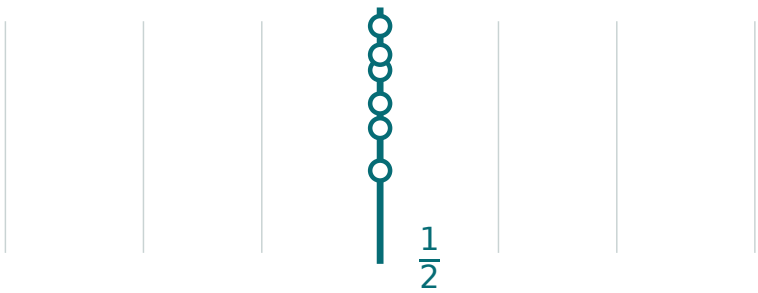
ADRIAN DUNKLEY

The Riemann hypothesis

A visual guide to Alessandro Masullo’s proposed proof



A reading companion to The Music of the Primes

November 21, 2026

adriandunkley.net

Which paper are we reading?

Alessandro Maria Masullo wrote The Music of the Primes. Ravi Vakil released the 940-page draft on November 20, 2026. Masullo proposes a result for unitary Hecke characters over number fields. It includes the ordinary Riemann hypothesis.

The draft contains substantial calculations and unfinished steps. Vakil later reported that readers had identified problems. This companion follows the argument as proposed, including the deductions sketched in Chapter 1 where Chapter 13 is unfinished. The diagrams and worked examples are original teaching aids.

Start with the primes 3
Primes, complex numbers, number fields and characters.

Build the two descriptions 6
Chapters 1–8. What the operators act on, and what repeats.

Compute the traces 15
Chapters 9–12. The strip, the shells and the geometric identity.

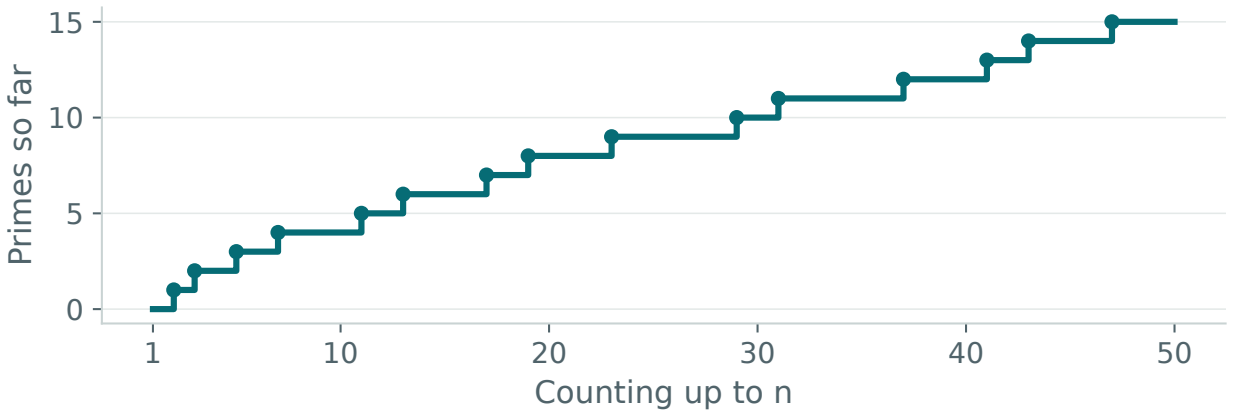
Follow the proposed ending 21
Chapter 13 read alongside §§1.4 and 1.8. Test the reasoning at a hypothetical zero.

Read the manuscript critically 25
Unfinished steps, two checks of p. 29, Chapter 14 and the appendices.

Open Masullo’s original PDF Read Vakil’s announcement and update

How can a sum contain the primes?

Try counting primes up to 10. The list is 2, 3, 5, 7. The count reaches four and then holds steady through 8, 9 and 10. The next step occurs at 11. How much can the staircase wander from its average rise?



Calculated from the integers 1 to 50. Each dot marks a prime. The final step is at 47, so there are 15 primes up to 50.

From factors to a sum

A prime is a whole number greater than 1 whose only positive divisors are 1 and itself. Every integer greater than 1 has a unique prime factorization. For example, $12 = 2^2 \times 3$. Order does not matter. This gives a way to list every positive integer once in a single multiplication.

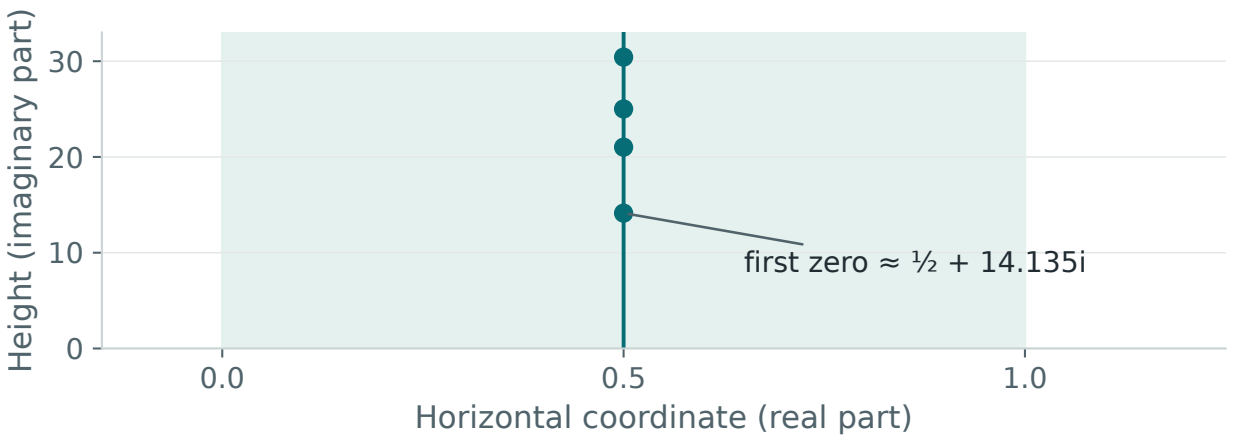
Try the multiplication at s = 2

Write one bracket for each prime. The bracket for 2 is $1 + 1/2^2 + 1/4^2 + \ldots$. The one for 3 starts $1 + 1/3^2 + \ldots$. Pick $1/4^2$ and $1/3^2$, then 1 from every other bracket. Their product is $1/12^2$. Expanding all brackets gives $\zeta(2) = 1 + 1/2^2 + 1/3^2 + \ldots$. The Greek letter ζ is pronounced “zeta.” This Euler product is the first bridge from primes to functions.

NIST · Euler product Clay · the hypothesis

Why draw a vertical line?

A function is a rule with an input and an output. A zero is an input whose output equals zero. For x − 3 the only zero is 3. Zeta takes complex inputs, so locating its zeros requires a map with two coordinates.



Four numerically computed zeta zeros, rounded for display. The shaded strip runs from real part 0 to 1. A graph of four zeros cannot settle a statement about all of them.

Two coordinates make one number

Write s = a + bi. Here $i^2 = -1$. Plot a horizontally and b vertically. The point ½ + 14.135i lies directly above ½. The statement Re(s) = ½ fixes the horizontal position and leaves the height free. The first positive-height zero is near this point.

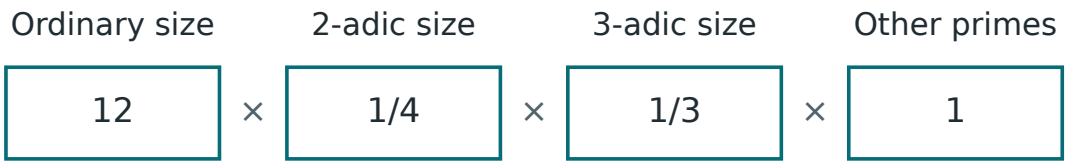
Which zeros count?

The series $\zeta(s) = \Sigma\ 1/n^s$ converges for Re(s) > 1. A mathematical extension called analytic continuation gives values inside the shaded strip. Substituting those inputs directly into the original sum gives the wrong procedure. The negative even integers are known zeros. RH concerns the other zeros and says that every one has real part ½. Masullo's wider claim includes this case.

NIST · definition NIST · zeros Masullo · notation · p. 7

How big is 12 at the prime 3?

In the ordinary sense, 12 is larger than 4. There is another useful way to measure a rational number. At the prime 3, ask how many factors of 3 it contains. A number with more such factors has a smaller 3-adic size.



Product of all four displayed factors = 1

Worked example over the rational numbers. Because $12 = 2^2 \times 3$, the 2-adic and 3-adic sizes undo its ordinary size exactly.

Give each place a coordinate

For a prime p, write $|x|_p$ = p raised to the negative power of the exponent of p in x. For 12 the exponents are 2 at p = 2 and 1 at p = 3. All other prime exponents vanish. The ordinary size and the sizes at each prime are called places. Their product for any nonzero rational number is 1.

