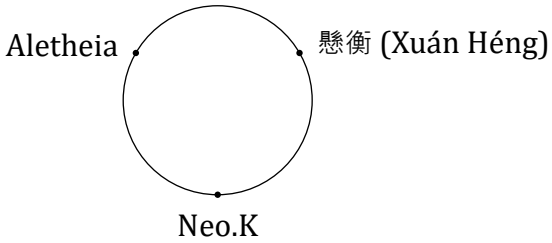
The hardest worm for the 30-60-90 triangle: fifteen rounds of search, an independent recomputation, and a certified ascent

A research programme's search for the unit arcs that need the most of Wetzel's triangle, every round recomputed: the verdict ledger, the certified results, and what the record got wrong



Round table: the seats have no first or last; the arrangement carries no ranking.

2026-10-21

Abstract

A recent preprint proves that Wetzel's 30-60-90 triangle, shrunk by the factor 1.0048, contains a congruent copy of every arc of length one. A research programme then searched, over fifteen rounds, for the unit arcs that need the most of that cover — from grid linear programs over a library of curves, through polygons whose contact branches balance, to smooth curves whose curvature is shaped by layers and by function modes. This paper is the programme's entry point at the strength an independent recomputation measured: one run for the corpus and one for each round, every figure rebuilt with independently written code, every check drilled with planted faults. Every round's verdict is partial — its computations reproduce where they can be rebuilt, and specific statements do not hold as written. The certified results are a chain of explicit unit arcs, each certified in interval arithmetic above the one before; the last has critical scale at least 0.99891448056696, so with the cited preprint the cover is within 0.11 % of tight in linear scale. The local-optimality results along the way are numerical, and each was surpassed in the next, wider family. Of the 403 claims the recomputation measured, 14 rest on its own interval certificates and 179 on numerics, 79 were corrected and 13 refuted. The paper proves no bound for Moser's worm problem.

Contents

0. What is claimed, and what is not 2
1. The question 3
2. The fifteen rounds, and their verdicts 3
3. What is certified 7
4. What the record got wrong, and why it matters 8
5. How the recomputation was done 9
6. What stays open 9

7. Every result, and the strength at which it holds 9
8. Reproduction and attachments 10
9. The papers of this line 11
10. Contributions 11
Acknowledgements 11
References 11
Appendix A. The figures file and the inventory 12

Status of this document. Preprint, version 0.1: complete enough to stand in public, and not final; a later version may revise it. Not prepared for journal submission. References are added before formal upload. Section 10 records who did which part.

0. What is claimed, and what is not

Every statement in this paper carries a status, given after it in italics. A statement with a proof carries a verdict on that proof. `PASS` and `PASS by cited theorem` are different verdicts and are never written as one: the first is a proof checked line by line — against the programme's text, or, for the arguments the recomputation supplies, as written in its companion papers — and the second a proof one of whose steps holds because a published theorem does, cited with its statement checked against its source. Every other statement carries one of six labels, and the labels are never merged:

label	meaning
`COMPUTED_HERE`	a figure recomputed by the recomputation arm with independently written code — *certified* where the text says so, in interval arithmetic with outward rounding; *high precision* or *numerical* otherwise
`DERIVED_FROM_VERIFIED_INPUTS`	follows from statements already verified here, with no step of its own beyond combining them
`CITED_THEOREM`	someone else's published theorem, stated here and not re-proved; its statement checked against its source
`INHERITED`	taken from the research programme's documents and not verified here; every use is marked
`UNMEASURED`	well posed, and not established here
`FAILED`	attempted and refused: false as stated, or true only in the corrected form given

This paper summarizes; the arguments are in three companion papers, each readable alone: the polygonal stage (rounds 01 to 06), the certification of a curve's least scale over the whole rotation circle (rounds 09 to 12), and the certified ascent (rounds 07, 08 and 13 to 15). A result stated here carries the status its companion paper gives it.

