PHI TECHNICAL MONOGRAPHS

Conformal Structures and Function Spaces

Algebra, topology, analysis and quantitative examples



PHI research collective

Abstract

This volume develops a connected sequence of algebraic, topological and functional-analytic constructions: finite symbolic coding, Cantor geometry, observable approximation, branched coverings, normal families and conformal distortion. The primary mathematical reference is Mikhail Lyubich's *Conformal Geometry and Dynamics of Quadratic Polynomials*, specifically §1.1.1, §3 and §4. Definitions, derived examples, repairs of ambiguous statements and quantified extensions are identified within that scope.

Finite representations are related to their limiting objects through explicit convergence estimates. Local multiplicity is connected to global degree and monodromy; orthogonal projection and operator norms quantify the information retained by finite observables. The worked examples examine degenerating branches, discrete symmetry, parameter-dependent continuation and finite-rank optimality. Standard prerequisites are declared where their full development lies outside the selected chapters.

The companion volumes treat algorithmic, statistical and physical applications with their own input and output specifications. Repeated appearances of a theorem in those applications are cross-references to a common result, not additional claims of mathematical novelty.

Reproduction material. The public source package contains the LaTeX manuscript and authored vector assets required to rebuild this volume. Its executable supplement covers selected numerical studies and certificate calculations. External textbook page reproductions and original externally acquired datasets are not distributed. References identify the scholarly works needed for independent source comparison.

Evidence convention. Established mathematical results retain their scholarly attribution. The exposition states the assumptions of each calculation and comparison. Recorded numerical results retain their original experimental scope. Formal specifications describe prospective proof developments unless an accepted proof object is explicitly identified. This public edition does not establish new proof-assistant kernel acceptance.

Contents

Abstract **i**

1 Symbolic coding, Cantor geometry and observable spaces **1**
1.1 Algebraic rederivation: words, affine maps, and exact coordinates 1
1.2 Analytical and geometric rederivation 4
1.3 Functional analysis: approximation spaces, measures, and operators 9
1.4 Topological rederivation and source Exercise 1.2 12
1.5 Algebraic topology: invariants that distinguish the right things 14
1.6 Computer-science application I: exact representations and decisions 16
1.7 Computer-science application II: automata, safety, and entropy 17
1.8 Computer-science application III: certified integration and optimization 19
1.9 Computer-science application IV: topology of finite observations 21
1.10 A bridge back to quadratic dynamics: a certified real repeller 22

2 Proper holomorphic maps, degree and lifting **24**
2.1 Local analytic models and proper maps 24
2.2 Riemann–Hurwitz from cells, divisors and energy 29
2.3 the topological argument principle 33
2.4 an explicitly supplied degree subsection 36
2.5 lifting is a compatibility condition 37
2.6 Galois coverings and controlled symmetrization 40
2.7 telescoping with explicit composition domains 44

3 Multiplicity, conditioning and compositional estimates **48**

4 Normal families, conformal mapping and distortion **56**
4.1 Reference §4.1: Little Montel Theorem 56
4.2 Reference §4.2: the Riemann Mapping Theorem 58
4.3 Reference §4.3: spherical normal families 61
4.4 Reference §4.4: normalized univalent functions and distortion 67
4.5 Annular collars and source Lemma 4.17 75
4.6 Additional functional and projective derivations 79
4.7 Completion of the radial-slit extremal argument 84

5 Constructive examples in algebraic geometry and topology **88**
5.1 Local multiplicity without the vocabulary barrier 88
5.2 Euler characteristic as a branch-count ledger 89
5.3 Boundary information can count an interior 90
5.4 Quotient algebras: an algebraic-geometric object you can compute 91
5.5 Following a root is a path problem 91
5.6 Composition: a formula is not its domain 92
5.7 MG1: finite fibres as algebras, including nilpotents 93
5.8 MG2: Hardy interpolation detects coalescing data 93
5.9 MG3: an explicit differential basis on a double cover 94
5.10 MG4: a finite cover defines an averaging adjoint 95
5.11 MG5: the complete data of a branched cover 95
5.12 MG6: a regular cover much smaller than the factorial bound 96
5.13 MG7–MG8: analytic operators and local inverse certificates 96
5.14 MP1: a discriminant margin gives a separation bound 97
5.15 MP2: a branch collision has a logarithmic period scale 98
5.16 MP3: Fredholm index survives norm perturbations, not arbitrary limits 98
5.17 MP4: a parameter-uniform continuation patch 100
5.18 MP5: the normal core and changes of branch generators 100
5.19 MP6: telescope certificates with nonuniform constants 101