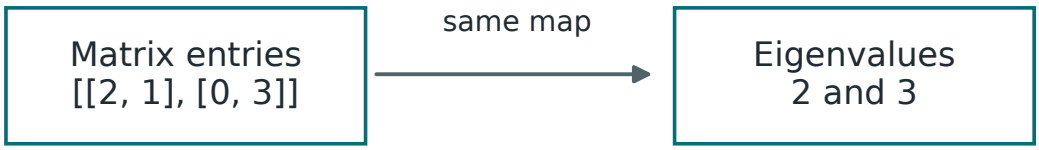
The setting of the proposed theorem

An adele collects a local number at every place, with an integrality restriction at almost all primes. An idele is invertible and is a unit at almost all primes. Identifying ideles that differ by one nonzero number from the field gives the idele class group. A unitary Hecke character is a multiplication-preserving continuous rule on that group, with values on the unit circle. Such rules label Masullo's family of L-functions. Taking the rational field and the constant character 1 gives the classical zeta case after accounting for its infinite-place factor.

Notation · p. 7 Appendix A.1 · p. 539 Chapter 3 · p. 117

Measure one operation twice

Hit a tuning fork and you hear frequencies. You could also describe how each point moves over time. Masullo asks whether two descriptions of a single mathematical operation can similarly connect arithmetic places to the zeros of an L-function.



Both calculations give a trace of 5.

Finite model. The upper triangular matrix has diagonal entries 2 and 3. Its eigenvalues are also 2 and 3. Either route gives the same number.

What a trace measures

A matrix is a table specifying how to change a vector. Its trace is the sum of the diagonal entries. An eigenvector keeps its direction and is multiplied by an eigenvalue. For the displayed matrix, (1, 0) is multiplied by 2 and (1, 1) by 3. Adding these multipliers gives the trace again.

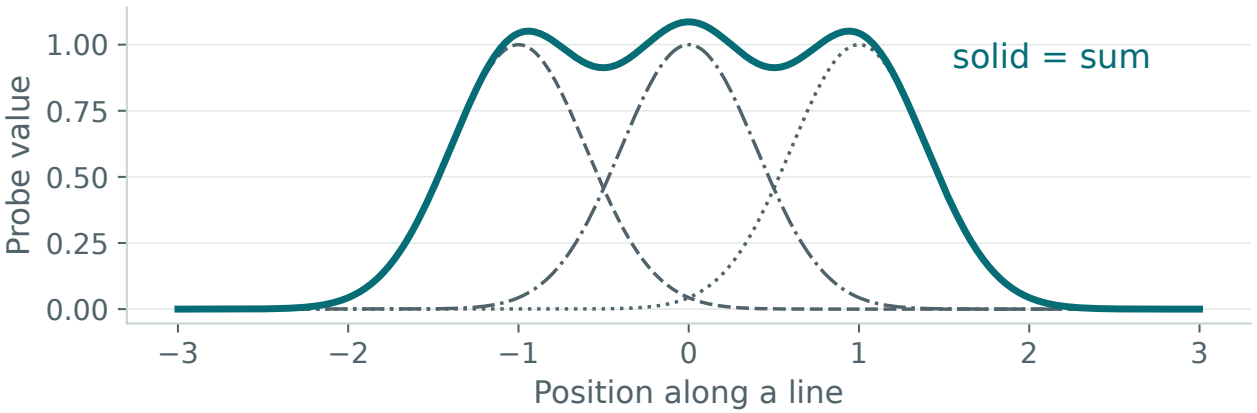
What Masullo proposes to compare

He averages multiplicative shifts of a function using a chosen weight, called a test function. One calculation describes the average using coordinates tied to zeros. The other uses places and repeating orbits. He changes the rule for convergence so the first calculation selects zeros in a narrow strip. The difficult claim is that the second calculation still gives the specified arithmetic answer for each strip width. Without that equality, restricting what the first calculation sees tells us nothing about zeros outside it.

§1.2.1 · p. 13 §1.3 · p. 21 §1.4 · p. 27

Add up copies of one probe

Imagine sliding a small bell-shaped probe across equally spaced marks. Near each mark its contribution is large. Far away it is tiny. Adding the readings detects the whole pattern, even though you used just one probe.



Finite analogy. Three translated Gaussian bells (thin lines) add to the solid line. Masullo sums multiplicative copies indexed by nonzero field elements.

The actual sum

Take a rapidly decaying test function f on the adeles. At an idele g, Masullo forms $\Phi(f)(g) = \Sigma\, f(\xi g)$, where ξ runs through all nonzero elements of the number field. Replacing g by ηg, with η from the same field, only reorders the sum. That invariance makes the result a function on idele classes.

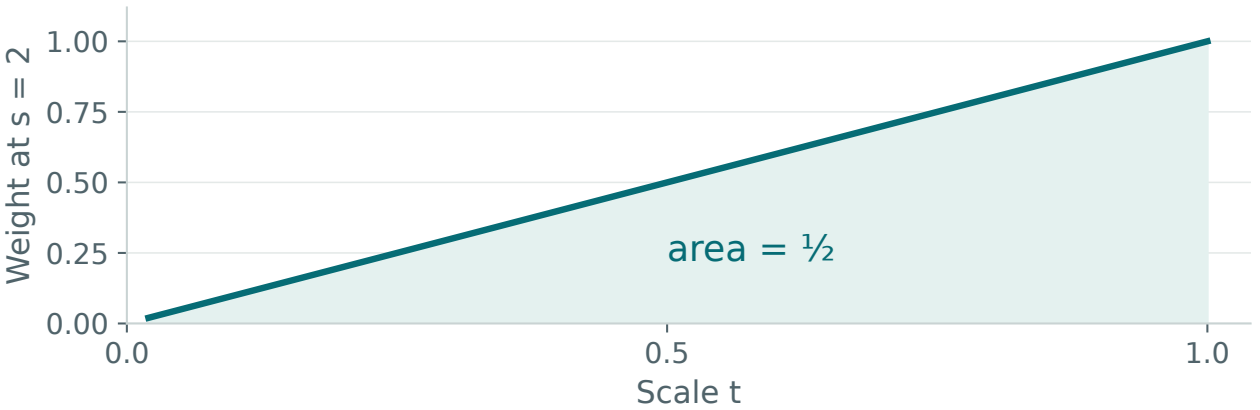
Why he varies f

One probe can miss a feature. A whole family gives a space of theta functions on which shifts and averages can act. Masullo also works with a finite set S of places and uses S-units, numbers that are units outside S. Bounds on the sums at small scales make later transformations possible. For a familiar comparison, at $x = 0$ the three plotted bells add to $1 + 2e^{-\pi}$, about 1.086.

§1.1.4 · p. 12 §§2.1–2.2 · p. 59 §§2.3–2.4 · p. 98

Hear a change of scale

A Fourier transform measures how a signal responds to frequencies. A Mellin transform measures how it responds to changes of scale. Multiplying by 2 always shifts log(t) by the same amount, whether t began at 3 or 300.



Worked Mellin integral at s = 2. The triangle beneath the line from (0, 0) to (1, 1) has area ½.

A transform you can calculate

Take a probe that equals 1 for $0 < t < 1$ and zero elsewhere. Its Mellin transform at real $s > 0$ is the area under t raised to the power $s - 1$ from 0 to 1. This equals $1/s$. At $s = 1$ you get the area of a unit square. At $s = 2$ you get the triangle in the drawing. This simple probe illustrates the transform; later chapters impose smoothness.

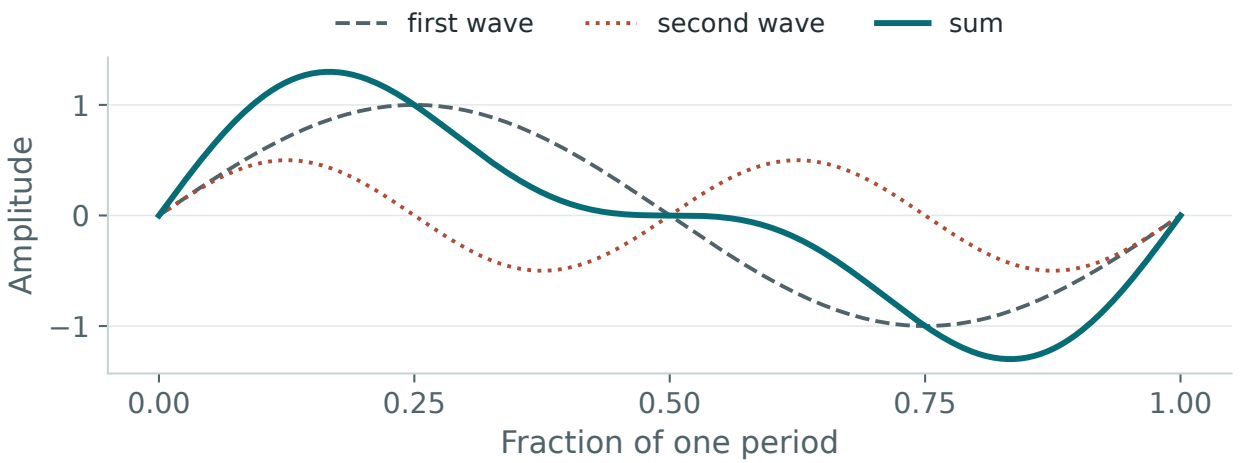
Where the L-function enters

Masullo transforms the theta functions from Chapter 2, weighting scales by a character and a complex power. A transform on the multiplicative group uses a measure unchanged by scaling, dt/t in the positive-real example. The choice of local probes produces local factors. The global integral relates them to the Hecke L-function. This lets the next chapters describe the same information using either functions on scales or functions of a complex variable.