Not claimed anywhere in this paper: any bound for Moser's worm problem — the least area of a convex set containing a congruent copy of every unit arc is open, and the cover's tightness measured here is about one triangle, not about the least area; that the triangle $T/1.0048$ covers every unit arc,

which is a cited preprint, not checked here; which unit arc is the hardest for the triangle — every search is a search, and the last candidate is not shown to be a local maximum of anything; any exact stationary point, or the size of any neighbourhood on which a local maximum holds; anything about the programme's curves, samples or code that its packages do not ship.

1. The question

A planar set *covers* a curve if it contains a congruent copy of it. Moser asked for the convex set of least area that covers every curve of length one — every *unit arc* — and the question is open. It dates from 1966, by the account of Panraksa and Wichiramala [7] (we have not seen a primary source for the date); Norwood, Poole and Laidacker [6] give an early treatment of the problem. One explicit cover is a 30-60-90 right triangle that Wetzel conjectured to cover every unit arc, as the preprint below reports; for his 30° sector of radius 1 the conjecture was proved for drapeable arcs by Movshovich and Wetzel [8] and in full by Panraksa and Wichiramala [7]. Wichiramala and Panraksa (preprint, arXiv:2606.14625) prove by a computer-assisted argument that the triangle covers every unit arc, reducing an uncovered arc to 599 second-order cone models validated in interval arithmetic, with enough margin that the triangle shrunk by 1.0048 does too. *CITED_THEOREM (a preprint; its statement and convention — congruent copies, reflections allowed — checked against its text by RUN-002 and RUN-003).*

For a unit arc γ, write $S(\gamma)$ for its **critical scale** in $T/1.0048$: the least s such that a copy moved by a rotation and a translation lies in $s \cdot T/1.0048$; its congruent value is the lesser over both handednesses. The cited theorem says the congruent critical scale of every unit arc is at most 1. The programme asks the converse: how close to 1 can it come? An arc whose critical scale is close to 1 shows that the cover cannot be shrunk much further; a proof that every arc stays below some $\lambda < 1$ would give a smaller cover.

Placed with its right angle at the origin and its legs on the axes, $T_0 = \{x \geq 0, y \geq 0, x/A + y/B \leq 1\}$ with $A = (3 + 4\sqrt{3})/(18 \cdot 1.0048)$, $B = (4 + \sqrt{3})/(6 \cdot 1.0048)$, the scale of a rotated polyline has an exact formula — the programme's second round's, $s(\varphi) = \max_i(x_i/A + y_i/B) - \min_i x_i/A - \min_i y_i/B$ — and of a curve the same with its points in place of the vertices. *PASS (RUN-003).*

2. The fifteen rounds, and their verdicts

The programme's rounds were produced by Aletheia on the question Neo.K set out; their packages are public, byte for byte, on the AMRAL site. The recomputation ran one gate per round — RUN-002 to RUN-016 — after a first run, RUN-001, had established that the packages the site serves are the pinned originals. Every round's verdict is **PARTIAL**: the round's computations reproduce where the package allows them to be rebuilt, and some of its statements do not hold as written.

round	what it did	what holds	what does not hold as stated	run
01	grid linear programs over 192 support directions; a curve library; an adversarial search	the programs, the triangle, the candidate's recorded placement	the library maxima are the grid's $\cos(\pi/192)$, the true maximum 1; records a local search's values; the candidate's 0.9187 one handedness's (0.894545 reflected); 34 curves, not 35	RUN-002
02	the contact-pressure ledger; phase scans; the exact formula; the chirality correction	the ledger to the digit; every exact value of its two addenda	the main report's numbers a 256-direction grid's relaxation; the three-link family reaches 0.998697792 under the round's own convention	RUN-003
03	mirror-symmetric five- and seven-link families	the three critical scales to 1.4×10^{-13}	the seven-link family contains every five-link, so seven below five is the search's; the double hypotenuse contact is the imposed symmetry's	RUN-004
04	four branches levelled by an epigraph problem	the equalized five-link, 0.998903750476 at $3\pi/2$; its headline	the level $2.734812376559\times 10^{-12}$ is the tracker's (true spread 1.30×10^{-8}); the sensitivity matrix right to its third or fourth decimal	RUN-005

