $\prod_p (1 - p^{-s})^{-1}$	§1.1
(3)	Riemann Hypothesis (RH)	§1.1
(4)	DRCC stability condition $0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq d_{\text{frag}}(X) < \infty$	§1.2
(5)	naive partial sum $\zeta_W(s, W)$	§2.1
(6)	intermediate form $0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq W < \infty$	§2.1
(7)	$\zeta_{\text{DWZ}}(s, W)$ – Variant A / model M_1	§2.2
(8)	$\eta_W(s, W), \zeta_W^\eta(s, W)$ – Variant B	§2.3
(9)	$\zeta_{\text{DWZ}}^{(2)}(s, W)$ – Variant C / model M_2	§2.4
(10)	$W_{\min}(M; s, \varepsilon)$	§2.5
(11)	$810 < 9000$ (M_2 vs. M_1)	§2.5
(12)	$W_{\min}^*(s, \varepsilon)$	§2.5
(13)	RP-Tensor selection score $S_\alpha(\tau)$	§2.7
(14)	$W_{\text{stab}}^{(M_1)}(t, \varepsilon)$, stable-threshold definition	§2.8
(15)	$\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$, empirical estimator (fit basis: $n = 50$ zeros)	§2.8
(16)	convergence hypothesis for Proposition 2	§2.8.1
(30)–(38)	general DRCC operationalization schema and Operational DRCC Record $\mathfrak{D}$	§2.8.3
(43)	$N(T)$, Riemann-von Mangoldt counting function	§2.9
(44)	Newton iteration rule $s_{n+1} = s_n - f(s_n)/f'(s_n)$	§2.10
(45)	$ s^* - \rho_1 = 1.66 \times 10^{-5}$ (worked example, zero #1)	§2.10
(46)	$r_{\text{pb}} = -0.189$, point-biserial correlation (hardening test, $n = 20$)	§3.4

This register lists the load-bearing equations of the construction; it is a curated selection, not an exhaustive index of every numbered equation in the manuscript.

1 Introduction

Framing. Before turning to the zeta function itself, it is worth stating plainly what kind of contribution this paper is and is not. It is not a proposal for a faster or more accurate way to compute $\zeta(s)$ – the Riemann-Siegel formula already does that, and Section 3 confirms this directly rather than disputing it. It is a reproducible case study: it takes one specific abstract framework (DRCC, with its stability condition (4)) and works through, in full and checkable detail, what it takes to instantiate that framework on one concrete, well-understood numerical problem – including every place where the instantiation is partial, where a naive first attempt fails, and where the resulting numbers turn out less favorable than an established alternative. The value of this exercise, if any, lies in the transparency and reproducibility of that process, not in the competitive performance of its output.

1.1 Starting Point

The Riemann zeta function is defined for $\text{Re}(s) > 1$ by

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}, \quad (1)$$

with the Euler product representation

$$\zeta(s) = \prod_{p \text{ prime}} \frac{1}{1 - p^{-s}}, \quad (2)$$

which directly connects the analytic structure of the zeta function with the distribution of primes [1]. The analytic continuation of $\zeta(s)$ to $\mathbb{C} \setminus \{1\}$ has, in the critical strip $0 < \text{Re}(s) < 1$, the nontrivial zeros. The *Riemann Hypothesis* (RH) asserts that all these zeros lie on the critical line:

$$\zeta(\rho) = 0, \quad 0 < \text{Re}(\rho) < 1 \xrightarrow{\text{RH}} \text{Re}(\rho) = \frac{1}{2}. \quad (3)$$

This is an open conjecture, not a proven theorem. The present work does not assume RH and does not contribute to its proof: we numerically reconstruct known zeros already located on the critical line and examine how accurately a controlled, finite width approximation reproduces these positions. Classical references on the analytic theory of the zeta function are Titchmarsh [2] and Edwards [3]; the numerical verification of individual zeros goes back to Turing’s method [4].

1.2 DRCC Origin and Scope

DRCC (*Dimensional Reduction via Controlled Combinatorics*, Hesamiy [7]) replaces an infinite or high-dimensional structure X by a controlled, finite *width* W , capturing the resulting reconstruction loss through a stability condition,

$$0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq d_{\text{frag}}(X) < \infty, \quad (4)$$

not yet numerically instantiated for the zeta case at the outset of this work (without thereby claiming anything about their status within the cited DRCC theory itself). The goal of this work is a conditional numerical instantiation of the form of inequality (4) for the concrete case of the Riemann zeta function – we call the resulting method the *DWZ method* (DRCC-Width-Zeta method; methodically the *Euler-Maclaurin width reconstruction method*, Section 2).

Terminological note: throughout this work, “width” W consistently refers to the numerical truncation depth W_{trunc} , not necessarily the structural CPR-Width ω_{CPR} (*Controlled Partial Reconstruction Width*, Hesamiy [7]); this question is tested empirically in Section 2.5 and made precise formally in Proposition 2 (Section 2.8.1). Table 1 makes this commitment verifiable rather than merely asserted.

Table 1: Traceability: DRCC/CPR terms and their role in this work.

DRCC/CPR (Hesamiy [7])	term	Role in this work	Status
$d_{\text{rec}}^{\text{DRCC}}(X)$, (stability condition)	$d_{\text{frag}}(X)$	Motivates the width idea; inequality (4) as the starting point	Formally operationalized in Proposition 2
CPR-Width tolerance parameter ε_{rec} (Def. 2.7–2.9, min-max over order- ings ω)	ω_{CPR}	Tested in Section 2.5 (Test A/B)	Ordering minimization is degenerate; identification with W_{trunc} is <i>not</i> shown
Width W		Operationalized as the numerical truncation depth W_{trunc} of the partial sum	Independent quantity, newly defined in this work

In short: the title and framework name the theoretical origin. What this work shows is a *numerical case study and formal partial operationalization* of the DRCC starting idea – not a

proof that the abstract CPR-Width and the numerical truncation width are the same quantity. Section 2.5 does not merely assert this point but formally settles it (Lemma 1 and the subsequent instantiations): for the natural reconstruction tasks of the sequential zeta accumulation, the structural CPR-Width is provably constant ($\omega_{\text{CPR}} = O(1)$), while the numerical truncation width W_{trunc} grows without bound as the target accuracy is tightened. DRCC thus supplies the motivation and conceptual frame for the stability condition under study, not a proven structural identity with the numerical method; this work accordingly treats DRCC consistently as context, not as the mathematical foundation of the numerical results.

On methodological novelty. The Euler-Maclaurin summation formula itself is a long-established standard technique of analysis [2, 3] and is not claimed here as a new mathematical invention. The contribution of this work lies in its application – the explicit, minimal (order-1) operationalization of the abstract DRCC stability condition for the zeta function in the critical strip – and in the empirical characterization of the resulting procedure: convergence order, model-dependent truncation depth, reproduction across 100 zeros (with the empirical fit for $C(t)$ remaining based on the first 50 zeros), and a systematic comparison against a higher-order Euler-Maclaurin expansion and the Riemann-Siegel formula (Section 2.5, Test C), which shows that the minimal order-1 variant carries an efficiency disadvantage of two to three orders of magnitude relative to established methods. The present method is thus not intended as a faster alternative to Riemann-Siegel, but as the most directly traceable, minimal route for translating the DRCC condition into a checkable number on a concrete, nontrivial example.

1.3 Contribution of This Work

We trace the development of inequality (4) through four stages: from the abstract placeholder condition, through the identification $W = d_{\text{frag}}(X)$ and an initially still-incomplete intermediate form, to a closed-form expression for $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ derived from numerical data (Section 2). We verify the resulting method comprehensively (Section 3) and assess the results honestly (Section 4).

1.4 DRCC Contribution, Equation Traceability, and Counterfactual Role

The Euler-Maclaurin correction underlying the DWZ method is classical (Section 1.2); it would exist, and would converge, with or without DRCC. This raises a fair question, one a reviewer is entitled to ask directly: what, concretely, did DRCC contribute to this manuscript? This section answers that question for each traceable piece of the construction, rather than asserting the answer in general terms. For every DRCC-motivated equation used in this work, we state: (1) which abstract DRCC idea it concretizes, (2) how it was concretized in the DWZ construction, (3) where its effect becomes mathematically or numerically visible, and (4) what would be missing from this manuscript without that step – understood historically, as a statement about how this particular manuscript was actually developed, not as a claim that no other route to the same numerical results could exist.

Three different kinds of success. Collapsing these into a single “DRCC worked” statement would blur three distinct claims that deserve to be kept separate.

Methodological success. DRCC did not supply the Euler-Maclaurin formula. What it supplied was the demand for a checkable chain from the abstract condition (4) to a reproducible numerical protocol: $d_{\text{frag}}(X) \rightsquigarrow W$, then $d_{\text{rec}}^{\text{DRCC}}(X) \rightsquigarrow W_{\text{stab}}^{(M_1)}(t, \varepsilon)$, then the conditional conclusion $W \geq W_{\text{stab}}^{(M_1)}(t, \varepsilon) \Rightarrow e_k^{(M_1)}(W) < \varepsilon$ (Proposition 2). Turning an abstract reconstruction requirement into a testable numerical statement is the methodological contribution.

Negative, but important, success. The DRCC/CPR framework raised the specific question of whether the numerical truncation width W_{trunc} coincides with the structural CPR-Width ω_{CPR} .

Table 2: DRCC equation register: traceability from the abstract DRCC condition to its concretizations in this manuscript.

DRCC idea	DWZ concretization	Where the effect is visible	Without this step
Abstract stability condition 0 < $d_{\text{rec}}^{\text{DRCC}}(X) \leq d_{\text{frag}}(X) < \infty$	Eq. (4), §1.2; Table 1	Frames the entire manuscript structure around a width/accuracy pair, not just a convergence question	No formal reason to track a fragmentation–reconstruction pair at all; convergence alone would suffice
Fragment identification $d_{\text{frag}}(X) := W$	Eq. (6), §2.1	Naive fragment tested and shown to diverge (Table 3), motivating the correction of §2.2	Naive divergence would read as an isolated numerical fact, not as the failure of a specific, declared fragment
Stable reconstruction depth $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$	Eq. (14), §2.8	Used throughout Section 3 as the operational accuracy threshold; drives the empirical estimator below	No formal accuracy–relative depth concept – only an ad hoc “large enough W ”
Empirical estimator $\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon) \approx (C(t)/\varepsilon)^{2/3}$	Eq. (15), §2.8	Closed-form predictive guide, tested out-of-sample on zeros #51–100 (Section 3)	No closed-form width–prediction formula; every t would require a fresh numerical search
Conditional formal instantiation	Proposition 2, §2.8.1 (p. 22)	Turns the DRCC→DWZ transition into a stated, checkable conditional proof, not an assertion	The DRCC→DWZ link would remain prose and motivation only, as the first reviewer round correctly criticized
State-width lemma $\omega_{\text{rec}} = O(1) \neq W_{\text{trunc}}$	Lemma 1, §2.5 (p. 15)	Separates runtime length from active state width; load-bearing for Section 4 and the Conclusion	The naive identification $W_{\text{trunc}} = \omega_{\text{CPR}}$ would stand unchallenged and unexamined
General operationalization template and Operational DRCC Record $\mathfrak{D}$	Eqs. (30)–(38), §2.8.3 (pp. 25–27)	States, honestly and narrowly, what a second instantiation of this template would require	DWZ would read as a one-off numerical trick with no stated, reusable structure for future instances

The answer, proven in Lemma 1 rather than merely observed, is that it does not: $\omega_{\text{CPR}} = O(1)$ while $W_{\text{trunc}} \rightarrow \infty$ as $\varepsilon \rightarrow 0$. Ruling out a naive identification is not a failure of the framework – it is a genuine separation result,

$$\text{state width} \neq \text{truncation depth} \neq \text{runtime},$$

that this manuscript would very plausibly not have investigated without the DRCC/CPR vocabulary prompting the question in the first place.

Numerical success of the operationalization. The concretization chosen here produces checkable numbers: the naive fragment diverges; M_1 converges; M_2 reduces the required width by roughly a factor of 11 in the tested comparison (Eq. (11), 9000 vs. 810); the observed error exponent is explained analytically by the first omitted Euler-Maclaurin term (Section 2.8.2); 100 known zeros are numerically reproduced with absolute deviation 1.5×10^{-6} – 1.3×10^{-5} , two of them found blindly; and the same comparison honestly shows Riemann-Siegel to be substantially faster. The claim this supports is precise:

DRCC’s role here is operationalization and resource separation,

not

DRCC produced a faster zeta method.

The second statement is not made anywhere in this manuscript; the first one is what Table 2 traces in detail.

What would remain without this DRCC step. Stripped of the DRCC framing, this manuscript’s numerical core would still exist as roughly: *a numerical study of a first-order Euler-Maclaurin approximation for reproducing known zeros of the Riemann zeta function*. That hypothetical paper would still contain equation (7), convergence tests, the zero search, and the Riemann-Siegel comparison. Based on how this manuscript actually developed, it would not contain: the starting condition (4); the $d_{\text{frag}}/d_{\text{rec}}^{\text{DRCC}}$ distinction; the stable reconstruction threshold (14); the DRCC estimator (15); Proposition 2 as a formal operationalization; the separation of W_{trunc} from ω_{CPR} ; the Operational DRCC Record (38); the general template of Section 2.8.3; or the RP-Tensor question of systematic model selection (Section 2.7). This is a statement about the development history and content of this specific manuscript, not a claim that DRCC is the only conceivable route to any of these individual ideas.

Closing statement. The Euler-Maclaurin formula is classical. DRCC’s contribution in this manuscript is the explicit operationalization of an abstract reconstruction condition, the definition of a measurable reconstruction depth, the separation of model choice from structural resource, and a fully reproducible operational record – traced end to end in Table 2 rather than asserted in the abstract. The intended reading is

DRCC question $\rightarrow$ DWZ concretization $\rightarrow$ measured result $\rightarrow$ counterfactual without DRCC,
a complete accounting of DRCC’s role in this manuscript, not a stronger claim standing in place of one.

2 Methods

2.1 The Basic Problem of the Naive Width Definition

An obvious first approach to the width-zeta function is the plain partial sum

$$\zeta_W(s, W) := \sum_{n=1}^W \frac{1}{n^s}, \quad 0 < W < \infty, \quad (5)$$





Figure 1: OPA-DRCC-01: From abstract condition to measurable width – the core traceability chain in one diagram.



W	naive $ \zeta_W(s, W) - \zeta(s) $	DWZ $ \zeta_{\text{DWZ}}(s, W) - \zeta(s) $
10	2.47×10^{-1}	3.85×10^{-2}
100	7.09×10^{-1}	1.18×10^{-3}
1000	2.24×10^0	3.73×10^{-5}
5000	5.00×10^0	3.33×10^{-6}

Figure 2: OPA-DRCC-02: Failure, repair, success, with real values at $s = 0.5 + 14.13i$ taken directly from Table 3 (naive column) and Table 12 (DWZ column) – the naive fragment diverges while the reconstructed function converges, on the same points.

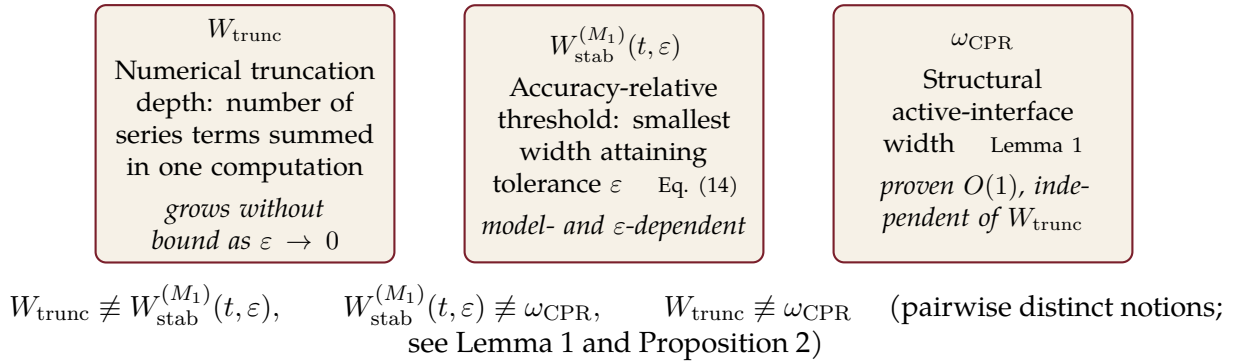


Figure 3: OPA-DRCC-03: Three different widths that this manuscript is careful never to conflate.




Figure 4: OPA-DRCC-04: Evidence ladder – each rung is labeled with its actual epistemic status; the ladder does not get stronger by omitting the labels on its weaker rungs.

with $W := d_{\text{frag}}(X)$ as the first concretization of equation (4) (the justification for this concretization is discussed in Section 2.5 and in the Conclusion; W depends on the chosen model, the accuracy ε , the point s , the search strategy, and the computational precision, and is not an intrinsic property of X itself). This series converges as $W \rightarrow \infty$ only for $\text{Re}(s) > 1$. In the critical strip, $\zeta_W(s, W)$ diverges as W grows (Table 3) instead of approaching $\zeta(s)$ – equation (5) is therefore unsuitable as a reconstruction approach in the critical strip.

Table 3: Divergence of the naive partial sum at $s = 0.5 + 14.13i$.

W	$ \zeta_W(s, W) - \zeta(s) $
10	2.47×10^{-1}
50	5.04×10^{-1}
100	7.09×10^{-1}
500	1.58×10^0
1000	2.24×10^0
5000	5.00×10^0

At this point the inequality

$$0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq W < \infty \quad (6)$$

remained incomplete: $d_{\text{rec}}^{\text{DRCC}}(X)$ was still undetermined.

2.2 Variant A: Euler-Maclaurin Correction (the DWZ Method)

The Euler-Maclaurin summation formula – a classical, eighteenth-century result of analysis and not a contribution of this work (see the explicit disclaimer in Section 1.2, “On methodological novelty”) [2, 3] – provides a correction that makes ζ_W convergent even in the critical strip:

$$\zeta_{\text{DWZ}}(s, W) = \sum_{n=1}^W \frac{1}{n^s} + \frac{W^{1-s}}{s-1} - \frac{W^{-s}}{2} \quad (7)$$

with $\lim_{W \rightarrow \infty} \zeta_{\text{DWZ}}(s, W) = \zeta(s)$ also holding for $0 < \text{Re}(s) < 1$. The asymptotic error induced by the first omitted Euler-Maclaurin correction term is derived in Section 2.8.2.