§3.1 · p. 118 §§3.2–3.3 · p. 122 §3.4 · p. 141

Unpack a pattern into waves

On a clock, one turn brings you back to where you started. For scales, decide that t and 2t represent the same point. Then 1, 2, 4 and 8 agree, and the interval from 1 to 2 supplies one full turn.



A worked two-wave expansion. At a quarter turn, the first wave is 1 and the second is 0. At half a turn, both vanish.

Coordinates for a function

A basis lets you describe a complicated object by simpler ones. In the graph, two coefficients, 1 and ½, specify the entire solid curve. A Schauder basis permits infinitely many coefficients, with the ordered partial sums required to converge back to the function. The meaning of convergence matters.

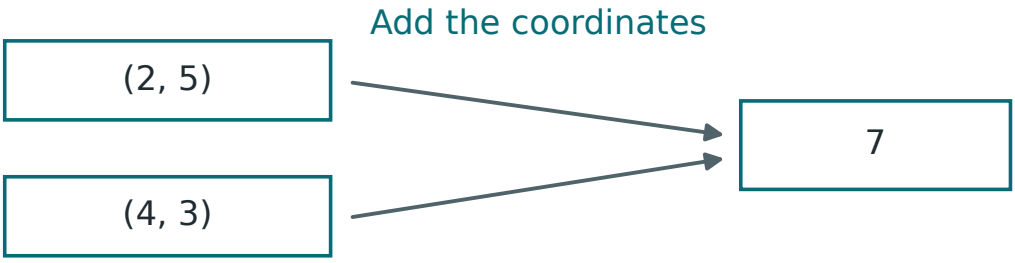
Why periodicity helps

Write u = log(t)/log(2). Doubling t adds 1 to u, so a function periodic in u is unchanged by doubling. Chapter 4 constructs periodic coordinate systems for more general groups and deals with repeated contributions, called multiplicities. These coordinates are used later to compute the effect of a convolution, an average over translated copies, one component at a time.

§§4.1–4.2 · p. 155 §4.4 · p. 188 §§4.5–4.6 · p. 197

Which differences can you discard?

Two inputs can give exactly the same reading. To study the instrument, it helps to identify such inputs. The rule “add the coordinates” gives the same reading for (2, 5) and (4, 3).



Difference = (−2, 2), which adds to 0.

Ordinary addition is a small example of a quotient. Inputs with equal totals represent a single output. The construction in the draft uses the theta operator.

Keep the observable part

A kernel, in this sense, is everything an operation sends to zero. A quotient treats two inputs as equivalent when their difference belongs to that kernel. Masullo uses this construction for theta sums. He then specifies the topology, the rule that says when a sequence of inputs or outputs converges.

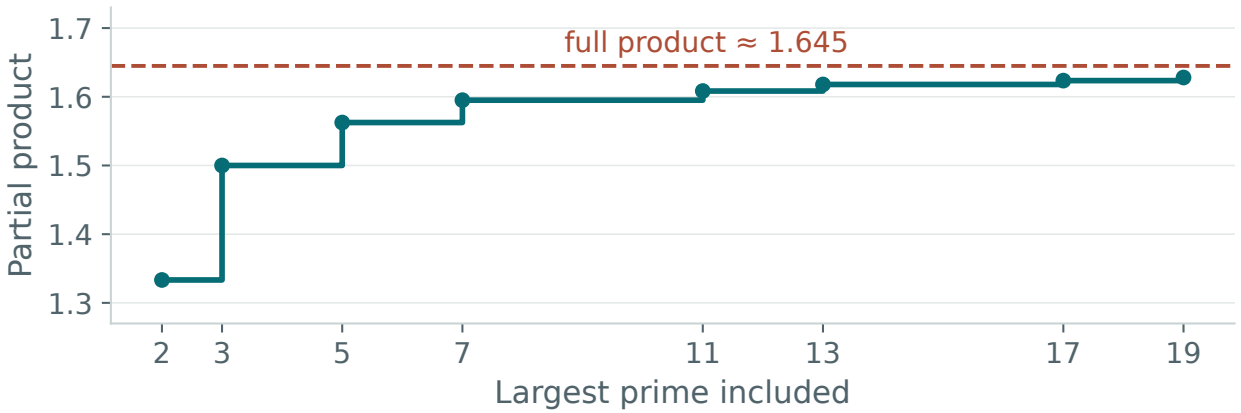
Three structure results

Chapter 5 first describes the spaces and their convergence. Principal-value operators then help separate singular contributions associated with finite places. A matching rule, the S-local Čech sequence, tells when local descriptions fit together. The third result turns to continuous measuring rules on these spaces, including rules that respond to translations by a fixed multiplier. These will later detect zeros and record repeated ones.

§§5.1–5.2 · p. 221 §§5.3–5.4 · p. 240 §5.5 · p. 277

How much information fits in S?

Start with the factors for 2 and 3 in the product for ζ(2). They give (4/3)(9/8) = 3/2. Including 5 gives (3/2)(25/24) = 25/16. Each additional factor changes the answer.



Calculated with Euler factors at s = 2, where the product converges. The running product approaches π²/6. A partial product is only a model of the chapter's transition maps.

Pass information between spaces

S names the places treated individually. The transition from S-local to global extends a probe by standard factors at the other places. Under the Mellin transform this inserts the Euler factors outside S. The chapter identifies what these maps can reach. For a chosen character component, it studies how large S must be for the maps to reach the required space.

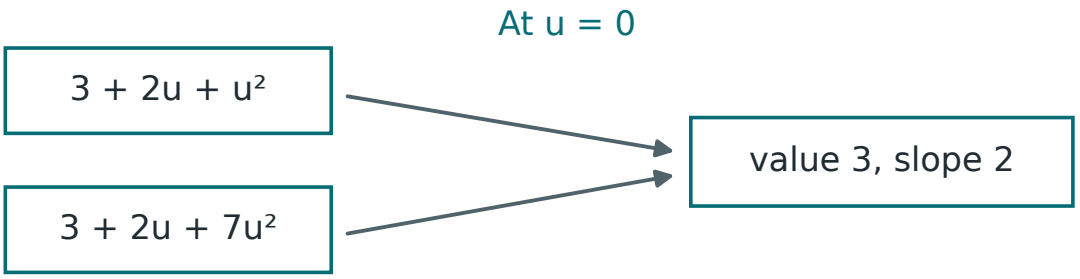
Why the inverse step is delicate

To undo multiplication you must know the multiplier is nonzero where you divide. Masullo's overview invokes established nonvanishing on the line Re(s) = 1 and bounds used in moving contours. This gives a route from local coordinates to global ones. Extending the elementary product in the drawing into the critical strip by ordinary multiplication would skip that work.

§1.4.7 · p. 30 §§6.1–6.2 · p. 287 §6.3 · p. 302

How a quotient can store a zero

On page 13 Masullo compares his construction to keeping the data of a function along its zeros. You can see what this means without knowing any geometry. Work with polynomials near u = 0 and agree to ignore every multiple of u^2.



Worked quotient by multiples of u^2. The polynomials give the same value and slope at 0. Their difference, $6u^2$, is invisible to both measurements.

Two numbers survive

Write a polynomial as $a + bu + cu^2 + \dots$. After removing multiples of u^2, only a and b remain. You read them from the value and the first derivative at 0. Try $3 + 2u + u^2$ and $3 + 2u + 7u^2$. Evaluating gives 3 for both. Differentiating gives slopes $2 + 2u$ and $2 + 14u$, so both slopes at zero are 2.