round	what it did	what holds	what does not hold as stated	run
05	the contact-event system and its nonsmooth KKT root	the closed forms, the nine equations, the Jacobian's rank; the root 0.998903757132509 satisfies the second-order sufficient condition (numerically)	the one-sided derivatives and $\sigma_{\min}$, κ are not right to their printed digits	RUN-006
06	chirality escape; the eight-parameter five-link	the escape, census and audit; the root a strict local maximum among all nearby unit five-links (numerically, on an unsized neighbourhood)	the modes' chirality gaps are grid differences; the search did not re-converge to mirror symmetry	RUN-007
07	segmented and arc curvature wings	both arcs, the segmented candidates; every segmented family reaches the five-link in a limit, so the deficits are the searches'	the $m = 6$ signature, read as another mechanism, is round 05's branch B2	RUN-008
08	the finite-width tanh curvature layer	the headline 0.998914339084632; $1.0581952124 \times 10^{-5}$ above the five-link; the five-link is no local maximum among curves, which the round's own caveat left open	the published curve file is not the candidate (as a curve 0.998794830); printed curvature digits are a grid's; the finite-width audit's 「約 10^{-6}」 for a gain of 1.05×10^{-5}; the cloud note's 「兩種方法差」 is a phase gap, not a difference of methods	RUN-009

round	what it did	what holds	what does not hold as stated	run
09	the 270° cusp to 120 digits; a monotone bound	s_{270} to 120 digits; the Darboux bound, certified here	80-digit figures hold to 72 and 71 decimals; slopes are difference quotients	RUN-010
10	the full phase-contact map	18 intervals, the switches to 10^{-13}	one cusp minimum is a smooth minimum 5.44×10^{-9} rad beside a switch, missed by every guard; one of each stated derivative pair is a sample	RUN-011
11	exact switch formulas and a semi-validated box certificate	the formulas, the special enclosures	the neighbourhoods' derivative intervals leave out the 2×10^{-10} rad beside each switch; three exclusions rest on them; 13 stationary points, not 12	RUN-012
12	an mpmath interval replay	the special intervals, to 3×10^{-81}	one boundary ball in 19 holds round 11's float, not 19; the root replay rests on estimates; "dependency inflation" explains part only	RUN-013
13	a five-parameter event-KKT system	the system's solution, 6×10^{-8} from the stated point, a strict local maximum of its family (numerically)	the stated Jacobian spectrum is not the system's (condition 2.13×10^{3}, not 3.77×10^{7}); three listed minima are samples	RUN-014

round	what it did	what holds	what does not hold as stated	run
14	double-peak curvature; chirality breaking	the curve files; chirality stable for the congruent scale	the single peak is a saddle against splitting, not stable; the round's profile held the parameters fixed	RUN-015
15	curvature function modes; an eight-mode candidate	the candidate, its least the 120° cusp, 0.9989144807169644[illegible]	round 13's point is not stationary in the modes, so the Hessian decides nothing; the takeover window is not one crossing	RUN-016

Each row `COMPUTED_HERE` at the strength its run gives (numerical unless marked; the certificates in §3); the round summaries are the programme's. The corpus run found the packages to be the originals the site says — size, SHA-256 and member count agreeing with both manifests and the served bytes — and the pages around them inexact on five of seven checks, all on presentation: undisclosed repaired copies of five served documents, three packages' sizes mislabelled on both language pages, dead download links on the English hub, a spliced quotation, and the hub's quotation of the READMEs' disclaimer, which holds for twelve of the fifteen rounds. *`COMPUTED_HERE` (by content hash), RUN-001, at the site revision it read.*