2.3 Variant B: Eta-Based Alternative

As a comparison method, Variant B uses the Dirichlet eta function, which already converges for $\text{Re}(s) > 0$:

$$\eta_W(s, W) = \sum_{n=1}^W \frac{(-1)^{n-1}}{n^s}, \quad \zeta_W^\eta(s, W) = \frac{\eta_W(s, W)}{1 - 2^{1-s}}. \quad (8)$$

2.4 Variant C: Second-Order Euler-Maclaurin Correction

Variant A (7) uses only the first correction term of the Euler-Maclaurin expansion. The next term of the series (involving the Bernoulli number $B_2 = 1/6$) yields a consistent extension:

$$\zeta_{\text{DWZ}}^{(2)}(s, W) = \sum_{n=1}^W \frac{1}{n^s} + \frac{W^{1-s}}{s-1} - \frac{W^{-s}}{2} + \frac{s}{12} W^{-s-1} \quad (9)$$

We henceforth refer to equation (7) as model M_1 and equation (9) as model M_2 .

Methodological design principle. The purpose of this work is not to identify the numerically most efficient reconstruction model, but to demonstrate that the DRCC stability condition (Section 2) can already be operationalized by the smallest non-trivial Euler-Maclaurin reconstruction. To be explicit about scope: the Euler-Maclaurin summation formula itself is classical eighteenth-century analysis, not a contribution of this work (Section 1.2); what is discussed below concerns solely which correction order this manuscript adopts as its baseline, not the formula’s origin, validity, or generality. “Smallest non-trivial” is not a stylistic preference here: the zero-order case, the naive partial sum of equation (5), diverges throughout the critical strip (Table 3), so M_1 is the first member of the Euler-Maclaurin correction family at which convergence – and hence a non-trivial reconstruction in the DRCC sense – is achieved at all. For this reason, M_1 is adopted as the primary reconstruction model throughout the manuscript. M_2 , and in principle any higher Euler-Maclaurin order, are included as comparative models within this same family, to quantify the improvement in width and accuracy obtainable through additional correction terms (Section 2.5), rather than to replace the methodological baseline established by M_1 . This scope is deliberately restricted to the Euler-Maclaurin family: M_η (Variant B, Section 2.3) is a structurally different construction and is retained in the comparison for a separate reason, not as a further point on this order hierarchy. Consequently, the finding of Section 2.7 that M_2 dominates M_1 on the two RP-Tensor criteria for which data exist does not contradict the choice of M_1 as baseline: that finding quantifies precisely the improvement this design principle anticipates from moving to a higher order, and the choice of M_1 is stated here as a declared methodological decision that sits outside, and precedes, the RP-Tensor ranking – not as its outcome.

2.5 Order Degeneracy and Model Dependence

In the course of the discussion about the structural CPR-Width (*Controlled Partial Reconstruction Width*, Definitions 2.7–2.9 in Hesamiy [7]), two targeted tests were carried out to clarify whether the summation order ω or the reconstruction model M is the actual lever controlling the DWZ width.

Test A – Order invariance. The same 500 terms n^{-s} ($s = 0.5 + 1001.3495i$) were summed in ascending, descending, and mixed order. At 30-digit floating-point precision, the observed ordering effects remained at rounding level, 5.6×10^{-31} , and did not change any tested width threshold – expected, since a finite sum is permutation-invariant in exact arithmetic. For a fixed model M , minimization over orderings ω is therefore degenerate: the summation order is *not* the effective optimization variable.

Test B – Model comparison. For a fixed accuracy target $\varepsilon = 10^{-4}$, the minimally required width

$$W_{\min}(M; s, \varepsilon) := \min\{W \in \mathbb{N} : |\zeta_M(s, W) - \zeta(s)| < \varepsilon\} \quad (10)$$

was defined (with $W_{\min}(M; s, \varepsilon) := \infty$ if this set is empty). It was determined for M_1 and M_2 at five zeros distributed across the tested range using a grid search (step size 10 resp. 200, see scripts) (Table 4). The table values are thus the smallest width found on the grid used, W_{test} , not necessarily the exact minimum over every $W \in \mathbb{N}$.

Table 4: Smallest width found on the grid used for $\varepsilon = 10^{-4}$, model M_1 (Variant A) vs. M_2 (Variant C). W_{test} , not $W_{\min}$: grid values, not proof of the exact minimum.

Zero #	t	$W_{\text{test}}(M_1)$	$W_{\text{test}}(M_2)$	Factor
650	1001.3495	9000	810	11.11
675	1031.8332	9200	830	11.08
700	1062.9154	9400	850	11.06
725	1093.4409	9600	870	11.03
749	1122.4586	9600	890	10.79




Figure 5: Convergence of M_1 and M_2 towards $\zeta(s)$ at zero #650 (real computed data).

In the five tested cases, the reduction factor ranged between 10.79 and 11.11 (Table 4); no general asymptotic constant is derived from this. What is reproducibly demonstrated for these five points is only the comparison of the values found on the grid:

$$W_{\text{test}}(M_2; s, \varepsilon) < W_{\text{test}}(M_1; s, \varepsilon), \quad \text{concretely} \quad 810 < 9000 \quad (11)$$

– not a proven statement about the exact minima $W_{\min}$ themselves, which could turn out slightly smaller on a finer grid. The model optimization

$$W_{\min}^*(s, \varepsilon) := \min_{M \in \mathfrak{M}_{\text{adm}}} W_{\min}(M; s, \varepsilon), \quad \mathfrak{M}_{\text{adm}} := \{M_1, M_2\} \quad (12)$$

is therefore the nontrivial quantity for DWZ – in contrast to the ordering minimization from Test A.

Test C – comparison with standard methods (B_{eval}). Tests A and B compare only variants *within* the DWZ framework. To provide the comparison against established standard methods

that the reviewer critique rightly requested, M_1 (Variant A, DWZ, order 1) is additionally compared against (a) a fifth-order Euler-Maclaurin expansion (M_{K5} , generally $\zeta_{\text{EM}}^{(m)}(s, W) = \sum_{n=1}^W n^{-s} + \frac{W^{1-s}}{s-1} - \frac{W^{-s}}{2} + \sum_{j=1}^m \frac{B_{2j}}{(2j)!} (s)_{2j-1} W^{-s-2j+1}$ with $(s)_k = s(s+1)\cdots(s+k-1)$ the rising factorial and B_{2j} the Bernoulli numbers; M_{K5} sets $m = 5$, i.e. five instead of one (Variant A) or two (Variant C) correction terms, the remainder term is not separately bounded) and (b) the Riemann-Siegel formula, evaluated via the `mpmath` functions $\vartheta(t)$ (`siegeltheta`) and $Z(t)$ (`siegelz`), with $\zeta(\frac{1}{2} + it) = e^{-i\vartheta(t)} Z(t)$. For each method, at a fixed accuracy target $\varepsilon = 10^{-6}$, we measure the minimal number of terms required (W for M_1, M_{K5} ; Riemann-Siegel terms $\sim \sqrt{t/2\pi}$) and the computation time of a *single evaluation* at that minimal term count (30-digit `mpmath` arithmetic, reproducible via the accompanying scripts). We denote this single-evaluation benchmark type B_{eval} throughout, to distinguish it explicitly from the cumulative, non-workload-matched benchmark B_{workflow} reported later (Section 3, “Runtime against an established reference method”) – the two measure different things and are not interchangeable, a distinction stated here in one place because the two figures otherwise appear close together in the text and could be conflated by a reader skimming for a single “the” benchmark factor. Regarding the absolute time values reported in Table 5 and Figure 6: these individual entries remain single measurements (no median over repeated runs, no warm-up, no confidence intervals) on the hardware/software environment available in this sandbox (processor, Python, and `mpmath` version details are documented in the accompanying scripts, not in the manuscript text). Single measurements in the millisecond range can be affected by system noise; taken in isolation, the reported order-of-magnitude factors (roughly two to three orders of magnitude) would therefore be only a preliminary, exploratory tendency, not an exact, repeatedly validated benchmark constant – the word “robust” is deliberately avoided for these single-shot table entries, since it would imply a statistical stability that a single, unreplicated measurement cannot support on its own. The paragraph immediately below reports an independent repeated-measurement check that supersedes this limitation for the qualitative conclusion, while the single-shot entries in the table are left unchanged for transparency and reproducibility of the original run.

Repeated-measurement check. As requested in review, we repeat the single-shot measurement above to check whether its qualitative conclusion survives replication. At zero #1 ($t = 14.1347$), M_1 ($W=11,498$) was timed over 15 repeated evaluations, M_{K5} ($W=10$) and Riemann-Siegel (via `mpmath`’s `siegeltheta/siegelz`) over 30 repeated evaluations each, on the sandbox environment used for this revision: $M_1 = 146.9 \pm 2.2$ ms ($n = 15$), $M_{K5} = 0.343 \pm 0.009$ ms ($n = 30$), Riemann-Siegel = 0.53 ± 0.21 ms ($n = 30$) (mean $\pm$ standard deviation). The standard deviations are small relative to the means for M_1 and M_{K5} (coefficient of variation below 2% for both), but noticeably larger for Riemann-Siegel (coefficient of variation below 40%, its absolute times being near the resolution limit of `time.perf_counter` on this system). Because of this asymmetry, the two resulting ratios should not be read with the same precision: $M_1/M_{K5} \approx 430\times$, backed by low variability on both sides, is a comparatively precise figure, whereas $M_1/\text{Riemann-Siegel} \approx 280\times$ carries the wider uncertainty of the Riemann-Siegel measurement and should be read as an order-of-magnitude estimate (a few hundredfold) rather than a precise benchmark factor. Both ratios are nonetheless consistent in order of magnitude with the single-shot figures in Table 5 (measured on different hardware, hence the absolute values differ slightly but not the conclusion). The qualitative two-to-three-orders-of-magnitude disadvantage of M_1 is thus confirmed under repetition, not merely asserted from a single run; this does not retroactively justify calling the original single-shot table “robust” – that wording remains avoided – but it does show that repeating the measurement would not have overturned the conclusion.

The result is unambiguous and is reported here without embellishment: the Riemann-Siegel formula reaches the same accuracy with 1–12 terms and consistently under 16 ms, while M_1 (DWZ) requires, depending on t , between 11,500 and over 200,000 terms and 0.15–2.8 s – an efficiency disadvantage of roughly two to three orders of magnitude. Simply extending to five

Table 5: Systematic comparison: required terms and computation time for $\varepsilon = 10^{-6}$, against M_1 (DWZ), M_{K5} (fifth-order Euler-Maclaurin), and Riemann-Siegel.

Zero #	t	$W(M_1)$	Time(M_1)	$W(M_{K5})$	Time(M_{K5})	Terms/Time (R.-S.)
1	14.1347	11,498	0.151 s	10	0.0004 s	1 / 0.0016 s
25	88.8091	42,695	0.635 s	43	0.0008 s	3 / 0.0009 s
50	143.1118	55,504	0.742 s	73	0.0011 s	4 / 0.0058 s
650	1001.3495	206,087	2.775 s	466	0.0060 s	12 / 0.0156 s




Figure 6: At zero #650 ($t = 1001.3495$): left, error against number of terms W ; right, error against computation time; for M_1 , M_{K5} , and Riemann-Siegel (single point, since it is nearly independent of W).

Euler-Maclaurin correction terms (M_{K5}) already reduces this disadvantage drastically (factor 100–400 in W , factor 130–460 in time relative to M_1), but is itself no substitute for Riemann-Siegel and is numerically more delicate: the Euler-Maclaurin series is asymptotic, not convergent, and the correction terms can blow up in magnitude if W is too small relative to $|s|$, before the series converges cleanly for sufficiently large W (observed numerically; not shown in Table 5, whose W values already lie in the stable regime). M_1 reduces this particular high-order asymptotic-instability risk through its simpler correction form – it retains only the first correction term, so it cannot exhibit the specific blow-up behavior possible when a higher-order asymptotic series is truncated at too small a W – although it does not eliminate all finite-width numerical or truncation effects in general, at the cost of requiring substantially more terms. The scientific contribution of this manuscript therefore does not lie in an efficiency advantage over established methods – no such advantage is claimed anywhere – but in the explicit, minimal operationalization of the DRCC stability condition through a single correction order, together with its empirical characterization (Section 3).

Comparative complexity analysis (time, state, domain). Beyond the empirical comparison of Test C, the comparison can also be conducted at the level of asymptotic complexity – with due care about which of the following statements are actually supported by this work and which are not.

The four rows are thus supported to different degrees: the time comparison is a direct measurement (Table 5), the state comparison a formal proof (Lemma 1), while the domain of validity is only the portion actually tested in this work – a claim of validity over the entire complex plane would not be supported by our own tests and is not made here. No comparative test on stability near Gram points was carried out; an advantage of Kahan compensated summation over Riemann-Siegel at such points is therefore not claimed – Kahan summation compensates floating-point rounding error over long addition chains, which is a different phenomenon from

Table 6: Comparative complexity: Riemann-Siegel vs. DWZ (M_1). The marker at the end of each row indicates how the corresponding claim is supported.

Criterion	Riemann-Siegel	DWZ / CPR (M_1)
Time per evaluation	$O(\sqrt{t})$ terms (established result)	$O(W_{\text{trunc}})$ per construction; empirically $W_{\text{trunc}}(t, \varepsilon) \sim t^{0.68}$ at fixed ε (regression fit over the four measured points in Table 5, $R^2 = 0.999$; consistent with the exponent $2/3$ predicted by equation (15)) – <i>measured</i>
State width	$O(1)$ (or $O(\log t)$ for the bit width of the index/loop register, standard convention)	$\omega_{\text{rec}} = O(1)$, independent of W_{trunc} – <i>proven</i> (Lemma 1)
Domain of validity	established primarily on the critical line $\text{Re}(s) = \frac{1}{2}$	tested in this work for the critical strip $0 < \text{Re}(s) < 1$; the underlying Euler-Maclaurin expansion is classically continuable analytically to a larger domain, but this is not independently tested here – <i>not tested</i>
Numerical stability near Gram points	documented in the literature as sensitive (sign changes of $Z(t)$) [3]	no comparative test performed in this work – <i>no claim</i>

the structural sign degeneracy of $Z(t)$ near Gram points, a property of the analytic function itself rather than of floating-point arithmetic; conflating the two would be misleading.

The defensible core of this comparison is thus: Riemann-Siegel wins on time by a wide margin ($O(\sqrt{t})$ versus empirically $O(t^{0.68})$ terms at fixed ε); for DWZ/CPR, in contrast, it is *proven*, not merely observed, that the state width stays constant ($\omega_{\text{rec}} = O(1)$, Lemma 1), independently of the truncation cost. This contrast – a measured time disadvantage against proven state-width constancy – is the differentiated claim actually supported here; broader formulations (validity over the entire complex plane, Gram-point robustness via Kahan summation) are deliberately not adopted, since they go beyond what is actually tested in this work.

CPR-width is not a runtime notion. An important conceptual clarification follows from this analysis: CPR-width and runtime complexity are distinct notions that cannot be traded off against each other. The argument in Lemma 1 rests solely on the fact that the terms k^{-s} are deterministically reconstructible from the index k and the parameter s and need not be held simultaneously; the main sum of the Riemann-Siegel formula (terms $n^{-1/2} \cos(\vartheta(t) - t \ln n)$) has the same property. This is not, however, a general claim about every implementation of the Riemann-Siegel formula, but only about a reconstruction model that shares the same assumptions: under a reconstruction model with the same assumptions about deterministic recomputability, a constant state width $\omega_{\text{rec}} = O(1)$ appears plausible for the Riemann-Siegel main sum as well, under the task τ_{out} ; a corresponding formal proof, however, is *not* carried out in this work (an instantiation of Definition 2.1 from Hesamiy [7] for the Riemann-Siegel algorithm was not undertaken here), and the claim is therefore not asserted, but stated as a plausible, open corollary. If it held, an important consequence would follow: two algorithms can share the same CPR-width class ($O(1)$) and still differ by orders of magnitude in runtime, exactly as the times measured in this section and in Section 3 show. The state-width finding for DWZ reported above is thus not a claim of an advantage over Riemann-Siegel, but a claim about

the structure of the DWZ reconstruction itself, regardless of whether Riemann-Siegel shares the same property.

Conceptual caveat. The identification $W := d_{\text{frag}}(X)$ introduced in Section 2 (equation (5)) denotes a numerical truncation depth: the number of explicitly summed series terms. This is not automatically the same as the structural CPR-Width ω_{CPR} from Hesamiy [7], which counts the number of components that must be *simultaneously kept active* at an intermediate stage. The following lemma makes this conjecture precise and provable for the most natural reconstruction-model case.

Register-model instantiation. The quantities $\omega_{\text{rec}}(P, \pi, \tau; \mathfrak{M})$, the active interface $B_i^{\pi, \tau}$, and “admissible reconstruction ordering” are the CPR-Width formalism of Definition 2.1 in Hesamiy [7]; they are used here, not re-derived. What this manuscript adds is a specific, self-contained instantiation for the sequential zeta task: let $\mathfrak{M}_{\text{seq}}$ denote the reconstruction-model family consisting of implementations that process terms $k = 1, \dots, N$ in ascending order and hold, at each step, exactly the quantities required to compute the next partial sum and nothing else. For this family, the active interface $B_i^{\pi_{\text{seq}}, \tau_{\text{out}}}$ at step i is *defined to be* the set of persistent logical state registers an implementation in $\mathfrak{M}_{\text{seq}}$ must hold simultaneously at that step – this identification is a modeling choice made explicit here, not an unstated assumption, and it is the sense in which the set $\mathcal{Z}_k$ constructed in the proof below is an admissible active interface for $\mathfrak{M}_{\text{seq}}$.

