The de Bruijn–Newman constant: three conditional upper bounds

READING LEVEL EVERYONE · EXPERTS

Three conditional upper bounds, computer-assisted, checked within this project; not refereed.

proved in writing 10 · conditional 3 · computed 3 · open 2 · cited 1

Overview

The Riemann hypothesis is a famous open claim about the zeros of the zeta function. A related single number, the de Bruijn-Newman constant, written Λ, measures how much a gentle warming of the zeta function is needed before all its zeros line up. Rodgers and Tao proved that Λ is at least 0, and Λ equal to 0 would be the Riemann hypothesis, which is open. This book works on the other side: how small can we show Λ to be? The smallest refereed upper bound is 0.2. The book gives three computer-assisted certificates, for Λ at most 0.170, 0.1715 and 0.172. Their status is conditional: each depends on published inputs, on an inherited chain of computer steps, and on changed proof steps, nine, five and four of them. Eight of those steps have written arguments, proved as scoped and read within this project, and the ninth is a program whose 13,808,993 rows a second program recomputed. Nobody outside this project has read them. Others have posted smaller figures that are not refereed, and the credits page lists them. Six written recipes let any reader recheck, in exact fractions, that the stored files agree and that all 76 final gates are positive. They do not repeat the long runs. Nothing in this book is proved in Lean.

CHECK IT YOURSELF

The check guide for this book · How to check a result yourself

Cite: Sonny (2026). The de Bruijn–Newman constant: three conditional upper bounds (preliminary). Go check yourself. https://gocheckyourself.org/read/riemann-debruijn-newman/everyone/

1. Start here
2. The simple version
3. Try it yourself
4. Articles

Start here

Sonny

What this book shows

- My result. Three computer-assisted certificates (computer-checked files that show a bound holds). Here Λ is the de Bruijn–Newman constant (how much a "warming" of the zeta function is enough to make all its zeros line up): $\Lambda \le 17/100 = 0.170$, $\Lambda \le 343/2000 = 0.1715$ and $\Lambda \le 43/250 = 0.172$. Each follows from Polymath15's criterion (Polymath15 was a large online collaboration; its criterion turns one positivity check at a given height into a bound on Λ), Platt and Trudgian's verified height (how far up the zeros of zeta have been checked by computer), the inherited pipeline (the chain of computer steps I inherited from Jude Gomila) and my changed proof steps (nine, five and four). CONDITIONAL on the listed conditions. Eight of the steps are PROVED as scoped, with written arguments; the ninth, a program, is COMPUTED.

- Every passed-on number checked. Every number one stage passes to the next is rounded in its safe direction from its exact binary value, and all 76 final gates (final checks) recompose positive in exact fractions. Six written recipes let any reader recheck this in seconds.
- Why it matters. Rodgers and Tao proved $\Lambda \ge 0$, and $\Lambda = 0$ is the Riemann hypothesis (the long-standing claim that the special numbers called the zeros of the zeta function all sit on one line); the smallest refereed upper bound (one checked by independent experts before publication) is 0.2. These certificates rest on the same published inputs as that 0.2, Polymath15's criterion and Platt and Trudgian's height, with each extra assumption named.
- In Lean: nothing in this book is Lean-proved (Lean is a proof-checking program that tests every step of a proof against fixed rules).
- Every theorem in this book is listed in the theorem library.

The question

How far can the heat flow of the Riemann xi function (a way to "warm up" the zeta function a little at a time, like heat spreading through a metal bar) be run before its

zeros are guaranteed to be real? That distance is a single number, the de Bruijn–Newman constant, written Λ.

Here is the setting in plain words. The Riemann hypothesis is about the zeros of the zeta function. In 1950 N. G. de Bruijn built a family of functions H_t, one for each real number t, from a Fourier transform (a way of breaking a function into waves) of the zeta function's completed form. At t = 0 the zeros of H_0 correspond to the non-trivial zeros of zeta, and the Riemann hypothesis says they are all real. De Bruijn showed that if all the zeros of H_t are real, then so are those of H_s for every s ≥ t. In 1976 Charles Newman showed that there is a single threshold: all the zeros of H_t are real exactly when t is at least that threshold. He conjectured that it is not negative. The threshold Λ is the smallest value such that all zeros of H_t are real for every t ≥ Λ.

The question this book works on is a bound: how small can we show Λ to be? Polymath15 (2019) proved Λ ≤ 0.22, and with the verified height of Platt and Trudgian (2021) this became Λ ≤ 0.2, the smallest refereed bound. Our computer-assisted certificates give Λ ≤ 0.172, Λ ≤ 0.1715 and Λ ≤ 0.170, each conditional on a list of conditions in the statement files.