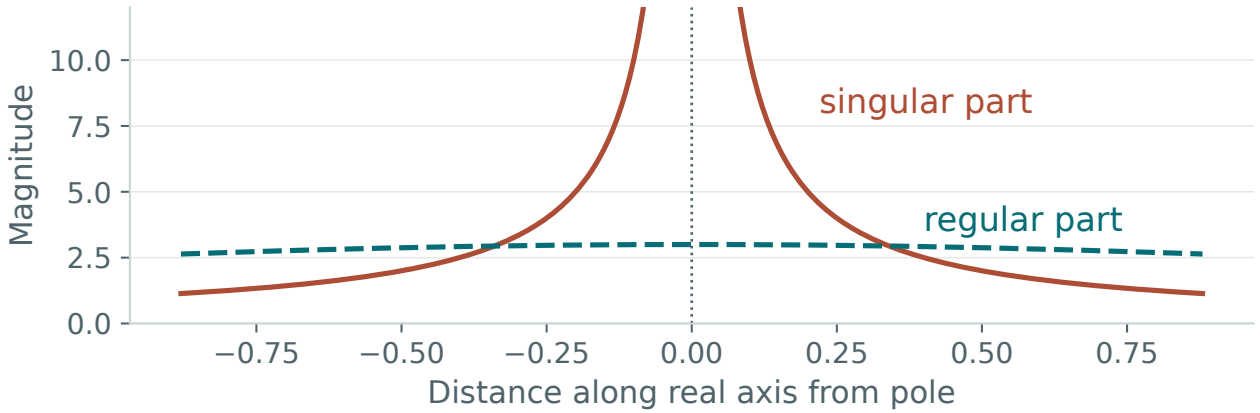
From the example to the manuscript

In this example the repeated factor u^2 leaves two independent readings. In Masullo's divisor sequence, taking a quotient is intended to retain values and derivative data at zeros of an L-function, with appropriate factors for its poles. Chapters 5–7 build function spaces and bases needed for this infinite-dimensional version. Later the choice of norm determines which readings are continuous enough to contribute to the trace.

Divisor sequence · p. 13 Strong duals · p. 277 Spectral basis · p. 438

Read what blows up near a point

If a formula behaves like 1/(s − a) near a, it has a pole there. Separate out the part that blows up. The remaining, regular part can then be studied near a without that singularity.



Teaching example on a real slice. The solid curve shows |1/x| near zero; the dashed curve is a bounded regular term. Poles of local factors play a different role from zeros of the global L-function.

Keep the principal parts

For $3/(s - a)^2 + 2/(s - a) + 5$, the principal part consists of the first two terms. The highest inverse power is 2, so the pole has order two. Chapter 7 organizes such data for S-idelic zeta functions. It groups test functions into Paley–Wiener levels, subspaces with controlled support or growth, and builds bases for the singular pieces.

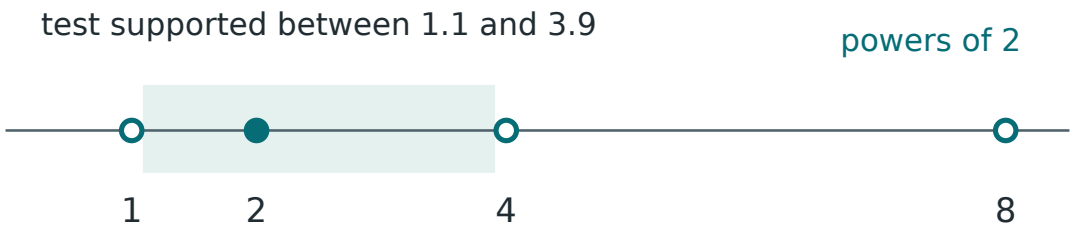
What must survive an infinite sum

A point-evaluation rule reads the value at a given input; derivative rules also see repeated vanishing. The chapter treats infinite sums of elementary operators and how contributions from different factors can be separated. Later chapters use this machinery to express operators in coordinates and justify taking a trace. A finite numerical sample of poles cannot establish those convergence properties.

§§7.1–7.2 · p. 311 §7.3 · p. 340 §§7.4–7.5 · p. 366

Why some primes drop out

Start at 1 and multiply by 2 on every step. You visit 1, 2, 4, 8, … . If a probe is active only between 1.1 and 3.9, it sees just one of the marked powers. For powers of 5, it sees none.



Worked scale example. The shaded interval is the support of a test. The filled point is the only displayed power of 2 inside it.

Distinguish fixed points from returns

A fixed point stays put after one action. A periodic orbit returns after some repetition, with the return read in the relevant quotient space. Chapter 8 identifies the arithmetic orbits and their stabilizers, the symmetries that leave a point unchanged. Forgetting the stabilizers would discard information used in the orbital traces.

A reusable calculation

On positive scales modulo multiplication by q, passing to log(t) gives period log(q). A periodic trace samples the probe at powers of q. A probe supported strictly between 1 and q misses all of them. Masullo uses this mechanism for the factors outside a suitably large finite set S. Chapter 11 must still connect that model calculation to the pieces of the actual global kernel.

§§8.1–8.3 · p. 381 §§8.4–8.5 · p. 389 Introductory orbit example · p. 40

Decide which readings stay stable

Suppose two signals look nearly identical under your chosen measure of distance. A stable measuring rule must give nearly identical readings on them. Change the measure and a reading that used to be stable may cease to be so.

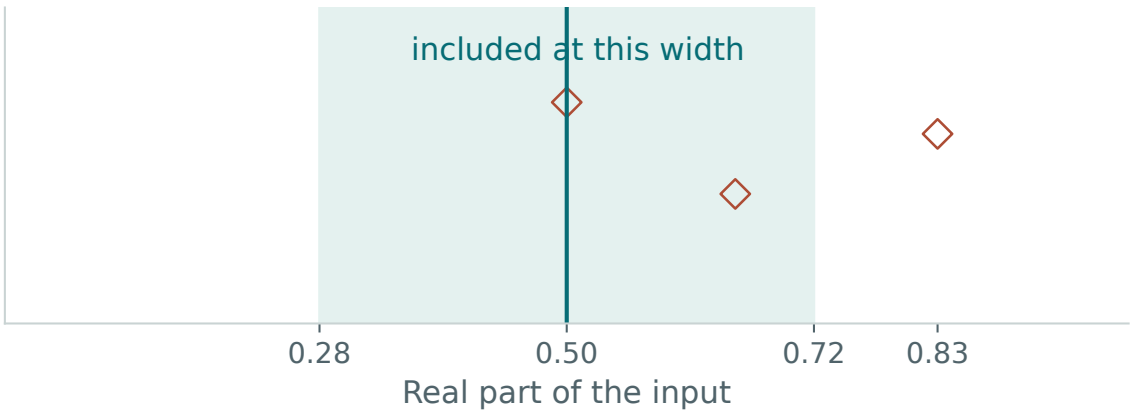


Illustration with invented inputs. Diamonds mark test positions, not known zeros. A strip of half-width 0.22 around 0.50 contains 0.65 and excludes 0.83.

The role of the norm

A norm measures size; it also gives a distance through the size of a difference. A completion includes limits of all sequences that converge in this distance. Masullo uses Hardy–Sobolev norms. They weight signals at small and large scales. In the transform picture, the weights correspond to the edges of a strip I = [−a, a] in coordinates shifted by ½.

A filter, with a burden of proof

At degree zero the proposed spectral calculation uses evaluations supported inside the strip. Higher degrees bring boundary and multiplicity qualifications. For a = 0.10 a hypothetical zero at real part 0.65 has shifted coordinate 0.15, outside. That observation alone cannot rule it out. The force of the argument would come from proving that the geometric calculation still works at each positive width.

§1.4.8 · p. 31 §§9.1–9.2 · p. 395 §9.3 · p. 419

Let a repeated zero count twice

For a finite matrix you can sum its diagonal. With infinitely many coordinate directions you also have to prove that the sum has a value independent of the approximation. This is what the trace-class estimates are meant to control.

3	1
0	3

one value: 3

two copies: 3 + 3 = 6

Finite Jordan block. The off-diagonal entry 1 affects how the map acts, but its trace is still 3 + 3 = 6. This illustrates why a multiplicity weights a trace.

How the zeros enter

The draft forms a quotient by the space supplied by theta functions. On the transform side, zeros constrain values and derivatives; the quotient retains those constraints as coordinates. A zero of order two can carry two pieces of information, a value and a first derivative. In these coordinates an averaging operator gives blocks with the same number down each diagonal.

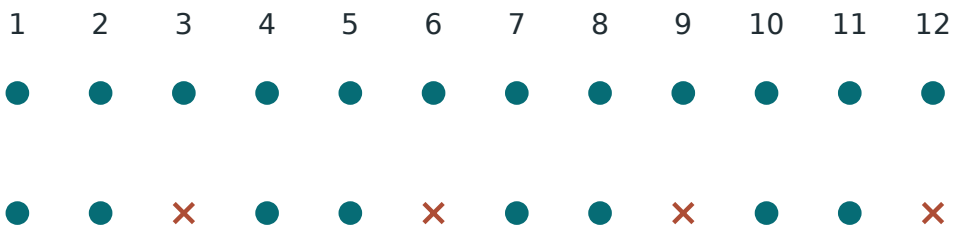
What the spectral sum says

Select a Hecke character. For each included zero ρ the proposed expression adds the transformed test value at ρ − ½, weighted by its multiplicity. The sign depends on which side of the trace identity the sum occupies. The author uses finite-rank approximations and trace-norm convergence to pass from finite blocks to an infinite trace. This is one side of the comparison.

§1.2.1 · p. 13 §§10.2–10.3 · p. 438 §10.4 · p. 455

Peel away one prime at a time

At the prime 3, the locally integral numbers include multiples of 3. The local units exclude them. Changing one condition to the other removes a recognizable slice of the input. Masullo makes this change successively at different primes.