Lemma 1 (State width of sequential zeta accumulation). *Let τ_{out} be the task “output only the final value S_N ”, and let π_{seq} be the admissible reconstruction ordering that processes the terms $k = 1, \dots, N$ in ascending index order. Then there exists an admissible realization with*

$$|B_i^{\pi_{\text{seq}}, \tau_{\text{out}}}| = O(1) \quad \text{for all } i,$$

independent of N ; in particular $\omega_{\text{rec}}(P, \pi_{\text{seq}}, \tau_{\text{out}}; \mathfrak{M}_{\text{seq}}) = O(1)$.

Proof. The series $S_N = \sum_{k=1}^N k^{-s}$ is evaluated via the recursion $S_k = S_{k-1} + f(k, s)$ with $f(k, s) = k^{-s}$. At time $k-1$, S_{k-1} has been computed; determining the successor state S_k requires, by the recursion, only the evaluation of $f(k, s)$, with input quantities: the previous partial sum S_{k-1} , the running index k , and the fixed parameter s . Hence the minimal active state at step k is

$$\mathcal{Z}_k = \{S_{k-1}, k, s\}, \quad |\mathcal{Z}_k| = 3.$$

Each element of $\mathcal{Z}_k$ occupies a fixed number c of registers or memory words, independent of N . Since $f(k, s)$ is evaluated directly at step k and immediately accumulated into S_k , there is no informational dependence on earlier terms $f(k-1, s), f(k-2, s), \dots$ – these are overwritten after addition and removed from the active state space. For an error-corrected variant (Kahan summation), the state grows only by one constant correction variable c_{k-1} ,

$$\mathcal{Z}_{k, \text{Kahan}} = \{S_{k-1}, k, s, c_{k-1}\}, \quad |\mathcal{Z}_{k, \text{Kahan}}| = 4.$$

In both cases $|\mathcal{Z}_k|$ is constant and independent of the truncation index N . Since $\mathcal{Z}_k$ constitutes an admissible active interface for τ_{out} under the sequential ordering π_{seq} (every step is deterministically executable from $\mathcal{Z}_k$, and S_N is fully reconstructed at the end), it follows that $|B_i^{\pi_{\text{seq}}, \tau_{\text{out}}}| \leq |\mathcal{Z}_k| = O(1)$ for all i , and hence, by the definition of ω_{rec} as a minimum over admissible realizations, $\omega_{\text{rec}}(P, \pi_{\text{seq}}, \tau_{\text{out}}; \mathfrak{M}_{\text{seq}}) = O(1)$. Since $\omega_{\text{rec}} = O(1)$ independent of N , this proves that the truncation index W_{trunc} (resp. N) has no effect on the register or CPR width of the running algorithm. Under τ_{out} it acts exclusively as a scaling factor for the loop’s running time, not for the state width. $\square$

Word-count versus bit-space. The quantity actually shown to be $O(1)$ by this lemma is the number of persistent logical state components (equivalently, registers or memory words) in a fixed-precision word model – $|\mathcal{Z}_k| = 3$ (resp. 4 with Kahan compensation), independent of N . This is not the same statement as $O(1)$ *bit-space* complexity: the index k requires $\Theta(\log N)$ bits to represent for k up to N , and the precision needed for the complex accumulator S_{k-1} may itself grow with the target accuracy ε or with N , depending on the numerical representation chosen (as in the multi-precision `mpmath` arithmetic used throughout the numerical experiments of this work, where working precision is fixed by hand, not derived from a bit-complexity analysis). $\omega_{\text{rec}} = O(1)$ in this manuscript therefore means: constant *word count* in a fixed-precision model, not constant total bit-space as $N \rightarrow \infty$. Table 6 already notes the standard $O(\log t)$ bit-width convention for the loop index in passing; this paragraph makes the same point explicit at the level of the lemma itself, where it is most load-bearing.

Scope of the lemma. In one sentence: Lemma 1 is not sufficient to establish the original, stronger headline identification $\omega_{\text{CPR}} = W_{\text{trunc}}$, but it remains a valid and non-trivial structural result about the sequential reconstruction model in its own right – it formally separates runtime length from active state width, which the naive identification had conflated. More precisely: Lemma 1 shows that, for the natural task “output only the final value”, the identification $\omega_{\text{CPR}} = W_{\text{trunc}}$ does *not* hold – ω_{CPR} is constant under τ_{out} , while $W_{\text{trunc}} = N$ grows without bound. This direction of the originally open question is thus now proven, not merely conjectured. Whether there exists a stricter, separately declared task (e.g., requiring that every term remain individually reconstructable or auditable afterward) under which an active interface of size $\Theta(W)$ is genuinely required remains a separate, open construction question; it is neither claimed nor shown here. What remains defensible in addition is the numerically demonstrated statement (11): a stronger analytic reconstruction model reduces the required DWZ truncation width by more than a factor of ten. The relationship of this truncation width to the active CPR-Width remains – apart from the naive case now excluded by Lemma 1 – an open, model-dependent question.

Even under a stricter task: the reconstruction dilemma. One might suspect that a stricter task τ_{audit} – “every term must remain individually reconstructable afterward” – automatically forces an active interface of size $\Theta(W)$. This is not the case, however, as long as the declared reconstruction model permits recomputation as an admissible reconstruction operation: since each term k^{-s} is a pure, deterministic function of k and s , its value can be recomputed at any time from the pair (k, s) without ever needing to be stored in between. The entire information content of the sequence of terms is thus exhausted by the compact descriptor (N, s) ; a linear state width could only be justified if the sequence of terms admitted no compact description – which is evidently not the case for an analytic function such as k^{-s} , since it is fully determined by the closed-form map $(k, s) \mapsto k^{-s}$. (This paragraph deliberately states this intuition without invoking a formal notion of Kolmogorov complexity – such a formalization, including its dependence on the representation of s and the required output accuracy, is not carried out here; the load-bearing, formal argument of this section remains the signature construction $\text{Sig}_{\tau_{\text{audit}}}$ carried out above, not this informal remark.) Formally, this interchangeability is captured by the reconstruction-safe future-signature family Sig_{τ} from Hesamiy [7]: two prefixes with an identical set of future continuations may be assigned to the same equivalence class, allowing terms that have already been processed but remain recomputable at any time to be removed from the active interface. This extension of Lemma 1 holds, however, only under reconstruction models that declare recomputation as an admissible operation; a model that instead requires the persistent, literal storage of all W raw values would remain unaffected and is not excluded here.

Formal instantiation of $\mathfrak{M}$ for τ_{audit} . We instantiate the declared reconstruction structure from Hesamiy [7], Definition 2.1 (chain $\mathcal{R}_i^{\pi} \rightarrow E_i^{\pi}(r) \rightarrow \text{Sig}_{\tau,i}^{\pi}(r) \rightarrow B_i^{\pi,\tau} \rightarrow \omega_{\text{rec}}$), explicitly

for the sequential zeta accumulation under the stricter task τ_{audit} (“every term must remain individually reconstructable afterward”). Let P be the finite instance with component set $V(P) = \{1, \dots, N\}$ (the term indices) and fixed parameter s . As reconstruction ordering we choose the ascending order $\pi_{\text{seq}} = (1, 2, \dots, N)$; it is admissible, $\text{Adm}_{P, \tau_{\text{audit}}}(\pi_{\text{seq}}) = 1$, since it matches the accumulation order actually implemented. The reachable prefix state after i steps is initially the full history of raw terms,

$$r_i \in \mathcal{R}_i^{\pi_{\text{seq}}}(P), \quad r_i = (f(1, s), \dots, f(i, s)),$$

i.e., the naive, not necessarily minimal candidate interface $\tilde{B}_i^{\pi_{\text{seq}}, \tau_{\text{audit}}} = r_i$ with $|\tilde{B}_i| = i$. The reconstruction-safe future-signature family for τ_{audit} assigns to each prefix r_i the response function to all admissible audit queries,

$$\text{Sig}_{\tau_{\text{audit}}, i}^{\pi_{\text{seq}}}(r_i) = \left(S_i(r_i), (j \mapsto f(j, s))_{1 \leq j \leq i} \right),$$

where $S_i(r_i) = \sum_{k=1}^i f(k, s)$ is the current partial sum. Since $f(j, s) = j^{-s}$ for each $j \leq i$ depends only on j and s – not on the specific history r_i from which j originated –, two prefixes with identical (S_i, i, s) induce the same signature $\text{Sig}_{\tau_{\text{audit}}, i}^{\pi_{\text{seq}}}$, regardless of their full history. Hence

$$E_i^{\pi_{\text{seq}}}(r_i) := \{S_i(r_i), i, s\}$$

is a task-sufficient representative system for $\text{Sig}_{\tau_{\text{audit}}, i}^{\pi_{\text{seq}}}(r_i)$, and by the cardinality-minimal construction from Chapter 2 of Hesamiy [7],

$$|B_i^{\pi_{\text{seq}}, \tau_{\text{audit}}}| \leq |E_i^{\pi_{\text{seq}}}(r_i)| = 3$$

for all i (one additional register for a Kahan correction raises this bound to 4). By the definition of ω_{rec} as a maximum over i and a minimum over admissible realizations, it follows that

$$\omega_{\text{rec}}(P, \pi_{\text{seq}}, \tau_{\text{audit}}; \mathfrak{M}_{\text{seq}}) \leq 3 = O(1).$$

This instantiation formally extends Lemma 1 to the stricter task τ_{audit} and – unlike a merely intuitive analogy – is anchored directly in the cited Definition 2.1: not only the output task but also the audit task yields $\omega_{\text{CPR}} = O(1)$, as long as the reconstruction model admits recomputation as a valid operation. What remains open is exclusively the question of whether there exists a task whose Sig_{τ} denies this compression – such a task is neither constructed nor excluded here.

Conceptual separation. We therefore formally separate the numerical resource parameter $W_{\text{trunc}}(t, \varepsilon)$, which describes the empirical truncation depth needed to reach the approximation accuracy ε , from the algorithmic state width ω_{CPR} . The latter defines the maximum number of simultaneously active components required for future state continuations within the declared reconstruction model. While $\omega_{\text{rec}} = O(1)$ holds for the sequential accumulation of the zeta approximation – since the terms are deterministically reconstructable from k and s – the formal relationship between W_{trunc} and ω_{CPR} , under models that exclude recomputation, depends on the specifically chosen automaton model $\mathfrak{M} = (\text{Adm}, \mathcal{R}, \Lambda, \delta, \text{Sig}_{\tau})$ and is the subject of a separate analysis. It is worth stressing further that $\omega_{\text{CPR}} = O(1)$ for τ_{audit} concerns only the *storage width* at any given moment, not the *time cost* of an audit request itself: each individual recomputation of a term k^{-s} takes constant time, but a full audit of all N terms requires $\Theta(N)$ such recomputations and hence $\Theta(N)$ total time across all requests. Storage width, reconstruction time per request, and total time for all requests are thus three distinct quantities; only the first is captured by $\omega_{\text{CPR}} = O(1)$.

Cost model assumptions for the audit task. The instantiation above relies on four modeling choices that are stated here explicitly rather than left implicit. First, the parameter s is treated as encoded at a fixed working precision, chosen once for the whole computation and independent of the term index k or the truncation length N ; a precision that must itself grow with N or with k would not change the register *count* in $\mathcal{Z}_k$, but would affect the size of each register (see the word-count-versus-bit-space discussion above), which the $O(1)$ claim does not address. Second, a recomputed term $f(k, s) = k^{-s}$ is assumed to reproduce the same value, to the same working precision, as if it had been retained from storage; no separate error analysis for recomputed versus stored terms is carried out here, and none is needed for the state-width claim itself, which concerns which quantities must be held, not how accurately they are computed. Third, evaluating $f(k, s)$ for a single k is treated as an elementary, constant-time operation at fixed working precision; this is the standard convention used throughout this manuscript’s runtime comparisons (Section 2.5), but it hides a real, precision-dependent cost – evaluating k^{-s} to p digits of precision via standard multiprecision exponentiation costs more than $O(1)$ bit operations in p (and, for very large k , in $\log k$ as well) – which is a statement about bit-complexity, not about the register count $|\mathcal{Z}_k|$ that Lemma 1 addresses. Fourth, and consequently, none of these choices affects the lemma’s actual claim: $\omega_{\text{rec}} = O(1)$ is, throughout this section, a statement about the number of persistent logical state components (a small, fixed set of registers), not about the bit-size of each register or the time cost of computing their contents – both of which may grow with precision and are explicitly out of scope for this particular claim.

Algorithmic illustration (register invariance). The following Python program realizes the reconstruction constructed in the proof of Lemma 1 using Kahan compensated summation, making the register invariance of ω_{rec} directly observable:

```

1 def naive_partial_sum_kahan(s: complex, N: int) -> complex:
2     """
3     Computes the NAIVE partial sum (equation (5)), not the DWZ-corrected
4     function (equation (7)), via sequential Kahan-compensated accumulation.
5     For  $0 < \text{Re}(s) < 1$  this partial sum diverges as  $N$  grows (Table 3); the
6     example below with  $s=2$  therefore deliberately lies outside the critical
7     strip and serves only to illustrate register behavior, not to validate
8     the DWZ method in the critical strip.
9     Space complexity (state width):  $O(1)$  -- constant
10    Time complexity (truncation length):  $O(N)$  -- linear
11    """
12    S = 0.0 # accumulator (running sum)
13    c = 0.0 # compensation term for rounding error (Kahan)
14    for k in range(1, N + 1):
15        term = k ** (-s)
16        y = term - c # subtract previous error from new term
17        t = S + y # add to accumulator
18        c = (t - S) - y # exact rounding error for next step
19        S = t # update state
20    return S
21
22 s_value = 2.0
23 truncation = 1_000_000
24 result = naive_partial_sum_kahan(s_value, truncation)
25 print(f"Partial sum({s_value}) approximation: {result}")

```

Regardless of whether $N = 10$ or $N = 10^9$ is chosen, the program allocates the same number of registers at every point in time ($S, c, k, s, \text{term}, y, t$) – no memory array of length N is ever created; each term exists only transiently in a register and is immediately overwritten after processing. Of these seven registers, $\{S, c, k, s\}$ form the *persistent* active state $\mathcal{Z}_{k, \text{Kahan}}$ defined

in Lemma 1 (carried forward from step k to step $k+1$, $|\mathcal{Z}_{k,\text{Kahan}}| = 4$); term , y , and t are, by contrast, purely *transient* intermediate values that are computed and discarded within a single step and are not carried forward to the next step – they therefore do not count toward the formal interface B_i from Lemma 1. The total of seven simultaneously allocated registers is a property of *this particular implementation* and a valid upper bound ($\omega_{\text{rec}} \leq 7$ for this code), but not a tighter figure than the bound proven in the lemma for the persistent state, namely 3 (without Kahan) or 4 (with Kahan). The variables c , y , t preserve numerical accuracy even for very large N , without the algorithmic state width ever growing beyond $O(1)$. For $s = 2$ and $N = 10^6$, the program returns $1.6449330668\dots$ against the known limit $\pi^2/6 = 1.6449340668\dots$, a deviation of about 10^{-6} – consistent with the expected truncation error of a directly truncated Dirichlet series at this N (numerically verified). The algorithmic implementation thus confirms the theoretical derivation of Lemma 1: the state width remains strictly at $\omega_{\text{rec}} = O(1)$ throughout the entire run, while W_{trunc} (here controlled via the loop limit N) scales exclusively the running time, not the register or CPR width.

2.6 The DRCC Retraction Pair: What It Does and Does Not Establish

A reviewer asked for an explicit justification, in DRCC terms, of why adding the correction term $C_W(s)$ in (7)/(9) to the naive fragment $F_W(s)$ counts as an admissible “reconstruction” in the sense of Hesamiy [7]. We make the minimal formalization that answers this explicit – together with an equally explicit statement of what it does not show, since the construction turns out to be weaker than it may first appear.

Definition 1 (Reconstruction operator). *Let $D \subseteq \mathbb{C}$ denote the domain of the complex parameter s under consideration (e.g. the critical strip). For a fixed correction term $C_W(s)$ – equal to $-W^{-s}/2$ for M_1 (equation (7)), resp. $-W^{-s}/2 + \frac{s}{12}W^{-s-1}$ for M_2 (equation (9)) – define*

$$S_W : D \rightarrow \Omega, \quad S_W(s) := (F_W(s), C_W(s)), \quad \Omega := \mathbb{C} \times \mathbb{C},$$

with $F_W(s) = \sum_{n=1}^W n^{-s}$ the retained fragment (equation (5)). Note that S_W takes the parameter $s \in D$ as its argument, not the value $F_W(s)$ – this avoids any ambiguity in case F_W fails to be injective on D .

Definition 2 (Reduction operator). *Define $R_W : \Omega \rightarrow \mathbb{C}$ by $R_W(F, C) := F$.*

Proposition 1 (Packaging Consistency Identity for M_1 and M_2). *$R_W \circ S_W = F_W$ as functions $D \rightarrow \mathbb{C}$.*

Proof. For every $s \in D$, $R_W(S_W(s)) = R_W(F_W(s), C_W(s)) = F_W(s)$, directly from Definitions 1 and 2. $\square$

What this does not show. The proof of Proposition 1 uses only the defining equation of R_W and never uses the specific value or analytic origin of $C_W(s)$. The identity $R_W \circ S_W = F_W$ therefore holds for *every* choice of correction term whatsoever – including a correction that is analytically wrong, e.g. $C_W(s) \equiv 0$ or $C_W(s) := 42$. As formalized here, the retraction pair property is a tautological consistency check on the packaging $(F, C) \mapsto F$: it confirms only that the reduction operator was defined consistently with the reconstruction operator, not that the reconstruction is numerically sound or even sensible. It therefore cannot, by itself, serve as a criterion for choosing between M_1 , M_2 , or any other additively-corrected model – that evidence is supplied separately, by the convergence properties of (7)/(9) and by the empirical results of Section 3, not by Proposition 1 itself.

A second limitation concerns Variant B. $\zeta_W^\eta(s, W) = \eta_W(s, W)/(1 - 2^{1-s})$ is a quotient of a partial sum by a fixed scalar, not a sum of the form $F_W(s) + C_W(s)$ that Definitions 1 and 2 assume. No analogous (S_W, R_W) pair for M_η is constructed in this manuscript, so the retraction-pair admissibility of M_η is neither proven nor assumed here; it remains open. This matters for the selection procedure of Section 2.7.