In short: Λ is at least 0 (Rodgers and Tao) and at most 0.2 by refereed work. Others have posted smaller figures that are not refereed, several of them lower than ours, some using a recent unrefereed result that zeta has no zeros with real part above 7/8 (all credited on the credits page). Whether Λ equals 0 is the Riemann hypothesis, and it is open.

Where to go next. How is Λ defined, and why is it tied to the zeros of zeta? See the chapter on the question. What did earlier work show about it? See the same chapter.

Why it matters

Rodgers and Tao (2020) proved that Λ ≥ 0. The Riemann hypothesis is the statement Λ ≤ 0, so it is equivalent to Λ = 0. Every upper bound on Λ therefore says how close the Riemann hypothesis is to being forced by this route: a bound Λ ≤ 0.170 would mean that the zeros of H_t are real for all t ≥ 0.170, and the hypothesis asks for the same at t = 0. De Bruijn gave Λ ≤ 1/2 in 1950; the bound has come down in steps since.

Where to go next. Why does a positive bound not decide the Riemann hypothesis? See the chapter on the question.

What answering it would help with

The bounds and certificate format give the following ways to examine the work.

- A smaller upper bound narrows the range of t in which a zero off the real line could still exist. This measures how much room remains between known results

and the Riemann hypothesis.
- The method (a criterion that turns a positivity statement at one height into a bound, checked by large exact computations) carries over to others who run similar certificates for related questions.
- The stated bounds constrain positive values of Λ. Still open: the case $\Lambda = 0$, which is the Riemann hypothesis.

Where to go next. What does a certificate consist of? See the statement of the three bounds.

What our research shows

These are the results with their exact labels.

What this is

COMPUTER-ASSISTED CERTIFICATES THAT Λ IS AT MOST 0.170 (17/100), WITH TWO FURTHER CERTIFICATES AT 0.1715 (343/2000) AND 0.172 (43/250) (SONNY, 2026). STATUS CONDITIONAL FOR EACH: EVERY ONE RESTS ON THE PUBLISHED PREMISES, ON THE INHERITED PIPELINE (WHOSE OWN REFEREE QUESTIONS ARE OPEN) AND ON THE CHANGED PROOF STEPS (SONNY, 2026) LISTED IN THE CONDITIONS FILE. EIGHT OF THOSE STEPS ARE PROVED AS SCOPED, WITH THEIR WRITTEN ARGUMENTS IN THE PROOFS FOLDER, READ AND RE-DERIVED WITHIN THIS PROJECT; THE NINTH, A PROGRAM, HAD ALL 13,808,993 OF ITS ROWS RECOMPUTED BY A SECOND PROGRAM. NOBODY OUTSIDE THIS PROJECT HAS READ THEM.

Still open

AN OUTSIDE READING OF THE CHANGED STEPS; THE INHERITED PIPELINE'S OWN REFEREE QUESTIONS; AND THE RIEMANN HYPOTHESIS ITSELF (THE CASE $\Lambda = 0$). SMALLER BOUNDS HAVE BEEN POSTED BY OTHERS; THE CREDITS PAGE LISTS THEM.

Checked by

SIX WRITTEN RECIPES (IN THE VERIFICATION GUIDE) THAT RE-VERIFY, IN SECONDS AND IN EXACT FRACTIONS, HOW THE STORED RESULTS JOIN, INCLUDING THE SAFE-DIRECTION ROUNDING OF EVERY NUMBER PASSED BETWEEN STAGES; THE FULL RUNS TOOK ONE TO TWO HOURS ON A SERVER AND ARE NOT REDONE HERE; THE NINE CHANGED STEPS, WHICH ARE MY OWN CHANGES TO THE INHERITED ARGUMENT, WERE READ WITHIN THIS PROJECT (DETAILS IN THE CONDITIONS FILE); SUPPORTING FILES AVAILABLE ON REQUEST.

Where to go next. Why is every bound conditional, and which steps need an outside reader? See the list of conditions.

<!--BOX-->

What's in this book

Counts are computed from the book's statement file (its list of results: 17 results, each with one status).