Lower row excludes the multiples of 3.

Finite arithmetic model. In the upper row all twelve integers are locally integral at 3. In the lower row the crosses mark precisely the four excluded multiples.

What a shell records

Think of two nested sets. Subtract their indicator functions and, up to sign, you pick out their difference. Masullo stratifies the ideles by divisibility conditions and applies differences of indicators to a lifted kernel, the two-variable function describing his operator. Successive layers can then be studied on coordinates tied to a prime.

Why this matters to the trace

The proposed decomposition separates a remainder associated with S and a collection of orbital terms outside S. For a fixed test of bounded support away from norm one, S can include all small prime norms. For larger ones the periodic sampling argument predicts zero contribution to the trace. Passing from that statement for individual shells to a statement about their combined diagonal is a further analytic step.

§1.3.2 · p. 26 §§11.1–11.2 · p. 475 §11.3 · p. 491

Can you add the diagonal?

Draw a square and its diagonal. The line has zero area inside the square. You can change every value on the line without changing a two-dimensional area integral. Yet that change can alter the integral along the line completely.

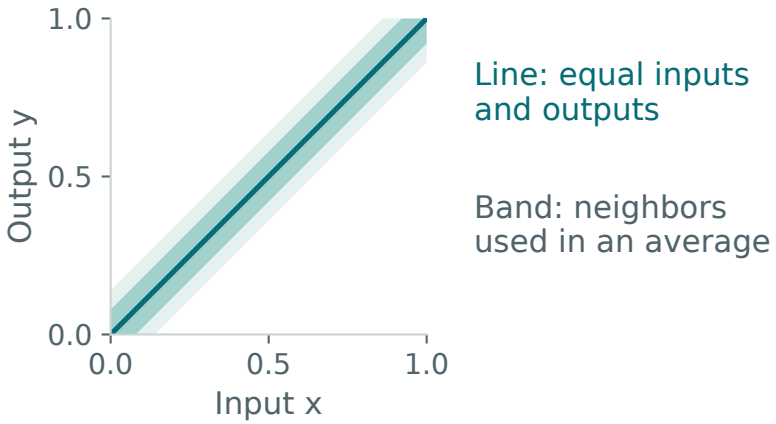


Illustration. Averaging kernel values in a narrow band around y = x is different from assigning values directly to the line. The limit needs a justification.

Why smooth the kernel

A kernel describes how much input at y contributes to output at x. Integrating K(x, x) is the continuous analogue of adding the diagonal of a matrix, when sufficient hypotheses hold. Masullo smooths and cuts off kernels, studies their diagonal averages and proposes to remove these restrictions in a limit. Two kernels equal almost everywhere need not have the same pointwise diagonal.

Where the global argument needs text

Section 12.3 on p. 514 is reserved for putting the global-to-S-local pieces together. It contains a heading and an instruction to add the section. Section 12.4 then gives an S-local calculation, with local integrals and a cutoff term involving φ(1). Reading the S-local calculation does not by itself supply the missing global justification. The introduction, especially pp. 26–28, gives the intended bridge.

§§12.1–12.2 · p. 498 §12.3 and §12.4 · p. 514 Intended bridge · p. 27

What the trace formula would buy

Choose a smooth test φ that is zero outside a bounded region of the idele class group and vanishes on the norm-one subgroup. That last condition means φ is zero at every class of norm one. The proposed formula compares two readings of the same test.



Masullo’s proposed equality for permitted test functions.

Schematic of Theorem 12.5.1.1, pp. 519–520. The special terms are Mellin-transform evaluations. The sum weights the zeros inside the chosen strip by multiplicity.

What is on each side?

On the left, the real, complex and finite places supply local integrals. For finite places their weights involve log(q), where q counts the elements of the residue field. On the right sit two special evaluations of the test, at powers +½ and −½ of the norm, followed by a subtraction for each included zero. Both sides depend on φ.

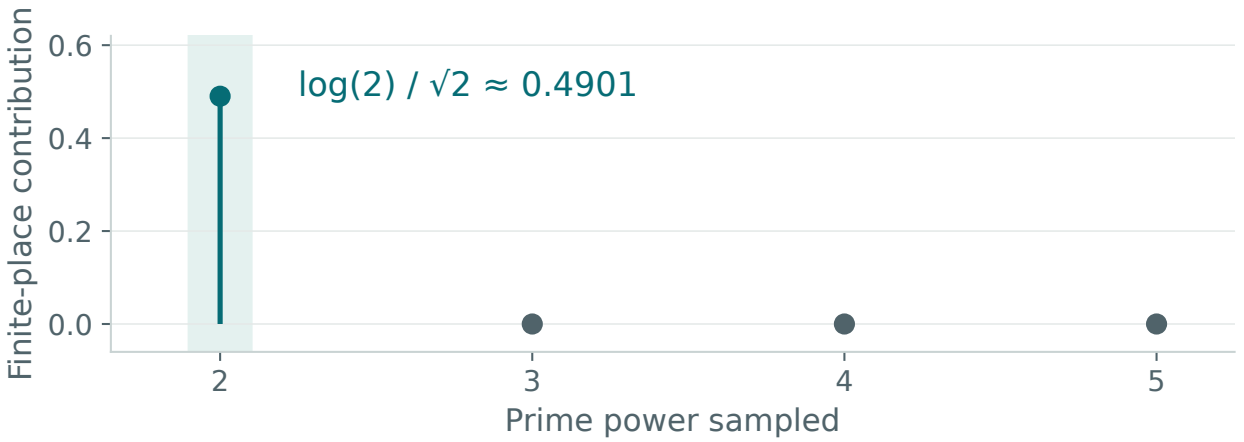
The hinge on page 520

The proof adds the shell integrals and the S-local remainder. It invokes Theorem 11.3.2.1 to remove the outside shells, Theorem 12.4.1.1 for the remainder, and Theorem 10.4.2.8 for the spectral trace. The result is stated for a strip centred at ½. Its use for arbitrarily narrow strips is decisive. The fixed ½ in the kernel normalization on p. 519 and the adjustable half-width α on p. 520 have different jobs.

Orbital traces · p. 495 Theorem 12.5.1.1 · p. 519 Assembly · p. 520

Try a probe that singles out 2

Set the number field to the rationals and the character to 1. Choose a smooth bump whose value at 2 is 1 and whose support lies between 1.9 and 2.1. Every other positive integer lies outside its support. So does every reciprocal prime power.



Calculated from the character-by-character formula on p. 28. The only surviving finite-place sample is p = 2, exponent 1. This evaluates one part of the formula.

Follow one term through the formula

For prime p and integer m ≥ 1, the finite contribution in this specialization is $(\log p)/p^{m/2}$ times $[\varphi(p^m) + \varphi(p^{-m})]$. The logarithm is natural. At p = 2, m = 1, the bracket is 1 + 0. At p = 2, m = 2 it is φ(4) + φ(¼) = 0. The full finite sum here is $\log(2)/\sqrt{2} \approx 0.4901$.

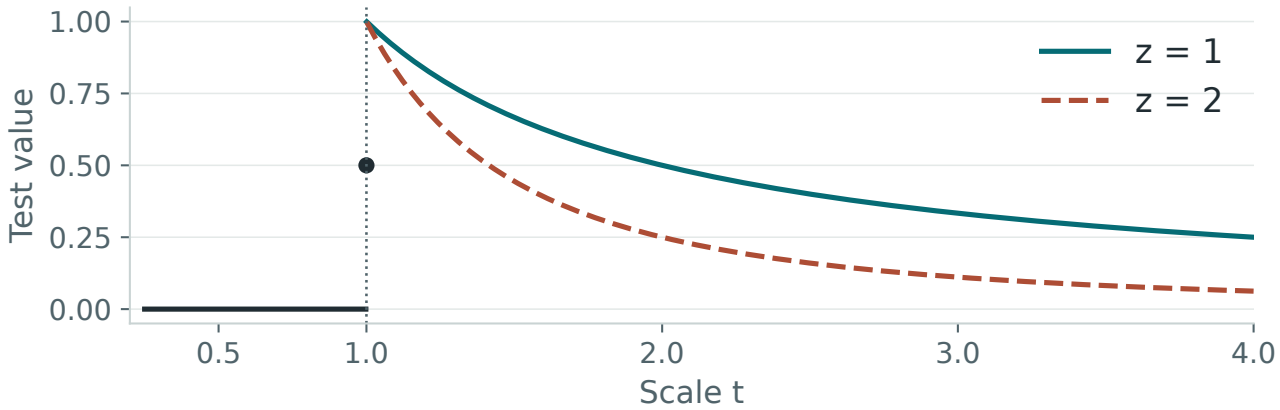
What would be needed to match it

The real-place integral still depends on the shape of the bump. The two exceptional Mellin values do as well. The proposed identity relates these three ingredients to a sum over zeros. Knowing the finite sample alone does not locate a zero. It does show, with an actual prime, what the geometric side is measuring.