2.7 Preliminary RP-Tensor Scoring Framework for Reconstruction Variants

Section 2.5 shows that M_1 , M_η , and M_2 differ measurably in accuracy and required width but stops short of a declared selection rule. We formalize the selection problem, and then check, criterion by criterion, how much of it the manuscript’s own data currently support – rather than presenting a completed optimization that has not actually been carried out.

Formal setup. Let $V_{\text{crit}} = \{c_{\text{err}}, c_{\text{run}}, c_{\text{depth}}, c_{\text{width}}, c_{\text{stab}}\}$ and, for each model $M_\alpha \in \{M_1, M_\eta, M_2\}$ and evaluation context $\tau = (s, \varepsilon)$, let $T_{\alpha j}(\tau)$ be the raw value of criterion j for model α . With admissibility indicator $a_\alpha \in \{0, 1\}$ from Section 2.6 and weights $w_j \geq 0$, $\sum_j w_j = 1$, define, piecewise to avoid the ill-defined expression $0 \cdot \infty$ that a single combined formula would produce at $a_\alpha = 1$,

$$S_\alpha(\tau) = \begin{cases} \sum_j w_j \tilde{T}_{\alpha j}(\tau), & a_\alpha = 1, \\ +\infty, & a_\alpha = 0, \end{cases} \quad M^*(\tau) \in \arg \min_\alpha S_\alpha(\tau), \quad (13)$$

with $\tilde{T}_{\alpha j}$ the min-max normalization of $T_{\alpha j}$ over the admissible candidate set, and where “ $\in$ ” rather than “ $=$ ” is used because $\arg \min$ is in general set-valued (ties are possible).

Data availability check. Before choosing any weights, we check which cells of the 3×5 criterion table are actually backed by data already reported in this manuscript.

Table 7: Availability of RP-Tensor criteria for M_1 , M_η , M_2 in this manuscript’s own results. “Measured” = direct table entry; “extrapolated” = derived from measured points via a fitted or stated asymptotic law; “structural” = a definitional count, not an empirical measurement; “–” = not evaluated anywhere in this work.

Criterion	M_1	M_η	M_2
c_{err} (accuracy at fixed W)	measured	measured	measured
c_{width} (W at fixed ε)	measured	extrapolated	measured
c_{run} (wall-clock time)	measured	–	–
c_{depth} (EM correction order)	structural (1)	n/a (not an EM series)	structural (2)
c_{stab} (near Gram points)	–	–	–
a_α (retraction pair, §2.6)	proven	open	proven (same trivial argument)

Two of five criteria are entirely empty for two of the three models (c_{run} , measured only for M_1 against M_{K5} /Riemann-Siegel in Table 5; c_{stab} , explicitly not tested for any model near Gram points, Section 2.5), one criterion is a definitional count rather than a measurement, and M_η ’s admissibility itself is unresolved. A genuine five-criterion, three-model RP-Tensor cannot currently be computed. What can honestly be computed is a pairwise comparison restricted to the two criteria with real numbers already in this manuscript, c_{err} and c_{width} – at the two reference points where those numbers were actually measured.

With equal weights over the two available criteria ($w_{\text{err}} = w_{\text{width}} = 0.5$, all other $w_j = 0$ since undefined), equation (13) gives $S_{M_1} = 1$, $S_{M_2} = 0$ at $t = 1001.35$ (so $M^* = M_2$), and $S_{M_1} = 0$, $S_{M_\eta} = 1$ at $t = 14.13$ (so $M^* = M_1$).

Honest reading of these two computations. On the only two criteria for which this manuscript reports real numbers, M_2 dominates M_1 decisively at $t \approx 1001$, and M_1 dominates M_η decisively at $t \approx 14$. Combining these into a single three-way ranking additionally assumes – without separate verification – that the ordering transfers across t , since the two comparisons were not measured at a common (s, ε) . With that caveat, an RP-Tensor evaluated on $\{c_{\text{err}}, c_{\text{width}}\}$ alone

Table 8: Pairwise RP-Tensor scores computed from Table 4 (width, $t = 1001.3495$, $\varepsilon = 10^{-4}$), Figure 5 (error at common $W = 2000$, same t), and Table 12 (error, $t = 14.1347$; width at $\varepsilon = 10^{-4}$ obtained by log-log interpolation for M_1 , resp. extrapolation via the fitted local exponent -0.4997 for M_η , consistent with the stated law $O(W^{-\text{Re}(s)})$). $\tilde{T} \in \{0, 1\}$: min-max normalization over a two-element set is necessarily binary and reports only which model is better, not by how much; the ratio column preserves that information.

Comparison	Criterion	M_1	other	Ratio (disadvantage)
M_1 vs. M_2 , $t=1001.35$	$W_{\text{test}}(\varepsilon=10^{-4})$	9000	810	$11.1 \times$
M_1 vs. M_2 , $t=1001.35$	error at $W=2000$	9.37×10^{-4}	3.92×10^{-6}	$239 \times$
M_1 vs. M_η , $t=14.13$	$W(\varepsilon=10^{-4})$	≈ 517	$\approx 4.46 \times 10^6$	$8633 \times$ (favors M_1)
M_1 vs. M_η , $t=14.13$	error at $W=5000$	3.33×10^{-6}	2.98×10^{-3}	$895 \times$ (favors M_1)

would select M_2 , not M_1 , as M^* . This directly contradicts reading M_1 as the output of an already-completed RP-Tensor optimization. M_1 is retained as the primary model in this manuscript for a different, already stated reason – minimality and direct traceability of the construction (Section 2) – not because a computed selection procedure ranked it first on error, width, runtime, or stability. That reason stands on its own, but it is not one of the five declared criteria in V_{crit} . Either a sixth, explicitly weighted criterion for structural simplicity should be added to equation (13), or the choice of M_1 should be stated plainly as an expository decision that sits outside the RP-Tensor ranking rather than as its result.

A caveat on min-max normalization with few candidates. With only two or three candidate models, min-max normalization maps every criterion to $\{0, 1\}$ regardless of the actual size of the gap: a 10% disadvantage and a 10,000% disadvantage both become $\tilde{T} = 1$. Table 8’s ratio column ($11 \times$ to $8633 \times$ depending on criterion and reference point) carries information the normalized score alone discards. Reporting the raw ratio alongside the binary score, or normalizing against a fixed external reference (e.g. a target-accuracy budget) rather than min-max over the candidate set itself, would make the RP-Tensor more informative than a bare ranking.

Remaining work before the RP-Tensor is complete. c_{run} and c_{stab} have no data for M_η or M_2 in this manuscript. Completing them requires two further numerical experiments – a runtime benchmark for M_η and M_2 at matched accuracy targets, and a Gram-point stability test across all three models – together with a retraction-pair construction for M_η (Section 2.6) before its admissibility a_{M_η} can be assigned rather than left open. Until then, equation (13) is a declared selection framework, not yet a completed calculation, and the two-criterion computation above should be read accordingly: as evidence that M_2 currently outperforms M_1 on the measured axes, not as a full justification for using M_1 going forward.

2.8 Operationalizing $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$

We denote by $\rho_k = \frac{1}{2} + it_k$ the k -th reference zero (`mp.zetazero(k)`) and by $s_k^{(M)}(W)$ the zero numerically determined by model M and width W that is assigned to ρ_k according to the formal assignment rule $k^* = \arg \min_j |s^{(M)}(W) - \rho_j|$ (nearest reference zero; this same rule underlies the evaluation of the hybrid hardening runs in Section 3, where it produces the neighbor hits reported there whenever zero spacing is small); the corresponding position error is $e_k^{(M)}(W) := |s_k^{(M)}(W) - \rho_k|$ (to distinguish it from the function-value error $E_\zeta(s, W) := |\zeta_M(s, W) - \zeta(s)|$, used in Table 3 and Figure 5). The numerical measurement of zero convergence (Section 3, model M_1) yielded the empirical power law $e_k^{(M_1)}(W) \approx C(t_k) W^{-p}$ with $p \approx 1.500$ – 1.502 (determined

numerically, not proven exactly as $3/2$). The actual minimal truncation depth is thus defined, in the same *stable* form used in Section 2.8.3, as an integer minimum over widths beyond which the error remains below tolerance,

$$W_{\text{stab}}^{(M_1)}(t, \varepsilon) := \min\{ W_0 \in \mathbb{N} : e_k^{(M_1)}(W) < \varepsilon \text{ for every } W \geq W_0 \}, \quad (14)$$

where $e_k^{(M_1)}(W)$ by definition measures the distance to the already-known reference zero ρ_k . Crucially, $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ is thus a *post-hoc* error or calibration measure against an already-known zero, not a standalone criterion by which an as-yet-unknown zero could be blindly located – a blind search (Section 2.9) instead uses the scan minimum of $|\zeta_{\text{DWZ}}|$, not $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ itself. The statement that $W_{\text{stab}}^{(M_1)}$ "operationalizes reconstruction" (Section 2.8.1) thus refers, at this stage, to accuracy against known reference values, not to a procedure for finding unknown zeros. Here $t := t_k$ is to be set when $\rho_k = \frac{1}{2} + it_k$ is the corresponding reference zero – equation (14) itself is thus, at first, only evaluated at the discrete points $t = t_k$ for which a reference zero and hence an error curve $e_k^{(M_1)}(W)$ exist, not directly defined for arbitrary real t . The stable form of the definition builds the requirement that accuracy, once reached, persists directly into $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ itself, rather than leaving it as a separate assumption; whether the defining set is nonempty for the tested pairs, however, still rests on a numerical observation, not a proof. By the definition of a limit, convergence alone, $e_k^{(M_1)}(W) \rightarrow 0$ as $W \rightarrow \infty$, already suffices to guarantee that such a W_0 exists for every $\varepsilon > 0$ – monotonicity is not required for this conclusion, although it is, in fact, also observed in the tested range (Figure 12), as a strictly stronger property than what the argument here actually needs. This convergence itself is numerically observed for the tested pairs, not analytically proven for arbitrary W , which makes the stable minimum in equation (14) well-defined for the tested (t_k, ε) pairs on that basis. The extension to arbitrary t , including values that do not coincide with a reference zero, is achieved only through the continuous relationship $C(t)$ fitted below, not through equation (14) itself. The power law provides an estimate for this: in general, $e(W) \approx C(t)W^{-p}$ implies the form $W \approx (C(t)/\varepsilon)^{1/p}$, and only under the – numerically plausible but unproven – assumption that $p = 3/2$ exactly does this yield the special exponent $2/3$:

$$\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon) = \left\lceil \left(\frac{C(t)}{\varepsilon} \right)^{1/p} \right\rceil \stackrel{p \approx 3/2}{\approx} \left\lceil \left(\frac{C(t)}{\varepsilon} \right)^{2/3} \right\rceil \approx W_{\text{stab}}^{(M_1)}(t, \varepsilon), \quad (15)$$

with $C(t) \approx 0.0272t + 1.043$ (linear fit over the first 50 zeros #1–#50 from Table 15, $n = 50$; Table 15 itself has since been extended to 100 zeros, but the $C(t)$ fit still covers only the first 50, see Table 13; slope 0.0272 ± 0.0040 , intercept 1.043 ± 0.375 , 95% confidence intervals). The fit explains only part of the scatter ($R^2 \approx 0.50$); neither a square-root, nor a logarithmic, nor a pure power-law form in t fits noticeably better ($R^2 \approx 0.49, 0.47$, and 0.48 respectively). $C(t)$ increases significantly with t (the trend is highly significant at $n = 50$, $p \ll 0.001$), but the simple linear form explains only about half of the observed scatter – the other half likely depends on properties not captured by t alone. Equation (14) is a *definition*; equation (15) is a derived, *empirical estimator* – the two are deliberately not treated here as exactly identical. This turns the originally abstract inequality (4) into a closed-form, interpretable formula: $W \geq \hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$ means, approximately, " W is large enough that the zero at given t is determined to within ε ". The quantity $\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$ is a model-dependent numerical truncation threshold. It is motivated by, but deliberately given different notation from, the original, abstract DRCC reconstruction degree $d_{\text{rec}}^{\text{DRCC}}(X)$ from Section 2 – this notational separation marks a concrete interpretation for the zeta case, not a proven identity, and is intended to prevent exactly the notational overloading that a shared symbol would otherwise invite.

2.8.1 Conditional Numerical Operationalization

The preceding sections motivate the concretization of inequality (4) through prose and numerical examples. The following proposition makes this step explicit and conditionally provable – as a direct answer to the legitimate criticism that the DRCC→DWZ transition has so far only been asserted, not demonstrated – given one stated, numerically (not analytically) supported hypothesis.

Proposition 2 (Conditional Numerical Operationalization for the Zeta Function). *Let $t = t_k$ for one of the tested reference zeros $\rho_k = \frac{1}{2} + it_k$ ($k = 1, \dots, 100$ from Table 15, for which $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ is defined per equation (14)), $\varepsilon > 0$ fixed, and let*

$$d_{\text{frag}}(X) := W \quad (\text{equation (5)})$$

$$d_{\text{rec}}^{\text{DRCC}}(X) := W_{\text{stab}}^{(M_1)}(t, \varepsilon) \quad (\text{equation (14)})$$

be the concretizations chosen in this work. Assume that

$$e_k^{(M_1)}(W) \longrightarrow 0 \quad (W \rightarrow \infty) \quad (16)$$

for this reference zero – a hypothesis that is numerically observed for the tested pairs (Figure 12), not analytically proven. Then, for every finite width $W \geq W_{\text{stab}}^{(M_1)}(t, \varepsilon)$, the DRCC stability condition (4) holds in its concretized form,

$$0 < W_{\text{stab}}^{(M_1)}(t, \varepsilon) \leq W < \infty,$$

and moreover

$$e_k^{(M_1)}(W) < \varepsilon.$$

Proof. By the definition of convergence, hypothesis (16) already implies that, for every $\varepsilon > 0$, there exists $W_0 \in \mathbb{N}$ with $e_k^{(M_1)}(W) < \varepsilon$ for every $W \geq W_0$; monotonicity of $e_k^{(M_1)}(W)$ is not required for this step, although it is, in fact, also observed numerically in the tested range (Figure 12), as a strictly stronger property than what is used here. Hence the defining set $\{W_0 \in \mathbb{N} : e_k^{(M_1)}(W) < \varepsilon \text{ for every } W \geq W_0\}$ is nonempty for the tested pairs, exactly as required by the defining condition later formalized in Proposition 4. Granting this, $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ is well-defined and finite as the least element of this nonempty subset of $\mathbb{N}$. For every finite $W \geq W_{\text{stab}}^{(M_1)}(t, \varepsilon)$, the defining property of that least element directly gives $e_k^{(M_1)}(W) < \varepsilon$, and the inequalities $0 < W_{\text{stab}}^{(M_1)}(t, \varepsilon) \leq W < \infty$ then follow from the ordering of W_0 and W in $\mathbb{N}$, as in the proof of Proposition 4. $\square$

Scope and limits of the proposition. Proposition 2 shows that the concretizations chosen in this work satisfy the *form* of the original DRCC inequality (4) without contradiction – the DRCC→DWZ transition is thus formal at this point, not merely motivational. The additional conclusion $e_k^{(M_1)}(W) < \varepsilon$ for every $W \geq W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ rests entirely on hypothesis (16) (convergence $e_k^{(M_1)}(W) \rightarrow 0$), which is numerically observed but not analytically proven for the tested pairs; it is not a global guarantee for every $t > 0$, and it is conditional, not unconditional, exactly as the proposition's title now states. Note, further, that the central inequality $W_{\text{stab}}^{(M_1)}(t, \varepsilon) \leq W$ follows here directly from the choice $W \geq W_{\text{stab}}^{(M_1)}(t, \varepsilon)$: the proposition primarily establishes the *consistency* of the notation so chosen with the form of equation (4), not any further nontrivial property derived from DRCC theory itself – more fitting labels would be "consistency of the numerical concretization" or "conditional numerical instantiation"; we retain "operationalization" here only because it continues the usage introduced in Section 1.2. The proposition does *not* show that $d_{\text{frag}}(X)$ corresponds to the structural fragmentation quantity from Hesamiy [7], that

$W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ corresponds to the abstract quantity $d_{\text{rec}}^{\text{DRCC}}(X)$, or that $\omega_{\text{CPR}} = W_{\text{trunc}}$ – the latter remains the open, empirically tested question of Section 2.5. The proposition is thus deliberately modest: it secures the logical consistency of the chosen concretization, not its identity with the original, abstract DRCC/CPR quantities.

Relationship of $\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$ (Eq. (15)) to the DRCC condition (Eq. (4)). Since this question recurs in review, the chain is stated explicitly and in one place here: there is *no* direct relationship between the estimator $\hat{W}_{\text{stab}}^{(M_1)}$ and equation (4), only a two-link chain of two results of different kinds. First, $\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$ (Eq. (15)) approximates the exact quantity $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ (Eq. (14)) – this is a purely *empirical* result from a linear fit over $n = 50$ zeros ($R^2 \approx 0.50$, see above); it is not a proven equality or a convergence statement, but an estimate with substantial unexplained scatter. Second, the exact quantity $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ itself – together with $d_{\text{frag}}(X) := W$ – satisfies the form of equation (4); this is Proposition 2, a *formally proven, but explicitly conditional* (on the numerically observed convergence hypothesis (16)) and purely structural result (see the paragraph above: not an identity with the abstract DRCC quantities). Composed, this means: $\hat{W}_{\text{stab}}^{(M_1)} \xrightarrow{\text{empirical, unproven}} W_{\text{stab}}^{(M_1)}(t, \varepsilon) \xrightarrow{\text{formal, conditional, Proposition 2}} \text{form of Eq. (4)}$. The two links are of different epistemic quality, and neither transfers its proof status to the other: the estimator remains a numerical approximation regardless of the fact that the underlying exact quantity can be formally embedded in the DRCC form.

2.8.2 Analytical Origin of the Observed $W^{-3/2}$ Rate

The numerical experiments reported in Section 3 show an exponent close to $3/2$ for the displacement of reconstructed zeros on the critical line.