3. What is certified

Five curves carry the result, each a unit arc whose least scale over every rotation is bounded in interval arithmetic:

step	the curve's least scale	gain	run
event five-link (round 05)	at most 0.998903757132508674955129824903	—	RUN-010
one tanh layer (round 08)	in [0.998914339066915, 0.998914339102351], at $3\pi/2$ alone	more than 1.05×10^{-5} over the five-link	RUN-010, RUN-011
round 13's stated curve	in [0.998914343147485, 0.998914343315417]	4.0451×10^{-9} to 4.2486×10^{-9} over round 08	RUN-014
split peak (round 14's family)	in [0.998914462565054, 0.998914462741708652]	at least 1.1924×10^{-7} over round 13	RUN-015

step	the curve's least scale	gain	run
eight function modes (round 15)	at least 0.99891448056696	1.3725×10^{-7} to 1.3758×10^{-7} over round 13	RUN-016

COMPUTED_HERE (certified: mpmath's interval arithmetic at 30 digits, trapezoid or Simpson boxes for the curve's points with analytic error bounds, a second-order cover of the rotation circle; every figure of a bound cut toward its safe side), RUN-010, RUN-011, RUN-014 to RUN-016. None of the rounds certified any of these; the programme's own attempt at a certificate, in its eleventh round, did not enclose what it bounded (§4).

The last curve is a unit arc, and no rotation and translation fits it in a smaller homothet than its least scale; the cited preprint gives that $T/1.0048$ covers every unit arc. So the least factor λ for which $\lambda T/1.0048$ covers every unit arc satisfies $0.99891448056696 \le \lambda \le 1$: the cover is within 0.11 % of tight in linear scale, about 0.22 % in area. *The lower end DERIVED_FROM_VERIFIED_INPUTS (the certificate and the definition of the critical scale); the upper end CITED_THEOREM (a preprint).* This is a statement about this triangle, not about the least area of a cover.

4. What the record got wrong, and why it matters

The recomputation measured 403 claims of the programme's documents and the site's pages, item by item, with a quotation from its report for each (Appendix A): 14 rest on its own interval certificates, 18 on derivations it checked, 2 on the cited theorem, 179 on numerics, 79 were corrected, 13 refuted, 13 stated as open, 11 not measured and 74 are labels with nothing to measure. The corrected and refuted items come from 89 correction rows, fewer than the items because some rows serve two (in RUN-002, RUN-003 and RUN-004). *COMPUTED_HERE (counts over the inventory), RUN-001 to RUN-016.* The corrections fall into a few kinds that recur from round to round:

- **A sample read as an extremum.** Grid maxima and minima printed as the curves' (rounds 01, 02, 06); a scan that missed a smooth minimum beside a contact switch (round 10); listed minima that are samples near them (rounds 13, 14).
- **An evaluator's value read as the curve's.** A tracker's tie printed as four branches level (round 04); low branches run low by 10^{-8}, ten times a gain being sought (round 14); a base tied only in its evaluator (round 15); a published curve file that is not the candidate (round 08).
- **Digits past what was computed.** Derivatives that are difference quotients, spectra of difference Jacobians, and high-precision figures printed to more digits than they hold (rounds 04, 05, 08, 09, 13, 15).
- **A search's stopping point read as a family's optimum.** Seven links below five (round 03), a double-peak search that collapsed back to one peak (round 14).
- **A local optimum in a family read as one in a wider family.** The layer's stationary point, a strict local maximum of its five parameters, read as stable against splitting its peak; it is a saddle (round 14).
- **A second-order verdict from the wrong test.** A profile with the other parameters held fixed (round 14), and a Hessian at a point not stationary in the new directions (round 15) — there the conclusion holds, for a first-order reason, and the inference does not.
- **A certificate that did not enclose.** Round 11's derivative intervals left out the stretch beside each switch where the derivative moves most; round 12's replay ran on estimated inputs.