§1.4.1 · p. 28 Finite-place sum · p. 29 Theorem 12.5.1.1 · p. 519

Choose a probe with a dial

One bump isolates a prime. To extract information about a whole L-function, Masullo turns a dial z and tests many scales at once. In §1.4.2 he chooses $\varphi_z(t) = t^{-z}$ above 1, zero below 1, and ½ at the jump.



The test for z = 1 and z = 2 on the positive real axis. At t = 2 the values are ½ and ¼. The dot at the jump uses the convention on p. 28.

Make it eligible for the formula

This test has a jump and extends to arbitrarily large t. The trace theorem calls for a smooth test of compact support that vanishes at norm one. The outline smooths the jump and multiplies by a cutoff supported strictly above 1. With the appropriate character weight, these admissible tests are then used in a limiting argument. Agreement at the jump itself is unnecessary for pointwise convergence away from it.

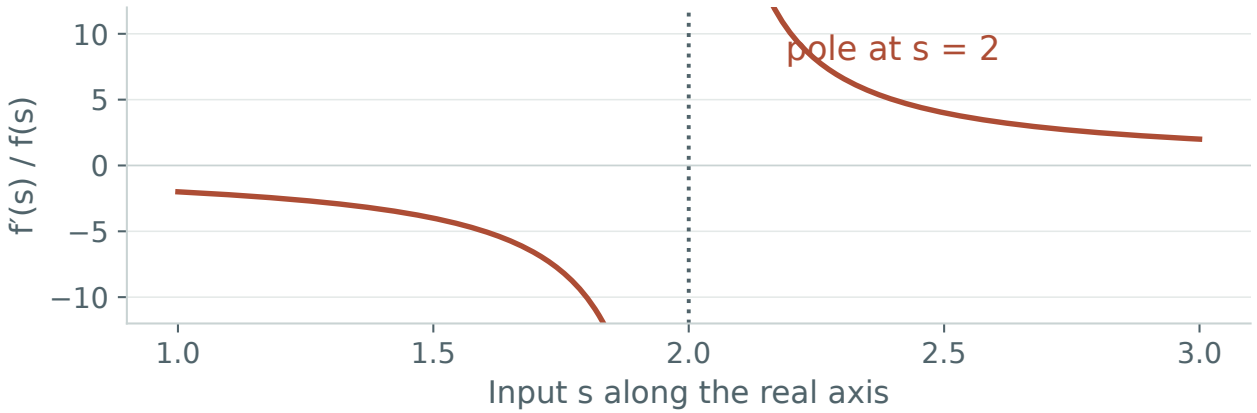
Why use a complex dial?

As long as t > 1, multiplying the finite-place weight $p^{-m/2}$ by $(p^m)^{-z}$ produces $p^{-m(z+½)}$. Summing over primes and exponents gives the arithmetic series behind a logarithmic derivative. For complex z, the same test also supplies oscillations that distinguish heights on the zero map. The analytical issue is whether the limiting operations preserve the identity in the required region.

§§1.4.2–1.4.4 · p. 28 Second account · p. 54 Chapter 13 · p. 522

How a zero announces itself

To find zeros of L, Masullo studies L′/L. The prime mark means the derivative, the local rate of change. Dividing by L magnifies the behavior where L is small. At a genuine zero, the resulting expression has a pole.



Calculated for f(s) = (s − 2)² on a real slice. Its logarithmic derivative is 2/(s − 2). Moving closer to 2 makes its magnitude arbitrarily large.

Work a double zero by hand

Differentiate f(s) = (s − 2)² to get 2(s − 2). Dividing cancels one factor, leaving 2/(s − 2). At 2.1 the result is 20; at 2.01 it is 200. In general, if L(s) = (s − ρ)^{m}h(s) with h(ρ) ≠ 0, the same calculation gives L′/L = m/(s − ρ) + h′/h. The second term stays regular near ρ.

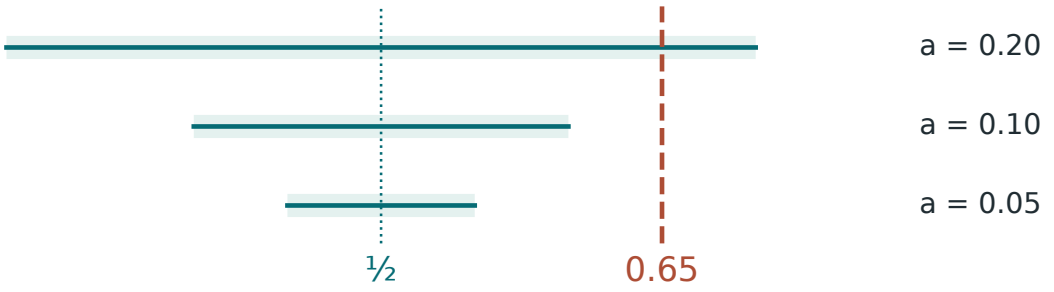
Recover it from primes

For ζ in the half-plane where its Euler product converges absolutely, −ζ′(s)/ζ(s) = $\Sigma_{p}\,\Sigma_{m}$ (log p)p^{-ms}, with m ≥ 1. The minus sign follows by differentiating p^{-s}. At s = 2, the p = 2 part is (log 2)(¼ + 1/16 + …) = (log 2)/3. Either sign of the logarithmic derivative detects the same poles. The proposed proof needs to carry this information beyond the initial half-plane.

§§1.4.3–1.4.5 · p. 29 Appendix B.4 · p. 606 NIST · Euler product

Why every strip width matters

Suppose someone proposes a zero with real part 0.65. It falls inside a strip of half-width 0.20 around ½. It falls outside one of half-width 0.10. Filtering it out is easy. Proving that it cannot exist is the hard part.



Nested strips in a horizontal cross-section. The dashed line is a hypothetical zero coordinate. The figure gives a choice of widths, not a numerical test of RH.

Continue the formula

In §1.4 Masullo starts where the prime sums converge. He seeks uniform bounds on smoothed tests as z varies over small closed regions. Montel's theorem can then extract a subsequence converging on compact subsets to a holomorphic function, one that is complex differentiable. This reasoning requires the bounds and the identification of the limit. Analytic continuation then extends an identity from an overlapping region where it is already known.

The contradiction, using 0.65

Write w = z + ½. The proposed consequence is that L′(w)/L(w), after allowing for the known principal pole, is holomorphic for Re(z) > a, for every a > 0. For Re(w) = 0.65 take a = 0.10. Then Re(z) = 0.15 exceeds a: the ratio would be regular there. A zero would force the pole on the previous page. This conditional contradiction is the step the narrowing-strip argument is designed to deliver.

§§1.4.3–1.4.5 · p. 29 §1.4.8 · p. 31 §1.8.5 · p. 55

Reflect the other half

A symmetry alone can leave two points on opposite sides of a line. It becomes decisive only after one side has been ruled out. That is why the right-hand nonvanishing argument comes first.

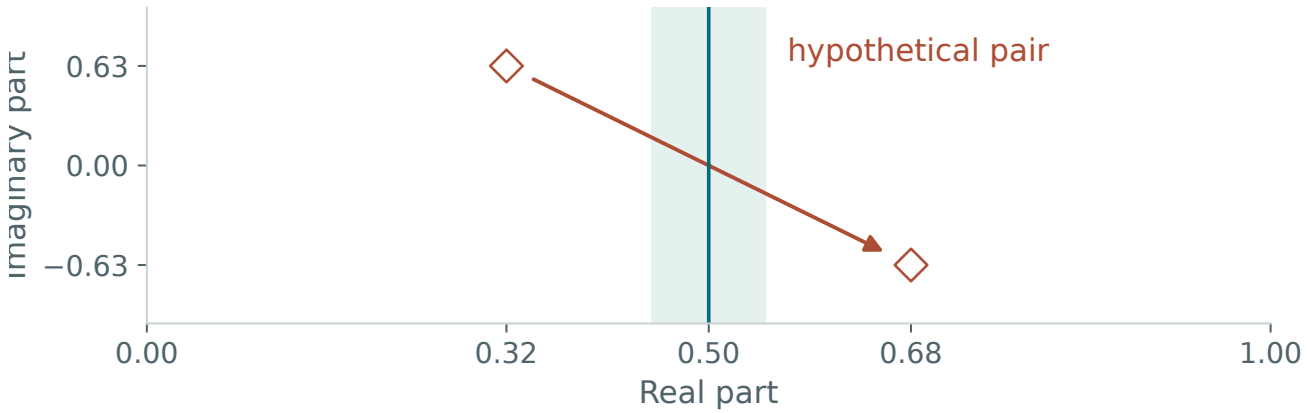


Illustration of s ↦ 1 − s. The two diamonds, at 0.32 + 0.63i and 0.68 − 0.63i, are invented to show the symmetry. The vertical line marks real part ½.