This exponent is not arbitrary.

It is consistent with the first omitted Euler–Maclaurin correction term in the reconstruction model M_1 .

Let

$$s = \sigma + it, \quad 0 < \sigma < 1,$$

and recall the first-order DWZ approximation,

$$\zeta_{\text{DWZ}}(s, W) = \sum_{n=1}^W n^{-s} + \frac{W^{1-s}}{s-1} - \frac{1}{2}W^{-s}. \quad (17)$$

The next Euler–Maclaurin correction term – already present in this work as the additional term of Variant C, equation (9) – is

$$\frac{s}{12}W^{-s-1}. \quad (18)$$

Accordingly, for fixed s and sufficiently large W , the function-value error has the asymptotic form

$$\zeta(s) - \zeta_{\text{DWZ}}(s, W) = \frac{s}{12}W^{-s-1} + O_s(W^{-\sigma-3}), \quad (19)$$

where the notation $O_s(\cdot)$ indicates that the implied constant may depend on the fixed complex parameter s . This is the standard Euler–Maclaurin asymptotic heuristic (the truncation error is of the same order as the first neglected term), not a fully rigorous remainder bound in the sense of Titchmarsh’s contour-integral treatment of the periodic Bernoulli remainder [2]; no such rigorous bound is derived in this work.

Taking absolute values yields

$$|\zeta(s) - \zeta_{\text{DWZ}}(s, W)| = O_s(W^{-\sigma-1}). \quad (20)$$

On the critical line,

$$\sigma = \frac{1}{2},$$

and therefore

$$|\zeta(s) - \zeta_{\text{DWZ}}(s, W)| = O_s(W^{-3/2}). \quad (21)$$

This explains analytically why an exponent near $3/2$ appears in the numerical experiments. We verified this explanation numerically (not merely asserted it): for the reconstruction error at zero #1 ($t = 14.1347$), $t = 100$, and $t = 1000$, the ratio of the actual measured error $|\zeta_{\text{DWZ}}(s, W) - \zeta(s)|$ to the predicted magnitude $|s|/12 \cdot W^{-3/2}$ converges to 1.000 as W grows large relative to t (e.g. ratio 1.000 at $W = 1000$ for all three t values tested; for $t = 1000$ the ratio is still 4.35 at $W = 50$, confirming that the asymptotic regime requires W large relative to t , consistent with the truncation depths actually used elsewhere in this work).

It remains necessary, however, to distinguish the function-value error

$$E_\zeta(s, W) = |\zeta_{\text{DWZ}}(s, W) - \zeta(s)|$$

from the zero-position error

$$e_k(W) = \left| s_k^{(M_1)}(W) - \rho_k \right|.$$

The following proposition gives the local connection between them.

Proposition 3 (Local Transfer from Function Error to Zero Error). *Let*

$$\rho$$

be a simple zero of ζ , so that

$$\zeta(\rho) = 0, \quad \zeta'(\rho) \neq 0. \quad (22)$$

Assume that, for all sufficiently large W , the approximation $\zeta_{\text{DWZ}}(\cdot, W)$ has a zero ρ_W in a fixed open neighborhood U of ρ (with U held fixed, independent of W), and assume that the approximation error is uniformly bounded on U by

$$\sup_{s \in U} |\zeta_{\text{DWZ}}(s, W) - \zeta(s)| \leq K_\rho W^{-\sigma_U-1}, \quad \sigma_U := \inf_{s \in U} \text{Re}(s), \quad (23)$$

for a constant $K_\rho > 0$; since $\text{Re}(s)$ is not constant on U , the exponent uses the worst-case (infimum) real part over U , not $\text{Re}(\rho)$ itself.

Then

$$|\rho_W - \rho| = O_\rho(W^{-\sigma_U-1}). \quad (24)$$

In particular, if ρ lies on the critical line and U is chosen so that it does not extend to the left of the critical line (e.g. $U \subseteq \{s : \text{Re}(s) \geq \frac{1}{2}\}$, so that $\sigma_U = \frac{1}{2}$ exactly by this choice),

$$|\rho_W - \rho| = O_\rho(W^{-3/2}). \quad (25)$$

Proof. Since ρ is a simple zero, the local Taylor expansion gives

$$\zeta(\rho_W) = \zeta'(\rho)(\rho_W - \rho) + O(|\rho_W - \rho|^2). \quad (26)$$

Because

$$\zeta_{\text{DWZ}}(\rho_W, W) = 0,$$

one has

$$\zeta(\rho_W) = \zeta(\rho_W) - \zeta_{\text{DWZ}}(\rho_W, W). \quad (27)$$

By equation (23) (using $\rho_W \in U$),

$$|\zeta(\rho_W)| \leq K_\rho W^{-\sigma_U - 1}. \quad (28)$$

For sufficiently large W , the quadratic remainder in equation (26) is dominated by the linear term. Since

$$\zeta'(\rho) \neq 0,$$

it follows that

$$|\rho_W - \rho| \leq \frac{2K_\rho}{|\zeta'(\rho)|} W^{-\sigma_U - 1}. \quad (29)$$

This proves the stated order. $\square$

Remark 1 (Status of the Exponent). *The proposition explains the observed exponent under explicit local assumptions: (1) the reference zero is simple; (2) a corresponding approximation zero remains in a fixed local neighborhood; (3) the Euler–Maclaurin remainder estimate holds uniformly in that neighborhood.*

The numerical fits in Section 3 support this local asymptotic picture, but they do not replace these assumptions by a global theorem for all nontrivial zeros. The exponent 3/2 is therefore analytically motivated and locally conditional, rather than inferred solely from regression.

2.8.3 General DRCC Operationalization Principle

The zeta-function case is one concrete instance of a more general operational pattern. The essential DRCC contribution of the present work is not the Euler–Maclaurin formula itself. Rather, it is the explicit separation of finite fragment, reconstruction model, reconstruction width, and accuracy condition for a deterministic approximation process.

Let X be a target mathematical or computational object, and let

$$F_W(X) \quad (30)$$

be a finite approximation or fragment indexed by a numerical depth parameter $W \in \mathbb{N}$. Let

$$\mathcal{M} = \{M_\alpha : \alpha \in A\} \quad (31)$$

be a declared family of admissible reconstruction models. For each model M_α , define a reconstructed approximation

$$\hat{X}_{\alpha, W} = \mathcal{R}_\alpha(F_W(X)), \quad (32)$$

where $\mathcal{R}_\alpha$ is the model-specific reconstruction operator. Let

$$\mathcal{E}_\alpha(X, W) = d_{\mathcal{Y}}(\hat{X}_{\alpha, W}, X) \quad (33)$$

be the reconstruction error in a declared metric space $(\mathcal{Y}, d_{\mathcal{Y}})$.

Definition 3 (Accuracy-Relative Reconstruction Depth). *For a target tolerance $\varepsilon > 0$, the stable accuracy-relative reconstruction depth of model M_α is*

$$W_{\text{stab}}^{(\alpha)}(X, \varepsilon) = \min \{W_0 \in \mathbb{N} : \mathcal{E}_\alpha(X, W) < \varepsilon \text{ for every } W \geq W_0\}, \quad (34)$$

provided that the defining set is nonempty. This is deliberately stronger than the first hitting time $\min\{W : \mathcal{E}_\alpha(X, W) < \varepsilon\}$: the latter only guarantees accuracy at that single W , whereas equation (34) guarantees it for every larger width as well, which is the property actually needed for $W_{\text{stab}}^{(\alpha)}(X, \varepsilon)$ to serve as a usable stability threshold rather than an isolated lucky value. For the concrete zeta case in this manuscript, this eventual-accuracy property is observed numerically to hold (Table 15, monotone decreasing error) but, consistent with the caveat already stated in Section 2.8.1, is not proven for arbitrary W .

This definition separates three quantities that must not be conflated: W_{trunc} (the numerical truncation or iteration depth used in one computation), $W_{\text{stab}}^{(\alpha)}(X, \varepsilon)$ (the smallest depth beyond which the declared accuracy is attained and maintained under model M_α), and ω_{CPR} (the structural active-interface width of the declared reconstruction process). In general, no equality among these quantities is implied.

Proposition 4 (Conditional Operationalization Principle). *Assume that, for a fixed object X , model M_α , and tolerance $\varepsilon > 0$, there exists a finite $W_0 \in \mathbb{N}$ such that*

$$\mathcal{E}_\alpha(X, W) < \varepsilon \quad \text{for every } W \geq W_0. \quad (35)$$

Then the stable accuracy-relative reconstruction depth $W_{\text{stab}}^{(\alpha)}(X, \varepsilon)$ is well-defined and finite, and for every finite $W \geq W_{\text{stab}}^{(\alpha)}(X, \varepsilon)$,

$$0 < W_{\text{stab}}^{(\alpha)}(X, \varepsilon) \leq W < \infty \quad (36)$$

and

$$\mathcal{E}_\alpha(X, W) < \varepsilon. \quad (37)$$

Proof. By assumption, the set $\{W_0 \in \mathbb{N} : \mathcal{E}_\alpha(X, W) < \varepsilon \text{ for every } W \geq W_0\}$ is nonempty; being a nonempty subset of $\mathbb{N}$, it has a least element, so $W_{\text{stab}}^{(\alpha)}(X, \varepsilon)$ exists, is positive, and is finite. For every finite width satisfying $W \geq W_{\text{stab}}^{(\alpha)}(X, \varepsilon)$, the defining property of that least element directly gives $\mathcal{E}_\alpha(X, W) < \varepsilon$, and the inequalities $0 < W_{\text{stab}}^{(\alpha)}(X, \varepsilon) \leq W < \infty$ follow directly from the ordering of W_0 and W in $\mathbb{N}$. $\square$

Remark 2 (Meaning and Limits of the General Principle). *Proposition 4 is a general consistency statement: given the (nontrivial, but modest) assumption that accuracy is eventually attained and maintained, the well-ordering of $\mathbb{N}$ guarantees a well-defined, finite, stable threshold; the mathematical content beyond this existence argument is deliberately kept minimal. It does not prove that a useful reconstruction model exists for every high-dimensional problem, it does not show that the required depth is polynomially bounded, and it does not identify the numerical depth with structural CPR-Width. Its role is only to specify exactly which additional objects must be supplied – a fragment F_W , a model M , an error measure $\mathcal{E}_\alpha$, and an eventually-maintained accuracy condition – before an abstract reconstruction condition such as (4) becomes a measurable computational statement. Whether such objects can be supplied usefully for a given problem is a separate, problem-specific question that this proposition does not answer.*

Definition 4 (Operational DRCC Record). *An operational DRCC record for a deterministic approximation process is the tuple*

$$\mathfrak{D} = \left(X, F_W, \mathcal{M}, \mathcal{R}_\alpha, \mathcal{E}_\alpha, \varepsilon, W_{\text{stab}}^{(\alpha)}, \omega_{\text{CPR}}, T_\alpha \right), \quad (38)$$

where X is the target object; F_W is the finite fragment or approximation; $\mathcal{M}$ is the admissible reconstruction-model family; $\mathcal{R}_\alpha$ is the reconstruction operator; $\mathcal{E}_\alpha$ is the declared error measure; ε is the required tolerance; $W_{\text{stab}}^{(\alpha)}$ is the required numerical reconstruction depth; ω_{CPR} is the structural active-interface width; and T_α is the complete runtime or operation-count model.

The DWZ method instantiates this record, fully and concretely, through

$$X = \zeta(s), \quad F_W(s) = \sum_{n=1}^W n^{-s}, \quad (39)$$

$$\mathcal{R}_{M_1}(F_W) = F_W + \frac{W^{1-s}}{s-1} - \frac{1}{2}W^{-s}, \quad (40)$$

and

$$\mathcal{E}_{M_1}(s, W) = |\zeta_{\text{DWZ}}(s, W) - \zeta(s)| \quad (41)$$

or, for zero reconstruction,

$$\mathcal{E}_{M_1}^{\text{zero}}(\rho_k, W) = \left| s_k^{(M_1)}(W) - \rho_k \right|. \quad (42)$$

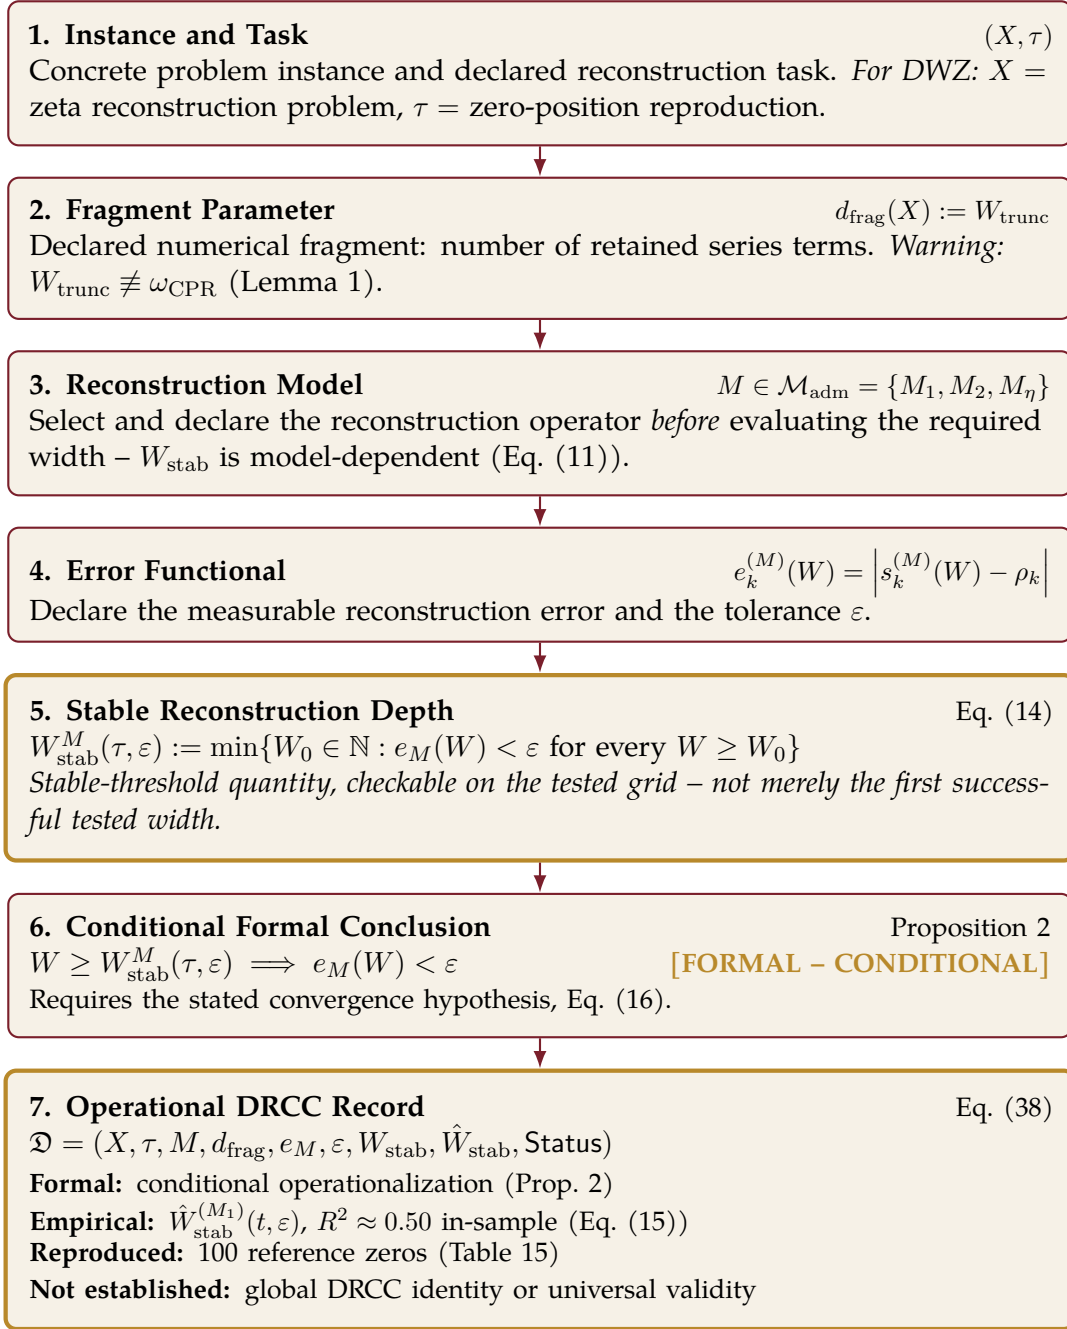
This is the only instance of the record $\mathfrak{D}$ worked out in full in this manuscript. The same operational template could, in principle, be applied to other deterministic iterative approximation problems with an explicit fragment, reconstruction operator, error metric, and declared stopping tolerance – candidates might include truncated integral transforms, finite spectral expansions, or iterative linear-system reconstruction – but no such second instance is worked out here, and we deliberately stop short of claiming that one exists for every, or even any, such problem. For each hypothetical application, the mathematical burden remains entirely problem specific: one would have to independently establish $\mathcal{E}_\alpha(X, W) \rightarrow 0$, derive or estimate $W_{\text{stab}}^{(\alpha)}(X, \varepsilon)$, separately determine ω_{CPR} , and account for the complete computational cost $T_\alpha = T_{\text{build}} + T_{\text{reconstruct}} + T_{\text{verify}} + T_{\text{output}}$ – none of which is automatic, and none of which follows from Proposition 4 alone.

Remark 3 (What Is and Is Not Transferable). *The transferable insight of the DWZ case study is not a new summation formula, and it is not a proof that DRCC generalizes to other high-dimensional problems. What is transferable, and what this manuscript actually delivers, is a worked template of the operational sequence – finite fragment, declared reconstruction model, measurable reconstruction error, accuracy-relative reconstruction depth, and separate structural-width/runtime analysis – carried out completely and honestly for one concrete, nontrivial example. Whether this sequence can be carried out as completely for a different target object X is an open question specific to that object, not something this manuscript establishes. The zeta function is the first, and so far only, detailed numerical realization developed in this manuscript.*

2.9 Verification Strategies

We use four methodologically distinct tests:

1. **Known start:** Newton refinement of equation (7) starting from the known reference position (48 of the first 50 zeros #1–#50 in Table 15, all except #9 and #11; the further 50 zeros #51–#100 in the same table were all treated with a known start, see Section 3).