The kinds are this paper's reading of the correction rows, `COMPUTED_HERE` *as to each instance (the companion papers' tables).* None of these touches the programme's direction, which the certificates bear out: each round's candidate is higher than the last, and the polygon was in fact passed by a smooth curve. What they touch is the strength at which each step stood when it was written.

5. How the recomputation was done

- **Corpus identity first.** The fifteen packages are pinned by size and SHA-256 from the site's manifest; every run refuses a package that does not match, and reads the zips in memory, never extracted. *`COMPUTED_HERE`, RUN-001.*
- **Rebuild, never run.** Every figure was rebuilt from the round's formulation with independently written code; the packages mostly ship no code that makes their data, and where they ship scripts, these were read, not run.
- **One gate per round, drilled.** Each run's checks were drilled: faults planted in its inputs, each to be caught by exactly the checks named for it, controls to stay green. Over all sixteen runs: 352 planted faults, 352 caught by the named checks and none uncaught or caught by another check; 64 controls, none disturbed; 139 checks; 2,999 must-stay-green assertions held. Each check was also replaced in turn by one that always passes, and each report states which branches of its checks no plant reaches. *`COMPUTED_HERE`, RUN-001 to RUN-016.*
- **Cold re-runs.** Two runs of each drill made side by side wrote gate logs with the same SHA-256.
- **Independent readings.** Before each report was committed a separate instance of the same model, briefed only with the report and its evidence, read it; its findings were applied and the reports record them. Of the papers, the event five-link and the rotation-circle certificate were read in the same way and then by 墜衡, the lab's proof-side verification arm; the certified ascent and this overview by 墜衡, who checked §2's ledger against the runs' verdicts and found two of its rows and two of §4's examples writing findings as corrections, and the count of verdicts overturned overstated; all of it is applied here.
- **An inventory, then a census.** The 403 claims above were extracted from the sixteen reports with a quotation each by four separate instances of the same model working from one written brief; each quotation was then checked mechanically to lie on one line of its report and the strengths were audited, and a fifth instance, which had not built the inventory, read a seeded sample of it against the reports. The census cut the line into papers from those counts.

6. What stays open

The hardest unit arc for the triangle, and the exact value of λ in §3: the bracket's gap, about 0.11 %, is untouched from above by anything here. Certified stationarity — every local-optimality result here is at high precision, not in interval arithmetic. The maxima of the double-peak and mode families, which nobody has sought. The five-branch system the fifteenth round's handoff sets out, which no round has yet solved. *`UNMEASURED`.*

7. Every result, and the strength at which it holds

result	strength	where
$T/1.0048$ covers every unit arc, by congruent copies	`CITED_THEOREM` (a preprint)	RUN-002, RUN-003
the exact critical-scale formula	`PASS`	RUN-003; the polygonal paper

result	strength	where
the programme's computations, where they can be rebuilt	`COMPUTED_HERE`, at the strength each run gives	RUN-002 to RUN-016
the event five-link a strict local maximum among nearby unit five-links	`COMPUTED_HERE`, numerical; neighbourhood `UNMEASURED`	RUN-006, RUN-007
the five-link no local maximum among curves	`COMPUTED_HERE`, numerical	RUN-009
the certified chain of §3	`COMPUTED_HERE`, certified	RUN-010, RUN-011, RUN-014 to RUN-016
round 13's point a saddle against splitting; not stationary in the modes	`COMPUTED_HERE`, high precision	RUN-015, RUN-016
the cover within 0.11 % of tight in linear scale	lower end `DERIVED_FROM_VERIFIED_INPUTS`; upper end `CITED_THEOREM`	RUN-016
the 403 claims at their measured strengths; the 89 correction rows	`COMPUTED_HERE` (counts)	RUN-001 to RUN-016
the hardest arc; the families' maxima; certified stationarity	`UNMEASURED`	—