6 Quantitative conformal examples **102**
6.1 MU01. Geometric scale in a derivative estimate 102
6.2 MU02. A finite polynomial with an explicit remainder 103
6.3 MU03. Local uniform convergence 103
6.4 MU04. Infinity as a target coordinate 104
6.5 MU05. The algebra behind the extremal mapping argument 105
6.6 MU06. Solving the Koebe inverse equation 106
6.7 MU07. Clearance and the derivative at the centre 107
6.8 MU08. Normalization prevents degeneration 108

7 Analytic spaces and extremal arguments **110**
7.1 MG01. The topology on a space of meromorphic maps 110
7.2 MG02. Jets, local algebras, and analytic realization 111
7.3 MG03. Attainment in the extremal proof 111
7.4 MG04. Cross-ratios and colliding target sections 112
7.5 MG05. Area and coefficient inequalities 113
7.6 MG06. The Schwarzian as a projective differential invariant 114
7.7 MG07. The Bergman pullback is unitary 115
7.8 MG08. A matching lower bound for finite-rank approximation 116

8 Parameter families and optimal approximation **117**
8.1 MP01. Quantitative total boundedness of analytic restrictions 117
8.2 MP02. Schatten scales and boundary degeneration 118
8.3 MP03. A length-area proof of an explicit collar bound 119
8.4 MP04. Moving-target separation and local conditioning 119
8.5 MP05. Global injection under analytic perturbation 120
8.6 MP06. A bounded observation model and an attained optimum 121
8.7 MP07. Conformal kernels and geometric observables 122
8.8 Undergraduate comparative workshop: calculate, explain, challenge 123
8.9 Graduate comparison: quotient, Hilbert geometry and covering data 126
8.10 Research comparison: theorem strength, optimality and nonuniform limits 129

References **134**

Chapter 1

Symbolic coding, Cantor geometry and observable spaces

1.1 Algebraic rederivation: words, affine maps, and exact coordinates

1.1.1 The construction as a free-monoid action

Let $\{0,1\}^*$ be the free monoid of finite binary words, with concatenation uv and identity $\varnothing$. Define affine contractions

$$S_0(x) = \frac{x}{3}, \qquad S_1(x) = \frac{x+2}{3}. \tag{1.1}$$

For $w = b_1 \cdots b_n$, set $S_w = S_{b_1} \circ \cdots \circ S_{b_n}$. Thus $S_{uv} = S_u \circ S_v$: the rightmost map is evaluated first, while the leftmost digit is the most significant address digit. This convention is important; reversing function-composition order reverses the address.

Define the integer numerator

$$A(w) = \sum_{j=1}^{n} 2b_j 3^{n-j}, \qquad A(\varnothing) = 0. \tag{1.2}$$

The recursion for appending one digit is

$$A(wb) = 3A(w) + 2b. \tag{1.3}$$

For example, $A(10) = 6$ and $A(01) = 2$, so the order of bits is not cosmetic.

Proposition 1.1 (Exact cylinder algebra)**.** *For every word w of length n and every $x \in \mathbb{R}$,*

$$S_w(x) = \frac{A(w)+x}{3^n}, \qquad I_w := S_w([0,1]) = \left[\frac{A(w)}{3^n}, \frac{A(w)+1}{3^n}\right], \tag{1.4}$$

$$A(uv) = 3^{|v|}A(u) + A(v), \qquad 0 \le A(w) \le 3^n - 1. \tag{1.5}$$

All $A(w)$ are even integers. The map $w \mapsto (n, A(w))$ is injective. The numerator $A(w)$ without its depth tag is not injective.

Proof. For $n = 0$, the affine map is the identity and the interval is $[0,1]$. Appending b gives

$$S_{wb}(x) = S_w(S_b(x)) = \frac{3A(w) + 2b + x}{3^{n+1}},$$

which proves the affine formula by induction. Splitting the sum in (1.2) at $|u|$ proves the concatenation identity. The range bound follows either by summing a geometric progression or by induction: if $A + 1 \le d$, then

$$(3A + 2b) + 1 \le 3d \quad \text{for } b \in \{0, 1\}.$$

Parity follows from the recurrence.

For injectivity at a fixed depth, suppose $3A(u) + 2b = 3A(v) + 2c$. Reduction modulo 3 distinguishes b and c, since 0 and 2 are distinct residues. Therefore $b = c$, and cancellation gives $A(u) = A(v)$. Induction removes all digits. The depth is essential: $A(\varnothing) = A(0) = A(00) = 0$. □

The same affine maps have homogeneous-coordinate representatives

$$M_b = \begin{pmatrix} 1 & 2b \\ 0 & 3 \end{pmatrix}, \qquad M_{b_1} \cdots M_{b_n} = \begin{pmatrix} 1 & A(w) \\ 0 & 3^n \end{pmatrix}.$$

This realizes the combinatorics as an exact matrix semigroup. These maps are invertible as affine maps of $\mathbb{R}$, but their restrictions into $[0, 1]$ are not surjective. The construction uses a semigroup action on the interval, not a group action by interval automorphisms.