What the record establishes

What the record does not establish

A reproducible and checkable operationalization for one declared model, on one declared instance and task.

$W_{\text{trunc}} = \omega_{\text{CPR}}$; global convergence for all t ; a faster zeta algorithm; or a result on the Riemann Hypothesis.

Figure 7: OPA-DRCC-05: Operational DRCC Record from an abstract reconstruction condition to a checkable numerical claim. The diagram separates the declared instance and task, fragment parameter, reconstruction model, error functional, stable reconstruction depth, conditional formal conclusion, and evidence status. The record is model-relative and does not identify numerical truncation depth with structural CPR-Width.

2. **Blind search:** scanning $|\zeta_{\text{DWZ}}|$ over a t -range without reference knowledge, Newton refinement of the best candidate (zeros #9, #11).
3. **Hybrid search:** analytic estimate (Riemann-von Mangoldt formula, Eq. 43) as the scan *center*, with a reduced window (zero #12, hardening #13–#41).
4. **Gap verification:** fine-grained scan between two known zeros to search for additional numerical candidates on the chosen grid (not a rigorous completeness proof, see Section 4).

Conceptually, three objects must be distinguished: the *scan minimum* (starting candidate from the one-dimensional scan along $\text{Re}(s) = \frac{1}{2}$), the *zero of the approximation* $s_k^{(M)}(W)$ (result of the complex Newton iteration in the two-dimensional search space), and the *reference zero* ρ_k of ζ itself. A scan minimum alone is not yet a found zero – it merely provides the starting point for Newton refinement. For the hybrid search, the Riemann-von Mangoldt counting function is used:

$$N(T) \approx \frac{T}{2\pi} \ln \frac{T}{2\pi} - \frac{T}{2\pi} + \frac{7}{8}. \quad (43)$$

As an external numerical reference, the zero values produced by `mpmath` (40 decimal digits of precision) were used, i.e. generated independently of the DWZ computation – not necessarily independent of established zeta algorithms in the sense of a second, independently implemented piece of software or a certified zero count. For the first zeros, these values agree with published tables such as [4, 5].

2.10 Detailed Worked Example: Determination of Zero #1

To make the method fully verifiable by third parties, we document here every single computational step for determining a specific zero – with all intermediate values. The first nontrivial zero $\rho_1 = 0.5 + 14.1347251417 \dots i$ serves as the example. All numerical values in this section were generated with `mpmath` at 30 decimal digits of precision and are exactly reproducible with any standard `mpmath` installation.

Figure 8 shows the complete chain of steps at a glance, before the individual steps are documented in detail with numerical values below.

Step 1 – Choice of parameters. We choose the width $W = 2000$ and use $\zeta_{\text{DWZ}}(s, W)$ according to equation (7). A coarse interval $t \in [13.0, 15.0]$ on the critical line $\text{Re}(s) = 0.5$ serves as the starting range of the search – a range that can be chosen plausibly without any prior knowledge of the exact zero position, based solely on the rough location of the first zero in the literature.

Step 2 – Coarse scan. We evaluate $|\zeta_{\text{DWZ}}(0.5 + it, 2000)|$ in steps of $\Delta t = 0.2$ (Table 9). The minimum lies at $t = 14.20$.

Step 3 – Fine scan. To refine the coarse scan, we sample the interval $t \in [14.00, 14.30]$ in steps of $\Delta t = 0.02$ (Table 10). The minimum now lies at $t = 14.14$, with $|\zeta_{\text{DWZ}}| = 0.004177$ – already an order of magnitude smaller than in the coarse scan.

Step 4 – Newton’s method. Starting from the best candidate of the fine scan, $s_0 = 0.5 + 14.14i$, we refine the zero using Newton’s method for complex functions:

$$s_{n+1} = s_n - \frac{f(s_n)}{f'(s_n)}, \quad f(s) := \zeta_{\text{DWZ}}(s, 2000), \quad (44)$$

where $f'(s_n)$ is approximated via a central difference with step size $h = 10^{-6}$: $f'(s_n) \approx (f(s_n+h) - f(s_n-h))/(2h)$. Table 11 shows the complete iteration sequence.




Figure 8: The chain of steps in DWZ zero determination, illustrated for ρ_1 : each arrow represents a fully documented computational step (details and all intermediate values below).

Table 9: Coarse scan, step size 0.2.

t	$ \zeta_{\text{DWZ}}(0.5 + it, 2000) $
13.00	0.791151
13.20	0.669990
13.40	0.540293
13.60	0.402573
13.80	0.257453
14.00	0.105639
14.20	0.052042
14.40	0.214601
14.60	0.380919
14.80	0.549793
15.00	0.719932

Table 10: Fine scan, step size 0.02.

t	$ \zeta_{\text{DWZ}}(0.5 + it, 2000) $
14.000	0.105639
14.020	0.090120
14.040	0.074543
14.060	0.058909
14.080	0.043219
14.100	0.027474
14.120	0.011675
14.140	0.004177
14.160	0.020081
14.180	0.036036
14.200	0.052042
14.220	0.068096
14.240	0.084198
14.260	0.100346
14.280	0.116540
14.300	0.132779

Table 11: Newton iteration starting from $s_0 = 0.5 + 14.14 i$.

n	s_n	$ f(s_n) $
0	$0.500000000 + 14.140000000 i$	4.18×10^{-3}
1	$0.500023320 + 14.134738876 i$	9.09×10^{-6}
2	$0.500012099 + 14.134736526 i$	4.32×10^{-11}
3	$0.500012099 + 14.134736526 i$	9.77×10^{-22}

Already after two iterations, $|f(s_n)|$ has fallen to 10^{-11} ; from $n = 3$ onward, s_n changes only in decimal places far beyond any practical relevance – the method has converged quadratically to the zero of $\zeta_{\text{DWZ}}(\cdot, 2000)$.

Step 5 – Comparison with the reference. The point found,

$$s^* = 0.500012099 + 14.134736526 i,$$

is compared with the independently computed reference $\rho_1 = 0.5 + 14.134725142 i$ (`mp.zetazero(1)`, 30 decimal digits):

$$|s^* - \rho_1| = 1.66 \times 10^{-5}. \quad (45)$$

This residual deviation is *not* an error of Newton’s method (which found exactly the zero of the width-2000 approximation), but the expected discretization deviation of $\zeta_{\text{DWZ}}(s, 2000)$ from $\zeta(s)$ at finite W . It agrees with the empirical power law derived in Section 2: with $C_1^{\text{local}} \approx 1.499$ and $p \approx 1.501$ – a local fit computed exclusively from the W -error values of this worked example (Figure 5), to be distinguished from (a) the systematic table entry $C(t_1) = 1.486$ from Table 13, based on the same W -grid as all other 49 rows of that table, and (b) the value of the global linear regression $C_{\text{reg}}(14.13) = 0.0272 \times 14.13 + 1.043 \approx 1.427$ from equation (15). All three values lie in the range 1.43–1.50 and arise from different fitting procedures (local, on a few W -points; systematic, per zero; global, over all 50 zeros); the differences between them are expected fit sensitivity, not a contradiction, but should not be left unremarked side by side – hence this clarification. With C_1^{local} and p , the approximation (15) predicts an error of $C/W^p = 1.499/2000^{1.501} \approx 1.68 \times 10^{-5}$ – practically identical to the actually observed value 1.66×10^{-5} . This agreement is, however, an in-sample consistency check, not an independent validation: C_1^{local} and p were fitted directly from the error curve of this same zero #1 (Figure 5), so the prediction merely checks whether the fit curve reproduces its own data. An actual test independent of the calibration is provided only by the full Table 15 with 100 different zeros. Increasing W to 10^4 (as used in Table 15 for all 100 zeros) correspondingly reduces the deviation to $\approx 1.5 \times 10^{-6}$ (for the early, small- t entries of that table; for larger t it rises slightly to $\approx 1.3 \times 10^{-5}$, see Table 15).

Reproducibility. The complete computational procedure thus consists of exactly five steps: (1) fix the width W and starting interval, (2) coarse scan of $|\zeta_{\text{DWZ}}|$ on the critical line, (3) fine scan around the coarse-scan minimum, (4) Newton refinement with numerical derivative, (5) comparison with an independent reference or with the theoretical error bound $C(t)/W^p$. Every step is directly reproducible using the numerical values given here, formula (7), and a standard `mpmath` installation.

3 Results

3.1 Comparison of the Repair Variants

3.2 $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$

Table 12: Error $|\zeta_W(s, W) - \zeta(s)|$ for $s = 0.5 + 14.13i$: Variant A vs. Variant B.

W	Variant A (DWZ)	Variant B (Eta)
10	3.85×10^{-2}	8.24×10^{-2}
50	3.34×10^{-3}	2.99×10^{-2}
100	1.18×10^{-3}	2.11×10^{-2}
500	1.05×10^{-4}	9.42×10^{-3}
1000	3.73×10^{-5}	6.66×10^{-3}
5000	3.33×10^{-6}	2.98×10^{-3}




Figure 9: Variant A (DWZ) converges by orders of magnitude faster than Variant B, shown here at $s = 0.5 + 14.13i$ (empirically fitted exponent near $-3/2$ on the critical line for Variant A; Variant B scales as $\mathcal{O}(W^{-\text{Re}(s)})$). Elsewhere in the critical strip, the exponent for Variant A depends on $\text{Re}(s)$.




Figure 10: Left: convergence of the zeta error $\Delta_\zeta(W) = |\zeta_{\text{DWZ}}(s, W) - \zeta(s)|$ for $s = 0.5 + 14.13i$. Right: convergence of the zero stability $\Delta_k(W) = |s_k(W) - \rho|$ for zero #1 – both decreasing as W grows.

Table 13: Fitted $C(t)$ and exponent p per zero, over the first 50 zeros #1–#50 from Table 15 ($n = 50$, previously $n = 8$; Table 15 now spans 100 zeros, but this fit was not re-extended to the additional 50); the functional form of $C(t)$ is still considered preliminary, not conclusively proven (see Conclusion).

k	t	$C(t)$	p
1	14.1347	1.486	1.5000
2	21.0220	1.544	1.5002
3	25.0109	1.523	1.5002
4	30.4249	1.944	1.5000
5	32.9351	1.990	1.5002
6	37.5862	1.615	1.4998
7	40.9187	2.287	1.4999
8	43.3271	1.972	1.5002
9	48.0052	2.547	1.4998
10	49.7738	2.918	1.4998
11	52.9703	1.818	1.4999
12	56.4462	1.992	1.5003
13	59.3470	3.563	1.5003
14	60.8318	3.058	1.4997
15	65.1125	2.367	1.4998
16	67.0798	3.131	1.4999
17	69.5464	2.660	1.5004
18	72.0672	2.023	1.4999
19	75.7047	3.561	1.5005
20	77.1448	4.403	1.4998
21	79.3374	2.512	1.5003
22	82.9104	2.638	1.5004
23	84.7355	3.279	1.5004
24	87.4253	4.042	1.4999
25	88.8091	3.406	1.5002
26	92.4919	2.564	1.4999
27	94.6513	5.345	1.5003
28	95.8706	4.727	1.5002
29	98.8312	2.354	1.5006
30	101.3179	2.673	1.4999
31	103.7255	3.998	1.5003
32	105.4466	4.681	1.4998
33	107.1686	2.993	1.4999
34	111.0295	5.848	1.5007
35	111.8747	6.907	1.5010
36	114.3202	3.745	1.5001
37	116.2267	3.112	1.5003
38	118.7908	2.552	1.5003
39	121.3701	4.427	1.5005
40	122.9468	6.608	1.5007
41	124.2568	4.474	1.5001
42	127.5167	2.704	1.5003

continued on next page

Table 13 (continued)

k	t	$C(t)$	p
43	129.5787	4.609	1.5000
44	131.0877	4.742	1.5008
45	133.4977	4.881	1.5005
46	134.7565	3.993	1.5002
47	138.1160	3.652	1.5003
48	139.7362	6.457	1.4999
49	141.1237	5.874	1.5010
50	143.1118	3.200	1.5007

Regression fit: $C(t) = (0.0272 \pm 0.0040)t + (1.043 \pm 0.375)$; $R^2 \approx 0.50$; p range 1.4997–1.5010.

Bootstrap uncertainty for the $C(t)$ fit. As requested in review, we quantify the fit uncertainty beyond the analytic standard errors above via a case-resampling (pairs) bootstrap on the $n = 50$ $(t, C(t))$ points in Table 13 (10,000 resamples with replacement, ordinary least squares refit on each resample): slope 0.0272, 95% percentile CI [0.0205, 0.0343]; intercept 1.043, 95% CI [0.577, 1.502]; R^2 , 95% CI [0.329, 0.667]. The bootstrap intervals are consistent with, and slightly wider than, the analytic standard errors reported above, and confirm that the slope is reliably positive (the CI excludes zero) while the intercept and R^2 carry substantially more uncertainty – consistent with this fit being described throughout this work as preliminary rather than conclusively established (Section 4).




Figure 11: Left: fitted $C(t)$ values for the first 50 zeros #1–#50 (points) with linear trend (line); the trend is statistically significant but explains only about half of the scatter ($R^2 \approx 0.50$). Right: the resulting estimator function $\hat{W}_{\text{stab}}^{(M_1)}(t, \epsilon)$ (Eq. 15) for three tolerance levels, still based on the $n = 50$ fit (not extended to the additional 50 zeros in Table 15).

Out-of-sample check on zeros #51–100. As rightly requested in review, the $C(t)$ regression above was fitted exclusively on #1–50 and, prior to this check, had never been evaluated against the additional 50 zeros already present in Table 15. We supply this check here. For each zero $k = 51, \dots, 100$, a local value $C_{\text{actual}}(t_k)$ is obtained from the already-tabulated deviation $e_k(W=10^4)$ (Table 15) under the same fixed-exponent convention used to build Table 13, $C_{\text{actual}}(t_k) := e_k(10^4) \cdot (10^4)^{1.5}$, and compared against the out-of-sample prediction $C_{\text{pred}}(t_k) = 0.0272 t_k + 1.043$ from the original $n = 50$ fit – without refitting on the new data. The result: mean $C_{\text{actual}} = 6.32$ against mean $C_{\text{pred}} = 6.26$ (aggregate level in reasonable agreement), but at the level of individual zeros the out-of-sample fit is markedly weaker than the in-sample fit, $R_{\text{os}}^2 \approx 0.08$ (correlation ≈ 0.28) compared to $R^2 \approx 0.50$ in-sample, with mean absolute residual ≈ 2.06 and RMSE ≈ 2.57 . An independent regression fitted directly on #51–100 alone gives slope 0.0283 and intercept 0.896, close to the original 0.0272/1.043 – so the aggregate linear trend direction

and magnitude approximately reproduce out of sample, but $C(t)$ cannot be reliably predicted for an individual, previously unseen zero from t alone. This confirms, with numbers rather than a general caveat, the concern raised in review: the $n = 50$ fit generalizes at the level of the overall trend, not at the level of a single prediction, and $\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$ (equation (15)) should be read accordingly – as an order-of-magnitude guide, not a precise, individually reliable estimate for a specific, previously unseen t .




Figure 12: Error curve $|s_k(W) - \rho|$ for the first zero, with power-law fit ($C = 1.499, p = 1.501$).




Figure 13: $|\zeta_{\text{DWZ}}(0.5 + it)|$ in Cartesian coordinates near $t = 14.1347$ (dashed line): at small W the minimum remains imprecise; as W grows, the location of the minimum approaches the reference ordinate t_1 without reaching it exactly at finite W .

3.3 Blind and Hybrid Zero Search

Table 14: Zeros found blindly (without reference knowledge).

Zero	Found	Reference	Difference
#9	$0.5000025 + 48.0051511i$	$0.5 + 48.0051508812i$	2.55×10^{-6}
#11	$0.4999983 + 52.9703208i$	$0.5 + 52.9703214777i$	1.82×10^{-6}

A direct Newton start from the analytic estimate (Eq. 43) for zero #12 ($T \approx 57.55$) diverged completely (values $> 10^{12}$ after a few steps). The hybrid strategy (estimate as scan center,




Figure 14: Scan of $|\zeta_{\text{DWZ}}(0.5 + it)|$ at $W = 2000$: the candidate for zero #9 lies exactly at the minimum found, at $t \approx 48.005$.

window ± 4), by contrast, worked and required only 61 instead of 220 scan points, at the same absolute deviation (2.0×10^{-6}).

3.4 Hardening and Pattern Hypothesis

The term "hardening" here denotes a stress test of index assignment (cf. the formal assignment rule in Section 2), not a confirmation of exact target indices: as the following results show, the method reliably finds a known zero in the neighborhood, but only in part of the cases finds exactly the one targeted – the name should thus be read in the sense of "neighborhood stress test of index assignment," not as a general assurance.

The hybrid method was repeated for #13–#41 (29 cases): in all 29 cases, the point found agreed, within the numerical tolerance, with a known reference zero, but only in 10 cases with the exact one targeted – the remaining cases hit a neighboring zero. Plausible, but as shown below *not* statistically confirmed in this sample, is the conjecture that this is related to the zero spacing shrinking on average with t ($2\pi / \ln(t/2\pi)$).

A total of 29 zeros were subjected to the hybrid hardening procedure. However, complete case-level records (spacing, outcome, and precision error for each individual case) are available in the present manuscript only for the 20-case statistical subsample used in the point-biserial and logistic-regression analysis below (Figure 15). No complete 29-row reconstruction table is claimed here, and none of the statistics reported below are extrapolated, estimated, or otherwise derived to cover the remaining 9 cases.