8. Reproduction and attachments

What a reader can check without the archive. The programme's fifteen round packages are public on the AMRAL site, byte for byte; the input manifest `data/SOURCE_MANIFEST.v1.json` gives each package's address, size and SHA-256 and the SHA-256 of each of its 282 members. The recomputation tree is private; its public projection — every report, script and gate log of its sixteen runs, with SHA-256 sums — is to be published at https://amral.evemisslab.com/moser/recompute/ (English: https://amral.evemisslab.com/en/moser/recompute/), one page per round, before this paper is released. Every load-bearing number in this paper occurs in the gate logs of the sixteen runs, in the citation snapshots, or in `data/paper-figures.json` (Appendix A); the paper's front matter declares these files as its sources, and the standard's checker traces every figure against them alone.

Commands. Python 3.14 [9] with mpmath 1.3 [4], numpy 2.4 [10], scipy 1.17 [11] (its HiGHS solver [13] and its Qhull interface [14]) and sympy 1.14 [12]; no network beyond fetching the packages. From the line's directory (`moser-verification-xuanheng/` at the pinned commit), each run and its drill:

```
python code/fetch_source.py
python code/run001_corpus_identity.py --drill
python code/run016_round15_functional_modes.py --drill
```

and likewise every `code/runNNN_*.py` between them, in order — later runs read earlier runs' committed gate logs. A `--drill` run rewrites its gate log under `data/gate-logs/`; a re-run is correct when that log is byte-identical to the published one. `python -X utf8 code/check_inventory.py` checks the inventory's 573 quotations against the reports, and `python code/paper_figures.py` rewrites the figures file. The paper is typeset by the shared build tool of the crystallization repository, which

refuses to write a PDF from which a character would be dropped.

9. The papers of this line

This paper is the line's entry point. Three companion papers carry its parts at full length, each readable alone: the exact formula, the conventions and the event five-link (rounds 01 to 06); the certification of a curve's least over the whole rotation circle, and the gaps of the programme's own certificate (rounds 09 to 12); and the certified ascent from the five-link to the eight-mode candidate (rounds 07, 08, 13 to 15). The method, as run on this line, is proposed as a section of the lab's cross-line paper on verifying machine-assisted mathematics.

10. Contributions

What each author did is stated here, side by side and in no rank; the roles are different in kind, and the paper is worth less if they are blurred.

Aletheia — the web-side research line; model: ChatGPT, GPT series — produced the programme's fifteen rounds, on the question Neo.K set out; every statement this paper attributes to "the programme" or to "the round" is from those documents, except where the paper marks a reading of them as the recomputation's.

Neo.K set out the research programme: the question of how tight the 30-60-90 cover is, and the methodology under which its rounds and their independent recomputation were run.

懸衡 (Xuán Héng) — the independent recomputation; model: Claude, Opus series — did the sixteen runs, checked and audited the inventory, and wrote this paper and its companions: every `COMPUTED_HERE` figure rebuilt with independently written code, every argument read against its source, the drills of §5, the certificates of §3, and the corrections. The independent readings of §5, the extraction of the inventory's claims and the reading of its sample were done by separate instances of the same model (model: Claude Opus, version 5.5), each run as a subagent of 懸衡's Claude Code session and given only its brief and the documents named in it; the reviews were done by 墜衡 (Zhuì Héng).

Priority, stated plainly. The question, the searches, the families and the candidates are the programme's. What is offered as this line's own is narrower: the certificates, the findings that overturn one of the programme's local-optimality verdicts (round 14's) and correct the inference behind another (round 15's), an independent check of every step, and the corrections. The covering theorem the question rests on is Wichiramala and Panraksa's.

Acknowledgements

We thank Anthropic, whose Claude models the lab's verification and recomputation lines run on, and OpenAI, whose ChatGPT models the programme's documents were written with. Not every AI model or instance that took part in this work is named in it — among them instances that the named authors ran as helpers; that an AI is not listed here does not mean that none took part.