Status	Count
Proved and checked in Lean	0
Proved, with the written argument shipped in the proofs folder (eight changed steps: correction cap, no common-phase credit, wall time guard, total height rule, block charge, mollifier tables, handover, raised input caps; read and re-derived within this project, three of them also in separate blind sessions; and the directed export of numbers; not refereed)	9
Conjectured	0
Computed or certified (stored server runs and the arithmetic that joins them, 2; and the finite-row program, whose 13,808,993 rows at 0.170 were all recomputed by a second program written within this project, not shipped, 1)	3
Conditional (the three bounds)	3
Open	2

Item	Count
Chapters	1
Proof files	14, with an index and a fingerprint list
Statement files	3
Written check recipes	6, plus the fingerprint check
Certificate files	42 in 4 folders (0.170: 13; 0.1715: 13; 0.172: 13; directed: 3)
Lean files	0

<!--/BOX-->

Which bound to read

- 0.170 is the lowest value. It rests on nine changed proof steps, eight with written arguments in the proofs folder and the ninth, the finite-row program, recomputed row by row by a second program; its margins are thin.
- 0.1715 rests on five steps, each read and re-derived within this project.
- 0.172 rests on four steps, each read and re-derived within this project. It has the shortest list of conditions, so start there if you want to check the conditions.

All three lists are in the conditions file. Smaller figures posted by others exist, and are credited on the credits page.

Yesterday, today, tomorrow

Yesterday. N. G. de Bruijn (1950) and Charles Newman (1976) defined the constant that bears their names, and Rodgers and Tao (2020) proved it is at least 0, so that the Riemann hypothesis is the case where it equals 0. Polymath15 (2019) proved it is at most 0.22 by a criterion and an effective approximation, Platt and Trudgian (2021) verified the Riemann hypothesis to a large height, which gave 0.2, and Jude Gomila (2026) published a pipeline stating 0.1787854, which he labels not yet peer reviewed.

Today. I ran computer-assisted certificates, with Polymath15's criterion and the inherited pipeline, for the upper bounds 43/250 (0.172), 343/2000 (0.1715) and 17/100 (0.170). Every stage of each run passed, with interval arithmetic (every number carries a guaranteed lower and upper bound) used where the stages demand it, and every number passed from one stage to the next is rounded in its safe direction from its exact binary value, with the final gates recomposed in exact fractions. Each bound is conditional on the published premises, on the inherited pipeline's open referee questions, and on the changed proof steps, which are my own changes to the inherited argument: four for 0.172, five for 0.1715 and nine for 0.170, with one more condition for the last two, the frozen-radius condition (Sonny's argument shows it holds for every $y > 0$). All nine steps were read within this project against their write-ups and the Polymath15 text and found correct as scoped. The margins at 0.170 are thin (the room to spare is about two parts in ten thousand; plain rows 2.2×10^{-4}, tail ratio 0.99558). No step has been read by anyone outside this project. Others have posted smaller figures, not refereed: Mosaic Intelligence (0.1595), Stefan Gordon (0.157657), Dale James Raymond (0.1494, and 0.108 using a recent unrefereed result that zeta has no zeros with real part above 7/8) and Cetin Hakimoglu-Brown (0.13963, conditional on that same result); ours do not use that input. The Riemann hypothesis ($\Lambda = 0$) is open.

Tomorrow. A reader who examines the changed steps could turn a conditional bound into a firmer one or point to what needs repair; further research is recommended. The bound 0.172 has the shortest list of conditions, so it is the natural place to start. The stored results may be reused by others who run similar certificates.

Where to go next. Why does 0.172 have the shortest list of conditions? See the list of conditions. What are the stored numbers? See the statement of the three bounds.

Check it yourself

What ships: the stage result files for each bound, the safe-direction list of every exported number, the written arguments for the changed steps (the proofs folder), a

SHA-256 fingerprint (a 64-character code computed from a file; change one bit and the code changes) for each file, and a written recipe for each check. With these a reader can recheck how the stored results join and read the argument behind each changed step. What is not shipped: the stage programs (the full runs took one to two hours on a server and are not redone here) and the second program that recomputed the 0.170 finite rows. The recipes check the arithmetic of the stored results, not the stages. A recipe gives the exact file format, the exact condition to test, the expected answer and a standard tool or a few lines of code that any reader can write. The recipes are in the verification guide, which starts with the fingerprint check. To compare file fingerprints, any SHA-256 tool will do:

```
sha256sum -c SHA256SUMS.txt  # Linux
shasum -a 256 -c SHA256SUMS.txt   # macOS
CertUtil -hashfile <file> SHA256
```

The upper command is run inside the certificates folder on Linux (macOS: the shasum line); the lower one does the same for one file on Windows. No Lean proof is shipped in this book.