1.1.2 Separation and finite-state invariants

Writing $a = A(w)$, the children of I_w have the common denominator 3^{n+1}:

$$I_{w0} = \frac{[3a, 3a + 1]}{3^{n+1}}, \qquad I_{w1} = \frac{[3a + 2, 3a + 3]}{3^{n+1}}. \tag{1.6}$$

The open gap is $((3a + 1)/3^{n+1}, (3a + 2)/3^{n+1})$. Consequently both children have length one third of the parent's length, their separation is one third of that length, and their outer endpoints equal the parent's outer endpoints. No floating-point division is required to check any of these facts: one checks the integer numerators in their common denominator.

Corollary 1.2 (Stage cardinality and retained length)**.** *Let $K_n = \bigcup_{w \in \{0,1\}^n} I_w$. Then $K_0 = [0, 1]$, $K_{n+1} \subset K_n$, the level-n intervals are pairwise disjoint, and*

$$\#\{I_w : |w| = n\} = 2^n, \qquad |I_w| = 3^{-n}, \qquad \lambda(K_n) = (2/3)^n. \tag{1.7}$$

Proof. Two distinct words first separate into opposite children of a common ancestor. Those children are disjoint by (1.6); all subsequent intervals remain within them. Thus there are precisely 2^n disjoint intervals, each of length 3^{-n}. Add their lengths. □

1.1.3 Infinite words and quantitative uniqueness

Put $\Sigma = \{0, 1\}^{\mathbb{N}_{\ge 1}}$ and define

$$\pi(b) = \sum_{j=1}^{\infty} \frac{2b_j}{3^j}. \tag{1.8}$$

The sum converges absolutely. Its prefix approximation $p_n(b)$ satisfies

$$p_n(b) = \frac{A(b_1 \cdots b_n)}{3^n}, \qquad 0 \le \pi(b) - p_n(b) \le 3^{-n}. \tag{1.9}$$

Indeed, the largest possible tail is $\sum_{j=n+1}^{\infty} 2/3^j = 3^{-n}$. Both endpoints of this interval are realizable, by an all-zero or an all-one tail.

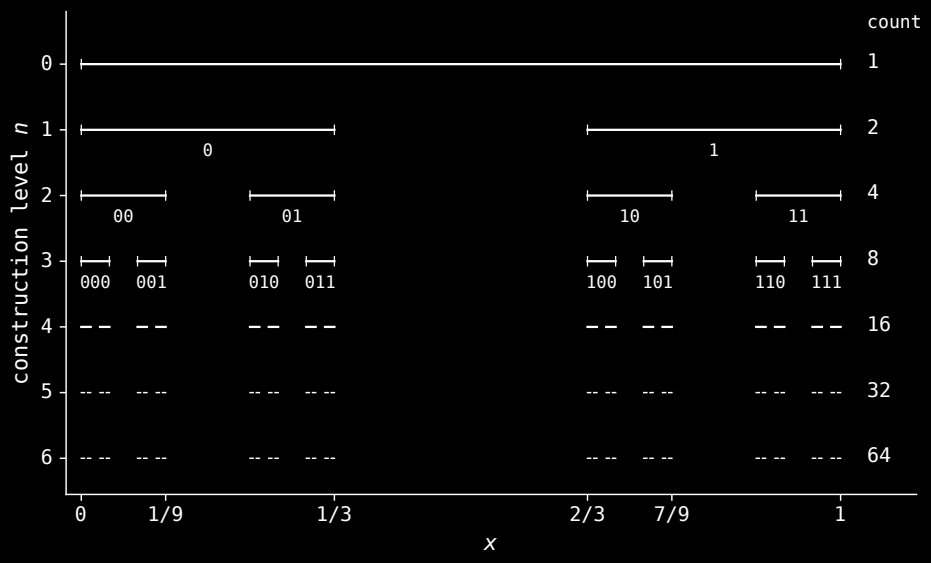


Figure 1.1: An exact-coordinate reconstruction of the hierarchy in source Figure 1.1, extended through level six. Addresses, corrected endpoints, interval counts, nesting, and separation are visible simultaneously. The plotted finite stages are not themselves perfect totally disconnected sets.