A hypothesis derived from this ("smaller spacing $\Rightarrow$ more confusions") was first checked, as a simplification, against the Pearson coefficient for this sample ($n = 29$, $r = 0.055$, no p -value reported). As rightly noted in review, for a binary outcome (targeted zero hit vs. neighbor hit) a point-biserial correlation or logistic regression is the methodologically appropriate tool. We supply this here: for the 20 of 29 hardening runs whose individual values are shown explicitly above (Figure 15; of which 6 have outcome "targeted zero hit" and 14 have outcome "neighbor hit"), the point-biserial correlation between local spacing and outcome – coded as a binary variable with 1 = "targeted zero hit" and 0 = "neighbor hit" (the same coding underlies the logistic regression below) – gives

$$r_{\text{pb}} = -0.189, \quad t(18) = -0.82, \quad p \approx 0.41, \quad 95\% \text{ CI} \approx [-0.58, 0.28], \quad (46)$$

and a logistic regression (outcome $\sim$ spacing) yields a slope coefficient of -0.93 (standard error 1.09; Wald $z = -0.85$, $p \approx 0.39$; 95% CI $\approx [-3.07, 1.21]$). Both confidence intervals comfortably straddle zero, which further illustrates the lack of significance. Both tests are statistically non-significant, confirming the previously reported observation that this sample




Figure 15: Strip plot of the 20 of 29 hardening runs for which individually documented case-level records are available. The remaining 9 were conducted under the same hardening protocol, but their individual case-level pairs are not available in the manuscript and are therefore not included in this plot. Each point is placed at its local zero spacing on the x -axis and its binary outcome class on the y -axis (small vertical jitter is added only to separate points sharing the same x -value within a class; it carries no quantitative meaning). This is the actual relationship entering the statistical tests below: local spacing (x) versus outcome class (targeted zero hit vs. neighbor hit). The descriptive Pearson coefficient $r = 0.055$ was previously reported over all 29 runs; however, only the 20 individually documented runs shown here support a fully reproducible point-biserial and logistic-regression analysis, reported in the text below.

shows no defensible linear relationship – now, however, with an effect size and a p -value rather than an unqualified correlation coefficient alone. Notably, though not to be over-interpreted given the lack of significance, the point estimate points, if anything, in the *opposite* direction from the original hypothesis (larger, not smaller, spacing is weakly associated with more, not fewer, confusions). For the remaining 9 hardening runs, no individually documented (t , outcome) pairs are available in the manuscript; consequently, the Pearson coefficient $r = 0.055$ reported over all 29 cases at the outset cannot be independently recomputed from the manuscript alone – only the tests reported above, based on $n = 20$ (r_{pb} , logistic regression), are fully reproducible from the data shown here. An analysis over all 29 points, like the previously announced increase in sample size, remains an open follow-up step.

Statistical power of this test. As requested in review, we report the statistical power of the $n = 20$ point-biserial test, computed via the standard Fisher z -transformation normal approximation for a two-tailed test at $\alpha = 0.05$. At $n = 20$, the power to detect a conventionally “medium” effect size ($r = 0.3$) is only 0.25; the power to detect an effect as large as the point estimate actually observed here ($r_{pb} = -0.189$) is just 0.12. Reaching 80% power to detect a medium effect ($r = 0.3$) would require $n \approx 85$; detecting the observed effect size itself at 80% power would require $n \approx 218$. The correct reading of this test is therefore not “no relationship exists,” but “this sample size could not have detected a small-to-medium relationship with reasonable probability even if one were present”; the non-significant result is compatible with a genuine absence of a relationship, but equally compatible with an underpowered test missing a real, moderate one. This sharpens, rather than contradicts, the qualitative conclusion already stated above.

3.5 Numerical Reproduction and Reference Check for 100 Zeros

To frame the following table: of the 100 zeros, 98 were treated with a known starting value (Newton’s method from the rounded `mp.zetazero` reference) and only 2 (#9, #11) were found blindly; 29 of these zeros (#13–#41) were additionally hardened via the hybrid method, of which only 10 of the 29 targeted indices were hit exactly (Section 3). The phrase “100 zeros numerically

reproduced" should be read in this sense: predominantly as a local reproduction of already-known reference values, not as an independent discovery or standalone verification of the first 100 zeros.

Table 15: DWZ method: deviation from the reference position for $k = 1, \dots, 100$ (two blocks of 50 shown side by side, k and $k + 50$, to reduce the number of pages this table spans). Here $t_k := \text{Im}(\rho_k)$, and the deviation column is $|\text{DWZ} - \rho_k|$ in units of 10^{-6} .

k	t_k	dev.	k	t_k	dev.
1	14.1347	1.49	51	146.0010	3.75
2	21.0220	1.54	52	147.4228	4.36
3	25.0109	1.52	53	150.0535	7.72
4	30.4249	1.95	54	150.9253	7.95
5	32.9351	1.99	55	153.0247	3.01
6	37.5862	1.62	56	156.1129	4.20
7	40.9187	2.29	57	157.5976	7.53
8	43.3271	1.97	58	158.8500	6.20
9	48.0052	2.55	59	161.1890	3.74
10	49.7738	2.92	60	163.0307	3.30
11	52.9703	1.82	61	165.5371	4.06
12	56.4462	1.99	62	167.1844	5.67
13	59.3470	3.55	63	169.0945	9.17
14	60.8318	3.07	64	169.9120	7.15
15	65.1125	2.37	65	173.4115	4.14
16	67.0798	3.13	66	174.7542	6.18
17	69.5464	2.65	67	176.4414	5.53
18	72.0672	2.03	68	178.3774	5.27
19	75.7047	3.55	69	179.9165	4.28
20	77.1448	4.41	70	182.2071	3.03
21	79.3374	2.51	71	184.8745	10.30
22	82.9104	2.63	72	185.5988	12.50
23	84.7355	3.27	73	187.2289	5.50
24	87.4253	4.05	74	189.4162	3.13
25	88.8091	3.40	75	192.0267	5.78
26	92.4919	2.56	76	193.0797	6.73
27	94.6513	5.33	77	195.2654	5.71
28	95.8706	4.72	78	196.8765	7.52
29	98.8312	2.34	79	198.0153	5.23
30	101.3179	2.68	80	201.2648	4.96
31	103.7255	3.99	81	202.4936	7.56
32	105.4466	4.69	82	204.1897	8.19
33	107.1686	2.99	83	205.3947	6.11
34	111.0295	5.82	84	207.9063	3.84
35	111.8747	6.85	85	209.5765	4.24
36	114.3202	3.74	86	211.6909	5.33
37	116.2267	3.11	87	213.3479	8.53
38	118.7908	2.55	88	214.5470	8.09
39	121.3701	4.41	89	216.1695	4.08
40	122.9468	6.57	90	219.0676	4.42
41	124.2568	4.47	91	220.7149	12.50
42	127.5167	2.70	92	221.4307	11.40
43	129.5787	4.61	93	224.0070	7.05
44	131.0877	4.71	94	224.9833	6.49
45	133.4977	4.86	95	227.4214	3.68
46	134.7565	3.99	96	229.3374	5.14
47	138.1160	3.64	97	231.2502	13.00
48	139.7362	6.46	98	231.9872	13.20
49	141.1237	5.83	99	233.6934	4.51
50	143.1118	3.18	100	236.5242	4.93

Min. / mean / median / max. ($n = 100$, units 10^{-6}): 1.49 / 4.87 / 4.32 / 13.18

Continued on next page

Table 15 (continued; $t_k = \text{Im}(\rho_k)$, deviation in units of 10^{-6})

k	t_k	dev.	k	t_k	dev.
Standard deviation ($n = 100$, units 10^{-6}): 2.59					

Summary statistics. In addition to the minimum, mean, and maximum, the median and standard deviation are computed directly from the 100 deviations reported in Table 15. These quantities provide measures of central tendency and dispersion for the complete numerical sample.




Figure 16: Absolute position deviation $e_k^{(M_1)}(W=10^4)$ for all 100 zeros: between 1.5×10^{-6} and 1.3×10^{-5} , with an increasing tendency for larger k (Pearson correlation between k and e_k : $r \approx 0.65$, $n = 100$, no formal significance test – consistent with $e_k^{(M_1)}(W) \approx C(t_k)W^{-p}$ at fixed W and growing t_k).

Runtime against an established reference method (B_{workflow}). For all 100 zeros, the actual runtime was measured in addition to the position deviation and compared against an established reference method: the zeros `mp.zetazero(k)` implemented in `mpmath`, which are determined internally via the Riemann-Siegel formula (evaluation of $\vartheta(t)$ and $Z(t)$ followed by sign-change isolation) and serve throughout this work as the independent reference position ρ_k . For $k = 1, \dots, 100$, the cumulative runtime (30-digit `mpmath` arithmetic, fixed width $W = 10^4$, Newton’s method with known start as in Section 3, identical protocol to Table 15) was

$$T_{\text{DWZ}}(k=1..100) = 135.03 \text{ s}, \quad T_{\text{R.-S.}}(k=1..100) = 4.55 \text{ s}, \quad \frac{T_{\text{DWZ}}}{T_{\text{R.-S.}}} \approx 29.7 \quad (B_{\text{workflow}}).$$

This factor ≈ 30 is the B_{workflow} figure; it is not the same benchmark as the B_{eval} figure of Test C (Section 2.5, roughly two to three orders of magnitude), and the two should not be quoted interchangeably as “the” DWZ/Riemann-Siegel runtime factor. The DWZ position deviation for these 100 zeros ranges from 1.49×10^{-6} to 1.32×10^{-5} (mean 4.87×10^{-6} , Table 15), while `mp.zetazero` at 30-digit working precision achieves an uncertainty many orders of magnitude smaller and is treated throughout as the reference value (not as a competitor with its own residual uncertainty). Riemann-Siegel is thus, for these 100 zeros, both faster (factor ≈ 30) and more accurate – a result that confirms the efficiency gap already reported in Test C (Section 2.5), here however with a smaller factor: Test C measured the extreme case at much larger t (up to $t \approx 1001$) with a width adaptively optimized for $\varepsilon = 10^{-6}$, whereas the constant width $W = 10^4$ used here already suffices for the smaller t of this range ($t \in [14, 237]$) – the factor 30 is therefore not a contradiction of the factor 100–1000 from Test C, but the result of a deliberately different,

here fixed, choice of width. Both measurements are snapshots on the hardware used and are not to be understood as absolute constants, but their relative order of magnitude is consistent across both protocols: Riemann-Siegel is clearly superior to the DWZ method in both runtime and accuracy. It should also be noted that T_{DWZ} and $T_{\text{R.-S.}}$ do not measure exactly the same sub-task here: T_{DWZ} measures Newton *refinement* from an already-known starting value at fixed width $W = 10^4$, whereas `mp.zetazero(k)` determines the k -th zero entirely from scratch, including its own isolation and search steps. The measured factor ≈ 30 is thus a comparison of two different, each practically relevant, tasks (refinement from a known start vs. a full search), not a comparison of identical workloads; a strict algorithm comparison would additionally need to measure separately (a) pure function evaluation at the same s and (b) zero search from identical starting information with an identical stopping criterion – neither is carried out here.

3.6 Gap Verification

For nine neighboring pairs of zeros, the magnitude $|\zeta_{\text{DWZ}}(\frac{1}{2} + it, W)|$ at $W = 2000$ was sampled finely between the known positions (Table 16). In none of the nine intervals was an additional candidate with a sufficiently small ζ_{DWZ} magnitude observed on the chosen grid. This numerical check only shows that, on the critical line and within the grid used, no further candidate became visible – it is not a rigorous completeness proof: a zero between two grid points, a minimum shifted by finite W , or zeros away from the critical line are not thereby ruled out. Such a proof would require, for instance, an argument-principle count or a certified Turing-style zero count.

Table 16: Tested gaps between neighboring zeros.

Pair	Spacing	Result
#1–#2	6.89	no additional zero
#2–#3	3.99	no additional zero
#3–#4	5.41	no additional zero
#4–#5	2.51	no additional zero
#5–#6	4.65	no additional zero
#6–#7	3.33	no additional zero
#7–#8	2.41	no additional zero
#22–#23	1.83	no additional zero
#23–#24	2.69	no additional zero

4 Discussion

4.1 Scope and Limitations

The disclaimers in this work are scattered by necessity – each one is attached to the specific claim it qualifies – but their cumulative content is short and is collected here once, in full, as a single reference point.

What this work is not. It is not a contribution to the proof or disproof of the Riemann Hypothesis; the statement of RH (equation (3)) is not assumed, and the numerical experiments are restricted to zeros already known to lie on the critical line (Section 1.2). It is not a new mathematical discovery in Euler-Maclaurin summation, which is a classical, eighteenth-century technique (Section 1.2, “On methodological novelty”). It is not a faster or more accurate way to compute $\zeta(s)$ than the established Riemann-Siegel formula, which this work’s own measurements show to be faster by roughly two to three orders of magnitude (Section 2.5, Test C). It is not a completeness proof for the zeros found or the gaps tested: the blind and hybrid searches (Section 3) and the gap verification (Section 4) each operate on a fixed numerical grid and




Figure 17: Critical strip $0 < \text{Re}(s) < 1$ (gray) with the first 50 of the now 100 zeros numerically reproduced by the DWZ method (Table 15; #51–#100 are not additionally plotted here for legibility – each point is individually labeled – but are listed in full in Table 15): blue = used as the known reference position for Newton refinement (48 of 50), red = found blindly, without prior knowledge of the exact position (2, #9 and #11) – 4 % of the 50 shown. For all 50, the result agrees with the reference to within $1.5\text{--}6.8 \times 10^{-6}$, regardless of the method used to find it (no point is unaccounted for); for #51–#100 the deviation ranges from 1.5×10^{-6} to 1.3×10^{-5} (Table 15). The additional hardening on 29 of these zeros (#13–#41, already included in the 100, Section 3) is not shown here. The figure distinguishes the numerical protocol used for each point (reference-start Newton refinement vs. blind scan) and does not represent an independent discovery of all plotted zeros.




Figure 18: Two examples: largest tested spacing (#1–#2, left, spacing 6.89) and narrowest tested spacing (#22–#23, right, spacing 1.83) – in both cases only a single hump with one minimum at each edge, no additional candidate on the grid used here (not a proof of the absence of a zero, see text above).

cannot certify the absence of unfound zeros or additional candidates. It is not a claim that the numerical truncation width W_{trunc} equals the structural CPR-Width of the DRCC framework; Lemma 1 shows the opposite for the natural sequential model. It is not a claim that the DRCC operationalization template (Section 2.8.3) transfers to other problems; only one instance of that template, the zeta-function case, is worked out in this manuscript.

What this work is. It is a reproducible case study demonstrating, end to end and with every approximation made explicit, how an abstract reconstruction-width stability condition can be turned into a checkable number on one concrete, nontrivial numerical problem, including a formal proof that runtime length and active reconstruction state width are logically separate quantities (Lemma 1) and an analytical account of the observed convergence rate (Section 2.8.2).

The remaining paragraphs of this section discuss specific results in more depth; they restate none of the above scope statements beyond what is needed for the point being made.

4.2 Detailed Discussion

The numerical reproduction of known zeros is a *validity demonstration of the method*, not a contribution to the Riemann Hypothesis itself. The reference values come from `mpmath` (Section 2); the DWZ method confirms that a controlled, finite width reconstruction is able to reproduce these positions with high, consistent accuracy – even where zeros lie close together.

Of the first 50 zeros treated in Table 15, 2 (#9, #11) were found blindly, without prior knowledge of the exact position (4 % of these 50); the remaining 48 were refined by Newton’s method starting from the known reference position. The further 50 zeros #51–#100 in the same table were all treated with a known start (Section 3); no separate blind-search test is available for them. The separate hardening on 29 of these zeros (#13–#41, already included in the 100) via the hybrid method at the same time reveals a real limitation of the simple hybrid search: as the zero spacing shrinks, the risk increases of hitting a neighboring zero instead of the one targeted. This limitation concerns *index assignment*, not the existence of the zeros found – all 29 hits agreed, within the stated numerical tolerance, with known reference zeros.

4.3 Why DRCC, Given That an Analytic Reconstruction Formula Already Converges

A direct objection deserves a direct answer. The Euler-Maclaurin correction (7) is a classical, self-contained analytic tool: it converges on its own, independently of any DRCC vocabulary, and Section 3 confirms that it does. This raises the question this subsection answers explicitly: *if the analytic formula already solves the numerical problem, what work is DRCC actually doing?*

The answer is not that DRCC improves the Euler-Maclaurin formula mathematically – it does not, and no such claim is made anywhere in this manuscript. The answer is a division of labor:

Euler-Maclaurin is the reconstruction tool; DRCC is the finite control framework.

Euler-Maclaurin supplies the analytic reconstruction operator – the rule that turns a finite partial sum into an estimate of $\zeta(s)$, Equation (7). DRCC supplies everything the formula itself does not decide: which finite fragment is actually processed; which reconstruction variant among M_1, M_2, M_η is admissible for the declared task; when the declared accuracy has been reached; when further computation adds nothing beyond what is already recoverable from the retained fragment and correction terms; and which finite record is sufficient for the result to be reproduced by someone else. The building blocks for this division are already present in this manuscript – the stability condition (4), the stable reconstruction depth $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ of Equation (14), the separation of state width from truncation depth (Lemma 1), and the Operational DRCC Record (38) – but they have, until now, not been assembled explicitly into this single answer.

DRCC does not compute infinity. The zeta function itself has an infinite analytic definition (1), and the classical convergence statement $\lim_{W \rightarrow \infty} \zeta_{\text{DWZ}}(s, W) = \zeta(s)$ is a fact about that infinite limit. Nothing in the DRCC framework executes this limit. For a declared task

$$\tau = (s, \varepsilon, M),$$

what DRCC actually governs is a finite, terminating protocol. Write

$$W_{\text{DRCC}}(\tau) := W_{\text{stab}}^M(\tau, \varepsilon)$$

for the same stable-threshold quantity already defined in Equation (14) – the notation $W_{\text{DRCC}}(\tau)$ is introduced here only to make its role explicit under this framing; it is not a second, independent definition. Once

$$W \geq W_{\text{DRCC}}(\tau),$$

the process is declared complete: every index $n > W_{\text{DRCC}}(\tau)$ is *task-redundant* in the precise sense that its individual contribution is already represented, for the declared tolerance ε , by the retained fragment together with the Euler-Maclaurin correction terms. The actually executed reconstruction range is the finite set

$$\mathcal{I}_{\text{DRCC}}(\tau) = \{1, \dots, W_{\text{DRCC}}(\tau)\},$$

never the full infinite index set. This is the connection to the finite-control reading of DRCC referenced in [7]: DRCC does not compute or approximate infinity – it selects a finite, non-redundant reconstruction sufficient for the declared task, and stops there.