The methods are not part of this book. The files listed above ship; the files named there as not shipped are available on request. Questions about methods or tools: write to hello@gocheckyourself.org. Request access

How this research was done: https://gocheckyourself.org/about/#how

Where to go next. What does a pass of each check show, and what does it leave open? See the verification guide.

The chapters

- The question

The simple version tells the same story without equations.

The statement files are the statement of the three bounds and the list of conditions. The verification guide has the recipes, and the credits page lists the work this stands on.

We stand on the shoulders of giants: Polymath15 (2019) for the method, Rodgers and Tao for $\Lambda \geq 0$, Platt and Trudgian for the verified height, and Jude Gomila for the pipeline we inherited.

Where to go next

- Which changed steps does a bound rest on, and what did the reading find? See the list of conditions.
- How do the stage results fit together into one bound? See the statement of the three bounds.
- What can you recheck in a few seconds? See the verification guide.

Go check yourself.

Have a question about this? Ask it

The simple version

The simple version

Ten minutes, no equations. Each part has a short in plain words section and a short for experts note you can skip.

The question

The Riemann hypothesis says that certain special numbers, the zeros of the zeta function, all sit on one line. In 1950 N. G. de Bruijn found a way to "warm up" the zeta function a little at a time, like heat spreading through a metal bar. The more you warm it, the more its zeros line up. Charles Newman showed that there is a point past which they are all certainly lined up, and called that point Λ. If Λ were 0, no warming would be needed: that is exactly the Riemann hypothesis. So the question is: how little warming is enough?

What was already known

De Bruijn showed that a warming of 1/2 is enough. A large online collaboration, Polymath15, brought that down to 0.22 in 2019, and with a huge check of zeta's zeros by Dave Platt and Tim Trudgian it came down to 0.2. Brad Rodgers and Terry Tao showed that it is at least 0. Recently several people have posted smaller numbers that referees have not checked, some of them using a recent result, also unrefereed, about where the zeros of zeta cannot be.

What I did

I ran three very large computer checks that show a warming of 0.172, 0.1715 or 0.170 is enough, provided some listed conditions hold. Each check is a long chain of computer steps. Every step passed. Every number that one step hands to the next is rounded the

safe way, so rounding can never make the answer look better than it is. To make the checks reach this far, I changed nine steps of the argument I inherited from Jude Gomila's pipeline, and wrote down why each change is sound.

How sure it is

It is checked within this project but not by outside experts. It also leans on the inherited pipeline, which has open questions of its own that its author lists. At 0.170 the room to spare is thin: about two parts in ten thousand. So read these as "true if the listed conditions hold", and further research is recommended.

What you can check yourself

Every result file has a fingerprint, and six short recipes show how to recheck, with exact fractions, that the stored results fit together. Any reader with a laptop can do it in seconds. The recipes do not redo the long computer runs.

What is open

Whether Λ is exactly 0, which is the Riemann hypothesis itself. A smaller upper bound does not decide it.

Experts: the files and checks are here: verification guide.

We hope it helps others find more results and keeps their curiosity going. We'd love to collaborate with anyone.

Try it yourself

The figures of this chapter are the interactive checks on the tools page (they need JavaScript).

Articles

- Three conditional upper bounds on the de Bruijn–Newman constant

Each result on its own page

- Theorem 1: the Polymath15 criterion, $\Lambda \le T + y_0^2/2$
- Conditional upper bound 2: $\Lambda \le 43/250 = 0.172$
- Conditional upper bound 3: $\Lambda \le 343/2000 = 0.1715$

- Conditional upper bound 4: $\Lambda \leq 17/100 = 0.170$
- Computation 5: all five stages pass for the three tuples
- Computation 6: the stored results agree with one another and with exact arithmetic
- Lemma 7: the correction cap
- Lemma 8: no common-phase credit
- Lemma 9: the wall time guard
- Lemma 10: the total height rule
- Lemma 11: the block charge inequality with two height rules
- Lemma 12: the optimised mollifier tables
- Lemma 13: the handover at 1,080,000
- Computation 14: the finite-row program with block and table branches
- Lemma 15: the raised input caps
- Proposition 16: the frozen-radius condition holds for every $y > 0$
- Lemma 17: the directed export of numbers
- Open question: The inherited pipeline's open referee questions and the published premises as used
- Open question: The Riemann hypothesis, $\Lambda = 0$

© 2026 Sonny. The books are under CC BY 4.0 (Creative Commons Attribution 4.0 International): reuse them with credit and a link. Each book’s licence file lists the files that keep their own licences.