Apply the functional equation

The completed L-function includes factors for the infinite places. Its functional equation pairs the value for a character at s with that for the conjugate character at 1 − s, up to a nonzero factor. A nontrivial zero with real part 0.32 would force a zero with real part 0.68 for the partner. Masullo's proposed right-half result applies to both characters. Under that result, only the middle line remains.

How this gives ordinary RH

Choose the rational field and the trivial character. The finite L-function is ζ. If the preceding argument held for all the claimed number fields and characters, it would include all nontrivial zeros of ζ. The singularity at s = 1 for the rational trivial character is a known pole. It is treated separately in the logarithmic derivative. The negative even zeros of the uncompleted ζ are outside the target.

§1.4.5 · p. 30 Overview and reflection · p. 55 NIST · functional equation

What is missing on page 522?

Turn to the end of Chapter 13. The theorem number is there. Its title is there. The statement that should follow is absent. That matters because a proof must say precisely what it establishes before a reader can check the deduction.

p. 514	pp. 519–520	p. 522
Missing section	Trace theorem	Missing statement

Three places to compare in the November 20 draft. The link beneath each source label opens the relevant page of the original.

What the manuscript actually contains

Page 514 leaves §12.3, the assembly of the global geometric kernel from local contributions, as a heading with an editorial note. Pages 519–520 do give a statement and an argument for Theorem 12.5.1.1, relying on earlier results. Page 522 gives a version of a trace identity marked for correction, cites an unresolved theorem reference and leaves Theorem 13.3.1.2 without its statement. A footnote asks for the main theorems and their proofs to be added.

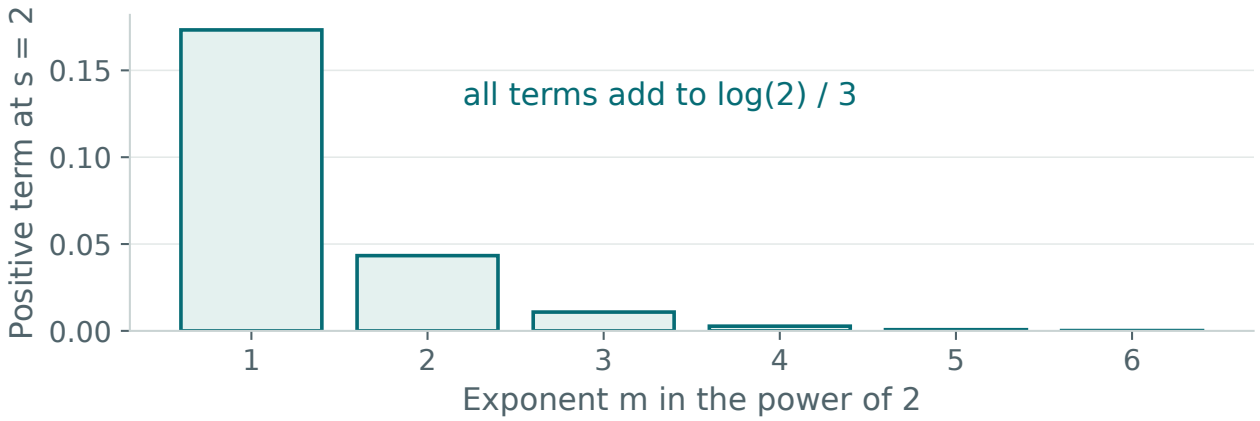
What a reader can conclude

The introduction supplies a mathematical route from the proposed trace formula to the critical line. It is possible to understand that route and to test its elementary inferences. A complete verification also needs the uniform estimates, the global decomposition and its limits, and a fully stated final deduction. In his update Ravi Vakil reported that readers had raised problems and that those issues required resolution. His announcement gives no technical verdict on each step.

Missing §12.3 · p. 514 The trace theorem · p. 519 Final page of chapter · p. 522 Vakil’s update

Check the sign at one prime

You can test a line of the outline by looking at just the prime 2. Page 29 identifies a positive prime-power sum with L′/L. For the trivial character and real inputs greater than one, the usual Euler product gives that sum the opposite sign.



Calculated for the p = 2 Euler factor at input 2. The six bars are the first terms of a positive geometric series. Its full sum is about 0.2310.

Do the sign check

The Euler factor is $E(s) = 1/(1 - 2^{-s})$. When real s increases, 2^{-s} decreases, the denominator grows and E(s) decreases. Its logarithmic derivative is therefore negative. Differentiation gives $E'/E = -(\log 2)2^{-s}/(1 - 2^{-s})$. At $s = 2$ this equals $-(\log 2)/3$. The positive sum in the figure equals $+(\log 2)/3$.

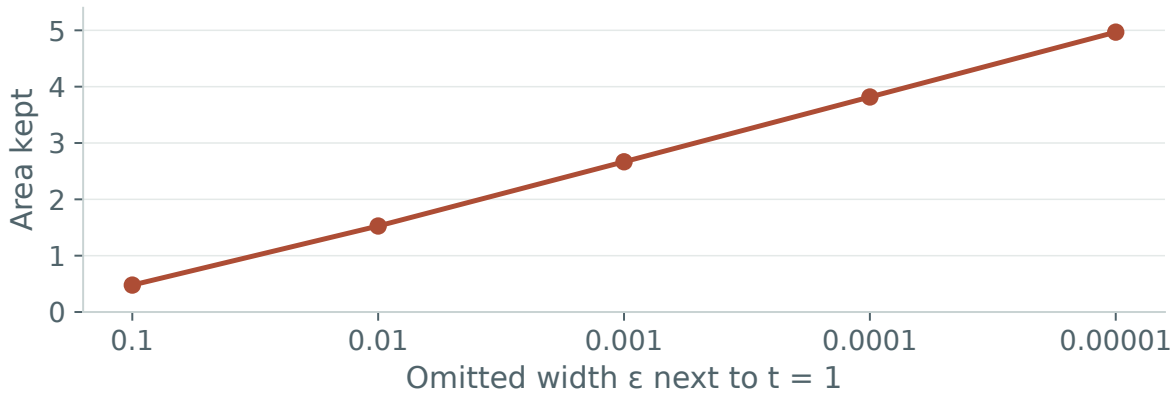
What to take back to the paper

In §1.4.4 the draft writes the limiting logarithmic derivative with a plus sign. The worked calculation on guide p. 22 uses the conventional minus sign for the positive sum. Changing the sign preserves the location of poles, but the signs of the individual terms must agree when the trace identity is assembled. This check is specific to the displayed identification.

§1.4.4 · printed sign · p. 29 Euler product

Zoom in on the jump at 1

Page 29 invokes a uniform bound before applying Montel’s theorem. Try the rational field, the trivial character and the real dial z = 2. The positive real-place term alone reveals a problem with that bound as printed.



Calculated area from 1 + ε to 1.5 beneath t raised to −5/2 divided by 2(t − 1). A hard cutoff makes the growth visible. The argument below also applies to nonnegative smooth approximations.

Where the growth comes from

Take the u = +1 term in §1.4.3. For t > 1, the limiting probe is t^{-2}. Multiplying by √t and by the measure dt/t gives t raised to −5/2 times dt. The denominator is t − 1 and the coefficient is ½. Near 1 the numerator tends to 1. The area grows like ½ log(1/ε). It rises from 0.478 to 2.667 as ε falls from 0.1 to 0.001.

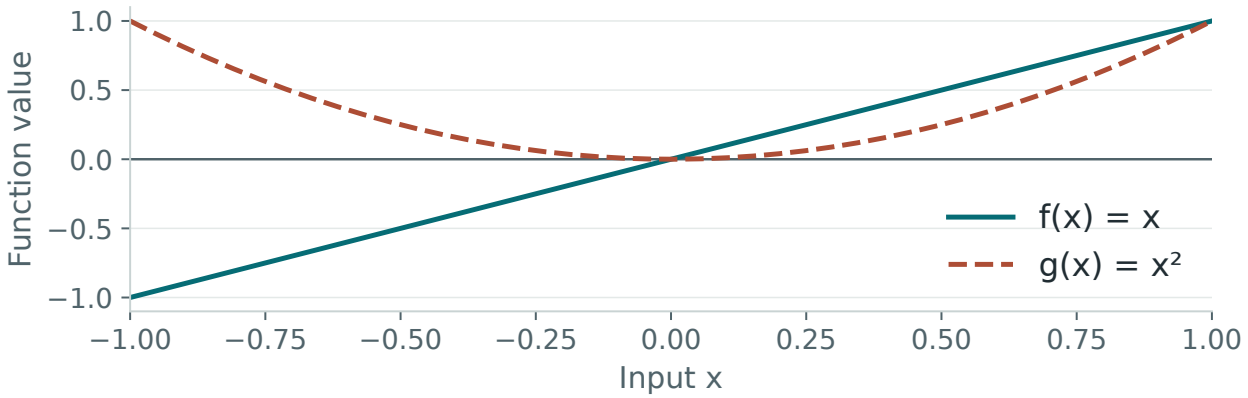
Why smoothing alone cannot fix this estimate

On every fixed interval strictly above 1, nonnegative approximating probes approach the stated probe. The limiting area on those intervals can be made arbitrarily large by moving the left end toward 1. Fatou’s inequality then rules out a common finite upper bound for these nonnegative integrals. The outline needs a further cancellation or subtraction to justify this use of Montel. This is a check of the written argument, independent of the issues reported to Vakil.