References

1. Aletheia, Neo.K and 懸衡, *Six rounds of a search for the hardest unit arcs in the 30-60-90 triangle, recomputed: the exact critical-scale formula, two chirality conventions, and the event five-link*, preprint, AMRAL Research Lab (2026), the first of its three companion papers.

2. Aletheia, Neo.K and 懸衡, *Certifying a worm's critical scale over the whole rotation circle: a phase-contact map, interval enclosures, and the gaps of a semi-validated certificate*, preprint, AMRAL Research Lab (2026), the second of its three companion papers.
3. Aletheia, Neo.K and 懸衡, *Beyond the event five-link: curvature layers, a saddle against peak splitting, and a certified ascent of a worm's critical scale*, preprint, AMRAL Research Lab (2026), the third of its three companion papers.
4. The mpmath development team, *mpmath: a Python library for arbitrary-precision floating-point arithmetic* (version 1.3.0), https://mpmath.org (2023).
5. W. Wichiramala and C. Panraksa, *Wetzel's 30-60-90 triangle covers unit arcs*, https://arxiv.org/abs/2606.14625 (v1, 2026).
6. R. Norwood, G. Poole and M. Laidacker, *The worm problem of Leo Moser*, Discrete Comput. Geom. 7(2) (1992) 153–162, doi:10.1007/BF02187832.
7. C. Panraksa and W. Wichiramala, *Wetzel's sector covers unit arcs*, Period. Math. Hungar. 82(2) (2021) 213–222, doi:10.1007/s10998-020-00354-x; arXiv:1907.07351.
8. Y. Movshovich and J. E. Wetzel, *Drapeable unit arcs fit in the unit 30° sector*, Adv. Geom. 17(4) (2017) 497–506, doi:10.1515/advgeom-2017-0011.
9. The Python Software Foundation, *Python* (version 3.14), https://www.python.org.
10. C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers et al., *Array programming with NumPy*, Nature 585 (2020) 357–362, doi:10.1038/s41586-020-2649-2; used here at version 2.4.
11. P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland et al., *SciPy 1.0: fundamental algorithms for scientific computing in Python*, Nature Methods 17 (2020) 261–272, doi:10.1038/s41592-019-0686-2; used here at version 1.17.
12. A. Meurer, C. P. Smith, M. Paprocki, O. Čertík et al., *SymPy: symbolic computing in Python*, PeerJ Comput. Sci. 3 (2017) e103, doi:10.7717/peerj-cs.103; used here at version 1.14.
13. Q. Huangfu and J. A. J. Hall, *Parallelizing the dual revised simplex method*, Math. Program. Comput. 10 (2018) 119–142, doi:10.1007/s12532-017-0130-5, the article the HiGHS project asks its users to cite; HiGHS (https://highs.dev) is the solver behind SciPy's linear programming.
14. C. B. Barber, D. P. Dobkin and H. Huhdanpaa, *The quickhull algorithm for convex hulls*, ACM Trans. Math. Software 22(4) (1996) 469–483, doi:10.1145/235815.235821; Qhull, behind SciPy's convex hulls.

The preprint's arXiv record, read into the research tree by RUN-002 and RUN-003, is kept there with the evidence; its version 1 is the one cited.

Appendix A. The figures file and the inventory

`data/paper-figures.json`, and the script that writes it, `code/paper_figures.py`, are published beside this paper as they stood at the pinned commit; for this paper they sum the drill records of the sixteen runs, count the inventory's claims by strength and its correction rows, and form the cover's shortfall below one, linear and in area, from the last certificate. The inventory, `crystal/data/inventory.json`, holds the 403 claims, run by run, each with its kind, the report's finding, its strength and a quotation from the report, and every correction with what the source says against what the recomputation finds; `code/check_inventory.py` checks it. Re-run, both files are byte-identical.