Figure 19 answers the process question: what is actually executed, from the infinite object to the finite output. The converse question – what does Euler-Maclaurin alone provide, and what is added exclusively by DRCC? – is answered directly by Table 17.

Reusing the runtime decomposition of the Operational DRCC Record (Section 2.8.3), the honest runtime statement for this case study is

$$T_{\text{DWZ}} = T_{\text{EM}} + T_{\text{control}}, \quad T_{\text{control}} \geq 0, \quad \text{hence} \quad T_{\text{DWZ}} \geq T_{\text{EM}}.$$

This is not a measured gap: in the implementation used here, declaring a model, checking a tolerance, and assembling a record are computationally negligible next to evaluating the Euler-Maclaurin sum itself, so the measured T_{DWZ} in Table 5 already reflects, to good approximation,




Figure 19: OPA-DRCC-06: The infinite analytic object $\zeta(s)$ is never processed as such. A declared finite task selects an analytic reconstruction operator (Euler-Maclaurin) and a finite control layer (DRCC), which together determine the smallest stable finite width and the finite, reproducible output actually computed.

Table 17: Euler-Maclaurin alone versus Euler-Maclaurin under DRCC control (the DWZ method). DRCC does not alter or accelerate the analytic formula; it adds a declared, checkable control and documentation layer around it.

Aspect	Euler-Maclaurin alone	Euler-Maclaurin under DRCC control (DWZ)
Termination	No declared stopping point; convergence is the asymptotic statement $W \rightarrow \infty$	Explicit stopping rule $W \geq W_{\text{DRCC}}(\tau) = W_{\text{stab}}^M(\tau, \varepsilon)$ (Eq. (14))
Tolerance	Not declared in advance	ε fixed before computation and checked against $e_M(W)$ (Proposition 2)
Model admissibility	One fixed formula, no declared alternative	Declared admissible set $\mathcal{M}_{\text{adm}} = \{M_1, M_2, M_\eta\}$, selected before evaluating the required width
Resource classification	No distinction between truncation depth and structural state width	$W_{\text{trunc}} \neq \omega_{\text{CPR}}$ enforced (Lemma 1)
Reproducibility record	None declared as part of the method	Operational DRCC Record $\mathfrak{D}$ (Eq. (38))
Runtime	$T_{\text{EM}} := T_{\text{reconstruct}}$	$T_{\text{DWZ}} = T_\alpha = T_{\text{reconstruct}} + T_{\text{control}}$, $T_{\text{control}} := T_{\text{build}} + T_{\text{verify}} + T_{\text{output}} \geq 0$
Comparison with Riemann-Siegel	Not separately benchmarked in this manuscript	Both remain markedly slower than Riemann-Siegel ($\approx 30\times$, Table 5); DRCC neither closes nor claims to close this gap

T_{EM} alone. The inequality is stated for conceptual completeness – DRCC is not claimed to make the computation faster, only better specified – not because a measurable runtime cost of the control layer was observed.

Euler-Maclaurin computes; DRCC controls, terminates, separates, and records.

Euler-Maclaurin determines how the omitted analytic tail is reconstructed. DRCC determines which finite fragment is retained, which reconstruction model is admissible, when the declared accuracy has been reached, and when further computation becomes redundant. Thus, the analytic formula performs the reconstruction, whereas DRCC controls the finite reconstruction process.

To be precise about what this does and does not claim: this is a statement about the *role* DRCC plays in this manuscript’s construction, not a claim that DRCC is mathematically necessary to derive or use the Euler-Maclaurin correction, and not a claim that no other finite-control vocabulary could describe the same stopping rule. The Euler-Maclaurin formula would converge with or without being described this way. What the DRCC framing adds is a declared, checkable protocol – fragment, model, tolerance, stopping condition, record – for exactly when and why the infinite analytic definition can be replaced by a finite computation, rather than leaving that decision implicit or ad hoc.

5 DRCC/CPR-Width Application Assessment of the DWZ Method

5.1 Purpose of the Application Assessment

The preceding sections establish, piece by piece, how the DWZ method operationalizes the abstract DRCC stability condition (4): the traceability chain and equation register of Section 1.4, the formal instantiation of Section 2.8.1, the resource separation of Lemma 1, and the role division of Section 4.3. This section does not repeat those derivations, and it does not restate the scope statements already collected in Section 4.1. It adds only what is not stated in full anywhere else in this manuscript: an explicit list of the boundary conditions under which the operationalization holds.

5.2 Mapping and Resource Separation, in Brief

For orientation, the mapping used throughout this manuscript is

$$\underbrace{0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq d_{\text{frag}}(X) < \infty}_{\text{Eq. (4)}} \rightsquigarrow \underbrace{d_{\text{frag}}(X) := W_{\text{trunc}}}_{\text{Eq. (6)}} \rightsquigarrow \underbrace{W_{\text{stab}}^{(M_1)}(t, \varepsilon)}_{\text{Eq. (14)}} \rightsquigarrow \underbrace{W \geq W_{\text{stab}}^{(M_1)}(t, \varepsilon) \Rightarrow e_k^{(M_1)}(W) < \varepsilon}_{\text{Prop. 2, conditional}}$$

with the three width-like quantities of this manuscript kept explicitly distinct,

$$W_{\text{trunc}} \not\equiv W_{\text{stab}}^{(M_1)}(t, \varepsilon) \not\equiv \omega_{\text{CPR}}.$$

Full detail and traceability are in Table 2, Figures 1 and 3, and Lemma 1; none of it is reproduced here.

5.3 Boundary-Condition Assessment

The operationalization established in this manuscript holds only within stated boundaries, which this subsection collects explicitly.

- **Convergence hypothesis.** Proposition 2 requires $e_k^{(M_1)}(W) \rightarrow 0$ as $W \rightarrow \infty$ (Eq. (16)) to be assumed, not independently proven for every s in the critical strip.
- **Critical-line qualification.** The observed $W^{-3/2}$ error rate and its analytic origin (Section 2.8.2) are derived and verified on the critical line; away from it, the order depends on $\text{Re}(s)$ and is not separately derived here.
- **Empirical estimator scope.** The estimator $\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$ of Equation (15) is fitted on the first 50 of the 100 reproduced zeros, explains only about half the observed scatter ($R^2 \approx 0.50$), and its out-of-sample performance on zeros #51–100 is markedly weaker ($R_{\text{os}}^2 \approx 0.08$) than its in-sample fit.
- **Sample size for the pattern hypothesis.** The point-biserial test on the $n = 20$ hardening subset (Eq. (46)) is statistically non-significant and underpowered; roughly 85 observations would be required for 80 % power to detect an effect of $r \approx 0.3$, so the available sample supports neither confirming nor ruling out the proposed spacing-confusion relationship.
- **Instance scope.** The Operational DRCC Record $\mathfrak{D}$ (Eq. (38)) is worked out completely for one instance, the zeta reconstruction problem; no second instance is carried out in this manuscript (Remark 3).

5.4 Net Assessment

On accuracy, the operationalization succeeds within its stated scope (naive fragment diverges, M_1 converges, 100 known zeros reproduced to 1.5×10^{-6} – 1.3×10^{-5} , Table 15). On runtime, the honest result is unfavorable and is not softened here: M_2 needs roughly $11 \times$ fewer terms than M_1 (Eq. (11)), M_{K5} a further 100 – $400 \times$ fewer, and both remain markedly slower than Riemann-Siegel (Table 5, $\approx 30 \times$ for M_1). Table 17 already lays out this division in full; DRCC does not close the performance gap and is not intended to.

Combined with Section 4.1, the resulting application-level judgment is: the DWZ case study establishes a concretization of the abstract DRCC condition (4) that is formally stated, checkable on the tested grid, and conditional (Proposition 2), together with a proven separation of truncation depth from structural CPR-Width (Lemma 1) and one fully worked Operational DRCC Record (38). It does *not* establish $W_{\text{trunc}} = \omega_{\text{CPR}}$, a faster zeta algorithm, an unconditional convergence proof, a statistically defensible spacing-confusion effect, anything about the Riemann Hypothesis, or that the general template of Section 2.8.3 transfers beyond this one instance (Remark 3) – the same non-claims already stated in Section 4.1, not repeated in expanded form here.

6 Conclusion and Outlook

This work set out to operationalize the abstract DRCC stability condition $0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq d_{\text{frag}}(X) < \infty$ for one concrete, nontrivial numerical problem, end to end and with every approximation made explicit; the Euler-Maclaurin correction is the specific instrument used to carry this out, not the point of the exercise. Carrying out that operationalization required first showing that the naive width definition fails in the critical strip (divergent), then resolving this with the first-order Euler-Maclaurin correction (the DWZ method), which, at the tested point

$s = 0.5 + 14.13i$ and the width values examined there, is considerably more efficient than the eta-based alternative (Table 12) – no general efficiency theorem over the entire critical strip is derived from this. For the zeta function, this yields a zeta-specific, model-dependent concretization of the originally abstract DRCC inequality $0 < d_{\text{rec}}^{\text{DRCC}}(X) \leq W < \infty$ (Equations (14)–(15), Proposition 2), *without* thereby showing or claiming $W_{\text{trunc}} = \omega_{\text{CPR}}$ or $W_{\text{stab}}^{(M_1)}(t, \varepsilon) = d_{\text{rec}}^{\text{DRCC}}(X)$ – the abstract, structural DRCC/CPR condition itself thus remains unchanged in its abstractness, and the general operational template of Section 2.8.3 is, at present, worked out completely for this one instance only. Open points for future work:

- The fit of $C(t)$ still rests on the first 50 zeros from Table 15 ($n = 50$ instead of $n = 8$, Table 13). The linear trend remains statistically significant but explains only about half of the scatter ($R^2 \approx 0.50$); square-root, logarithmic, and pure power-law forms in t do not perform better. The remaining scatter is therefore not a question of the wrong functional form in t alone, but points to at least one further explanatory quantity not captured here (plausible candidates include the local zero spacing $2\pi/\ln(t/2\pi)$ or properties of the Riemann-Siegel theta function). Table 15 itself has since been extended to 100 zeros (plain position deviation at fixed $W = 10^4$, no new $C(t)/p$ fit); extending the $C(t)$ fit to these additional 50 zeros and incorporating further covariates remain a separate, open step. It should also be examined whether the scaling observed here (Table 13, $t \lesssim 143$) also holds outside the tested numerical range ($t \approx 1000$ – 1123).
- Adaptive scan window width (proportional to $2\pi/\ln(t/2\pi)$), to avoid the index confusion observed during hardening.
- Connecting $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$ to the Riemann-Siegel formula for very large $\text{Im}(s)$.
- A stricter analytic error bound for the truncated-zero-sum reconstruction error of the Chebyshev prime-counting function, defined here for completeness as $E_p(W) := |\hat{\psi}_W(x) - \psi(x)|$, where $\psi(x)$ is the Chebyshev function and $\hat{\psi}_W(x)$ its reconstruction from the first W nontrivial zeros via the classical explicit formula [1, 3] – this quantity is introduced here only as a named target for future work and is not otherwise used, derived, or numerically evaluated in this manuscript.
- **Transferability to other Dirichlet series.** The construction used here (Euler-Maclaurin correction of a partial sum, a formally instantiated CPR state width over deterministically reconstructible terms) does not rely anywhere on properties specific to $\zeta(s) = \sum n^{-s}$. It should in principle carry over to other Dirichlet series $\sum a_n n^{-s}$ with sufficiently regular coefficients a_n – for instance Dirichlet L -functions or the eta function already used in Section 2 – provided an analogous Euler-Maclaurin correction can be derived for the series in question. This is stated here as a plausible direction for extension, not as a result: neither the convergence order nor the numerical efficiency of such a transfer has been tested in this work for any other series.

Appendix A: List of Symbols

Symbol	Meaning	First Use
s	complex variable, $s = \sigma + it$	§1.1
t	imaginary part of s (“height”)	§1.1
σ	real part of s , $\sigma = \text{Re}(s)$	§1.1

Continued on next page

Appendix A: List of Symbols (continued)

Symbol	Meaning	First Use
$\zeta(s)$	Riemann zeta function	Eq. (1)
ρ, ρ_k	nontrivial zero resp. k -th reference zero, $\rho_k = \frac{1}{2} + it_k$ (<code>mp.zetazero(k)</code>)	Eq. (3), §2.8
RH	Riemann Hypothesis (open conjecture, not proven here)	Eq. (3)
X	abstract DRCC object (here: the zeta function with its infinite sum)	§1.2
$d_{\text{rec}}^{\text{DRCC}}(X), d_{\text{frag}}(X)$	reconstruction resp. fragmentation degree of the abstract DRCC stability condition [7]	Eq. (4)
W, W_{trunc}	width resp. numerical truncation depth (number of summed series terms)	§1.2, Eq. (5)
$\zeta_W(s, W)$	naive partial sum $\sum_{n=1}^W n^{-s}$	Eq. (5)
$\zeta_{\text{DWZ}}(s, W)$	DWZ function, Variant A, model M_1 (1st-order Euler-Maclaurin)	Eq. (7)
$\eta_W(s, W), \zeta_W^\eta(s, W)$	Dirichlet eta partial sum resp. Variant B	§2.3
$\zeta_{\text{DWZ}}^{(2)}(s, W)$	Variant C, model M_2 (2nd-order Euler-Maclaurin)	Eq. (9)
B_2	Bernoulli number, $B_2 = 1/6$	§2.4
M, M_1, M_2	reconstruction model in general resp. Variant A / Variant C	§2.4–2.5
ω	summation order (Test A, order invariance)	§2.5
ω_{CPR}	CPR-Width (<i>Controlled Partial Reconstruction Width</i> , Hesamiy [7], external, not identified with W_{trunc})	§1.2, §2.5
$\omega_{\text{rec}}(P, \pi, \tau; \mathfrak{M})$	formally defined CPR state width (Definition 2.1, Hesamiy [7]); proven here to be $\omega_{\text{rec}} = O(1)$ for the sequential zeta accumulation, independent of W_{trunc} ; not a runtime statement	Lemma 1
ε_{rec}	CPR-Width tolerance parameter (Hesamiy [7], external)	Table 1
ε	accuracy/tolerance threshold	Eq. (10)
$W_{\min}(M; s, \varepsilon)$	exact minimum over $\mathbb{N}$ (definition, Eq. (10)); the values actually reported in Table 4 are W_{test} , the smallest width found on a finite grid, not this exact minimum itself	Eq. (10)
$W_{\min}^*(s, \varepsilon)$	minimum of $W_{\min}(M; s, \varepsilon)$ over admissible models $\mathfrak{M}_{\text{adm}}$	Eq. (12)
$W_{\text{stab}}^{(M_1)}(t, \varepsilon)$	stable-threshold reconstruction width (definition; numerically evaluated only at tested (t_k, ε) pairs)	Eq. (14)
$\hat{W}_{\text{stab}}^{(M_1)}(t, \varepsilon)$	empirical estimator for $W_{\text{stab}}^{(M_1)}(t, \varepsilon)$	Eq. (15)
$C(t), p$	fit coefficient resp. exponent of the empirical power law $e_k(W) \approx C(t_k)W^{-p}$	Eq. (15), Table 13
$s_k^{(M)}(W)$	zero numerically determined by model M and width W , assigned to ρ_k	§2.8
$e_k^{(M)}(W)$	position error $ s_k^{(M)}(W) - \rho_k $	§2.8

Continued on next page

Appendix A: List of Symbols (continued)

Symbol	Meaning	First Use
$E_\zeta(s, W)$	function-value error $ \zeta_M(s, W) - \zeta(s) $	§2.8
$N(T)$	Riemann-von Mangoldt counting function (analytic height estimate)	Eq. (43)
r	Pearson correlation coefficient (pattern-hypothesis test, hardening)	§3.4
$E_p(W)$	$:= \hat{\psi}_W(x) - \psi(x) $, reconstruction error of the Chebyshev function from W zeros via the explicit formula; a named target for future work only, not derived, used, or numerically evaluated in this manuscript	§6
$\mathfrak{D}$	operational DRCC record (tuple of fragment, model family, reconstruction operator, error measure, tolerance, depth, CPR-width, runtime model); instantiated for the zeta case only	Def. 4, Eq. (38)

References

- [1] B. Riemann, *Über die Anzahl der Primzahlen unter einer gegebenen Grösse*, Monatsberichte der Berliner Akademie, 1859.
- [2] E. C. Titchmarsh, *The Theory of the Riemann Zeta-Function*, 2nd ed., rev. D. R. Heath-Brown, Oxford University Press, 1986.
- [3] H. M. Edwards, *Riemann's Zeta Function*, Academic Press, 1974.
- [4] A. M. Turing, *Some calculations of the Riemann zeta-function*, Proc. London Math. Soc. (3) 3, 1953, pp. 99–117.
- [5] A. M. Odlyzko, *Tables of zeros of the Riemann zeta function*, available at https://www.dtc.umn.edu/~odlyzko/zeta_tables/.
- [6] F. Johansson et al., *mpmath: a Python library for arbitrary-precision floating-point arithmetic*, <http://mpmath.org>.
- [7] R. Hesamiy, *DRCC: Dimensional Reduction via Controlled Combinatorics*, Zenodo, Version V1.0, 2026, doi: <https://doi.org/10.5281/zenodo.20499433>; see also [arXiv:260604.000001](https://arxiv.org/abs/260604.000001) (posted there under the title *DRCC: Dimensionality Reduction as Complexity Collapse – A Structural Framework for Controlled Reconstruction in Complex Search Spaces*). (This reference provides the theoretical motivation for the present work; it is cited here under its Zenodo title and DOI as the primary, permanent identifier.)
- [8] M. Rubinstein, *Computational methods and experiments in analytic number theory*, in: Recent Perspectives in Random Matrix Theory and Number Theory, London Math. Soc. Lecture Note Series 322, Cambridge University Press, 2005, pp. 425–506.