Test family in §1.4.2 · p. 28 Bound in §1.4.3 · p. 29 Alternate overview · p. 54

Where a zero lies and how often it occurs

Two functions can vanish at exactly the same point and do so to different orders. The location answers “where?” The multiplicity answers “how many factors disappear there?” They are separate questions.



Calculated examples at the origin. The line x has a simple zero. The parabola x^2 has a double zero. Their logarithmic derivatives have coefficients 1 and 2 at the pole.

Why the draft goes further

Chapter 14 discusses expected bounds on multiplicities and their relationship with conjectures about simple zeros. The cap involving $[F : \mathbb{Q}]$ in Chapter 13 depends on the degree of the number field, the number of independent rational coordinates needed to describe its elements. For $\mathbb{Q}$ that degree is 1. A cap in a trace expression is not by itself a proof that every actual zero is simple.

Where Masullo draws the boundary

The proposed main argument concerns number fields and unitary Hecke characters. Section 14.2 discusses why extending the same multiplicity conclusions to function fields in positive characteristic would run into examples with double zeros. Section 14.3 sketches a possible higher-rank direction and ends with an editorial note. Neither extension is needed to specialize the proposed Hecke result to the classical zeta function.

§14.1 · p. 524 §14.2 · p. 534 §14.3 · p. 536

Where do the arithmetic tools enter?

A. Places, measurements and orbits

Appendix A builds the arithmetic setting used on pages 5, 7 and 14 of this guide. Its Haar measures let the calculation average without giving one translate more weight than another. On positive real scales, an interval from 1 to 2 has the same multiplicative measure as one from 3 to 6. Both have length log(2) when measured with dt/t.

A. Why averages near a point matter

Sections A.4–A.5 develop counting and differentiation results. The latter concerns when averages over smaller neighborhoods recover a function’s value. This supports Chapter 12’s attempt to recover diagonal information from regularized kernels; the reader must still check the conditions for the actual kernels.

B. The earlier explicit formula

Appendix B.4 relates sums over the arithmetic places to sums over the zeros. This established kind of identity uses the zeros without the new arbitrary-width restriction. Masullo’s proposed formula in Chapter 12 would retain only the zeros in a chosen strip when tested on specified functions. That extra restriction carries the new mathematical claim.

B. How far an existing result reaches

Sections B.1–B.3 and B.5–B.7 collect growth bounds, contour and residue tools, recurrence, zero-free regions, and control of reciprocals near the line Re(s) = 1. They supply estimates for constructing bases and transferring data from local to global. A zero-free boundary at 1 leaves the region between ½ and 1 to be dealt with.

A.3 · measure · p. 555 A.5 · differentiation · p. 566 B.4 · explicit formulas · p. 606 B.7 · reciprocals · p. 643

Where do the analytic tools enter?

C. Which limits and traces are allowed?

Appendix C deals with infinite-dimensional function spaces. An ordinary convergent numerical series lets you add infinitely many numbers; an operator series needs a rule for how close two operators are. Section C.3 develops nuclear operators. Chapter 10 needs a sufficiently strong form of convergence to pass traces from finite approximations to the proposed spectral operator.

D. What makes a good test function?

Appendix D constructs the classes of rapidly decaying functions used on locally compact groups. On the real line, the bell e to the power $-x^2$ illustrates smooth rapid decay. At a finite place a probe can be locally constant on small arithmetic neighborhoods. The chapters on theta sums and Mellin transforms need these two kinds of behavior to coexist.

E. What does a quotient forget?

A group action can bring different points to the same class. It can also leave a point fixed. A list of classes alone loses this second piece of information. A groupoid records the moves as well as the points; the topological stacks in Appendix E organize this information across local descriptions. Chapter 8 uses it to distinguish orbit types and stabilizers.

How to use these appendices

Read A.1 for “idele”, B.4 for “explicit formula”, C.3 for “nuclear”, D.1 for “test function” and E.4 for “groupoid”. Each gives background on a particular step. The question of whether the new global trace identity follows belongs to Chapters 11–13.

C.3 · operators · p. 693 D.1 · test functions · p. 745 E.4 · groupoids · p. 883 E.6 · stacks · p. 920

Five checks on the argument

1. Prime products

Expand just the brackets for the primes 2 and 3 at $s = 2$. Which selections give $1/36^2$? Can any other selections from those brackets give it?

2. An arithmetic orbit

A probe is supported strictly between 1.5 and 4.5. Which positive powers of 2 and 3 can it sample? What about powers of 5? State what knowing the support tells you and what further information you need to compute the sum.

3. A repeated zero

For $f(s) = (s - 4)^3$, calculate $f'(s)/f(s)$ and evaluate it at 4.01. What property of the zero can you recover from the answer?

4. A zero just off the line

Suppose a nontrivial zero had real part 0.5004. Choose a positive half-width that would exclude it, then give its reflected real coordinate. Which claim in the manuscript is supposed to turn this drawing into a contradiction?

5. Audit a citation

Trace the dependency from Theorem 12.5.1.1 to §12.3 and from Theorem 13.3.1.2 to §1.4. Which two parts of the released document would you use to distinguish a stated theorem from an outlined deduction?

Euler product Trace calculation · p. 519 Chapter 13 · p. 522

Check your reasoning

1. Unique factors

$36 = 2^2 \times 3^2$. Choose the $1/4^2$ term in the first bracket and the $1/9^2$ term in the second. Their product is $1/36^2$. Unique prime factorization rules out another pair of exponents.

2. Sampling depends on the probe

The possible samples are 2, 4 and 3. The first positive power of 5 already lies outside. Support tells you which terms must vanish. To compute a surviving term, you need the probe's actual value there as well as its arithmetic weight.

3. Multiplicity appears in the numerator

Differentiating gives $3(s - 4)^2$. Cancel the two common powers to get $3/(s - 4)$, which equals 300 at 4.01. The pole has coefficient 3, the order of the zero.

4. Exclude the point from the proposed formula

The distance from ½ is 0.0004. Choose $a = 0.0002$. Reflection gives 0.4996. The necessary input is a valid trace identity at that arbitrarily small width, with enough control on the limiting probes to make the logarithmic derivative regular in the indicated half-plane. A drawing establishes neither estimate.

5. Follow the two trails

Section 12.5 states the trace formula and gives a proof using earlier results. The assembly section on p. 514 is unwritten. Page 522 leaves the divisor theorem's statement to be inserted. Pages 27–30 explain the proposed route from the trace formula to the conclusion and can be read as an outline of that missing final stage.

§1.4 · p. 27 §12.3 · p. 514 §12.5 · p. 519 §13.3 · p. 522

SOURCE NAVIGATION

Open the original at the right page

The page numbers below belong to Masullo’s PDF. Each entry links to the start of a chapter or appendix. Some mobile PDF readers open a file at its cover; enter the printed page number in that reader.

1	Introduction and letter to Peter Sarnak	11
2	Adelic theta functions	57
3	Idelic zeta functions	117
4	Periodic bases	155
5	Spaces of theta and zeta functions	221
6	Local and global structures	287
7	Bases of idelic zeta functions	311
8	Periodic orbits and fixed points	381
9	Hardy–Sobolev topologies	395
10	Spectral side of the trace	433
11	Shell structure	475
12	Geometric side of the trace	497
13	Zeros of Hecke L-functions	521
14	Complements	523
A	Algebraic number theory	537
B	Analytic number theory	587
C	Locally convex spaces	655
D	Schwartz–Bruhat functions	743
E	Topological stacks	823
	Bibliography	933

Open original manuscript Open existing interactive explorer

What you can verify

The manuscript

Alessandro Maria Masullo, The Music of the Primes. The 940-page PDF circulated on November 20, 2026. All links to numbered pages in this guide refer to this release.

Open source ↗

The release and the subsequent warning

Ravi Vakil, “Hurried Announcement: Alessandro Masullo GRH,” November 20, 2026, with a later update reporting concerns raised by readers. Vakil describes his role and Brian Conrad’s feedback.

Open source ↗

The classical statement

Clay Mathematics Institute, Riemann Hypothesis. The institute’s description gives the classical conjecture and explains its connection to the fluctuations in the prime count.

Open source ↗

Functions, zeros and reflection

NIST Digital Library of Mathematical Functions, Chapter 25. See §§25.2, 25.4 and 25.10 for the zeta function, its Euler product, its functional equation and its zeros.

Open source ↗

Edition note

Adrian Dunkley wrote the reading guide. The proposed proof is Masullo’s. The charts use direct arithmetic, numerically evaluated zeta zeros, or explicitly labelled teaching examples. The guide follows the route through the released chapters. Its checks of the sign and the endpoint bound on manuscript p. 29 are independent readings of those expressions. The unfilled sections are identified at the points where the argument needs them.