Lemma 1.3 (First-difference estimate)**.** *If $a, b \in \Sigma$ first differ at position k, then*

$$3^{-k} \le |\pi(a) - \pi(b)| \le 3^{1-k}. \tag{1.10}$$

In particular, the restricted ternary representation with digits 0 *and* 2 *is unique.*

Proof. Interchange a, b if necessary so that $a_k = 0$, $b_k = 1$. The common prefix cancels and

$$\pi(b) - \pi(a) = 2\,3^{-k} + v - u, \qquad 0 \le u, v \le 3^{-k}.$$

The minimum occurs at $u = 3^{-k}, v = 0$, and the maximum at $u = 0, v = 3^{-k}$. This proves both bounds and positivity. □

There is no contradiction with the familiar nonuniqueness of unrestricted base-three expansions. For example,

$$\frac{1}{3} = 0.1_3 = 0.02222\ldots_3.$$

Only the second representation uses the permitted digits exclusively. A program that rejects 1/3 on encountering the digit 1 in its terminating expansion has confused a representation choice with a membership proof.

Define $d_3(a, b) = 0$ for $a = b$ and $d_3(a, b) = 3^{-k}$ at the first difference. If a, b agree through position r, and b, c do also, then a, c agree through position r. Hence

$$d_3(a, c) \le \max\{d_3(a, b), d_3(b, c)\}.$$

This is the ultrametric inequality. Equation (1.10) gives $d_3 \le |\pi(a) - \pi(b)| \le 3d_3$: the coding is bi-Lipschitz for this particular symbolic metric.

For the more usual binary metric $d_2(a, b) = 2^{-k}$, set $\alpha = \log 3/\log 2$. Then

$$d_2(a, b)^\alpha \le |\pi(a) - \pi(b)| \le 3d_2(a, b)^\alpha. \tag{1.11}$$

Thus the numerical regularity of a coding depends on the metric assigned to its symbolic domain. A topological conjugacy does not, by itself, specify a Euclidean condition number.

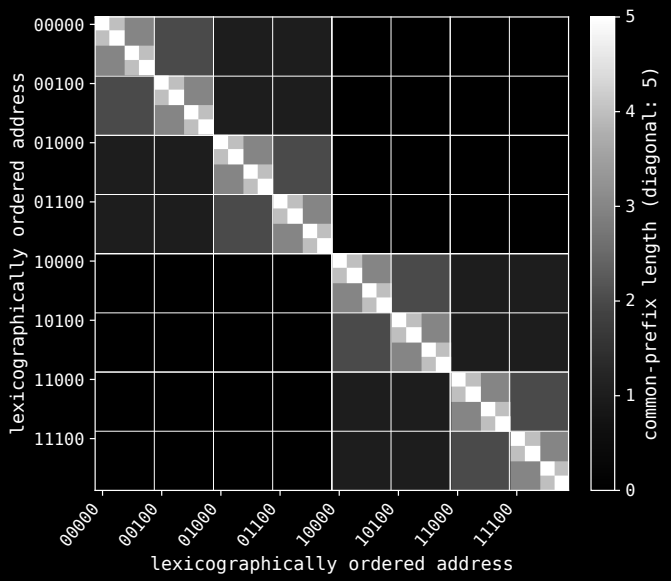


Figure 1.2: The common-prefix matrix for the 32 depth-five addresses. Its nested blocks are a radix-tree partition and a finite model of the ultrametric. For off-diagonal entries r, the corresponding infinite-address distance is $3^{-(r+1)}$ when that first difference is fixed. A diagonal entry records only agreement of the five displayed bits, not equality of infinite tails.

Computational interpretation. A cylinder is stored as the exact pair (n, A) instead of as two binary floating-point endpoints. Appending a bit applies the two integer updates $n \leftarrow n + 1$ and $A \leftarrow 3A + 2b$. The invariant $0 \le A < 3^n$ is suitable for a loop proof. Random access to one cylinder requires n such updates and $O(n)$ output bits; storing the entire stage requires 2^n cylinders. The difference between one-path cost and whole-tree cost must not disappear behind a casual claim of "linear complexity."

1.2 Analytical and geometric rederivation

1.2.1 Existence of the limit and solution of source Exercise 1.1

Define $K_{1/3} = \bigcap_{n\ge 0} K_n$. Each K_n is a nonempty compact subset of $[0, 1]$, and the sequence is decreasing. Alternatively to invoking compact intersection directly, fix any $b \in \Sigma$: the intervals $I_{b_1 \cdots b_n}$ have nested endpoints and diameters 3^{-n}, so their intersection consists of exactly the real number $\pi(b)$. In particular, $K_{1/3}$ is nonempty, closed, and bounded.

Conversely, each $x \in K_{1/3}$ lies in one and only one interval at every level. Nesting makes their words compatible, so they are the prefixes of one infinite word b. The diameter estimate forces $x = \pi(b)$. Therefore

$$\pi(\Sigma) = K_{1/3}. \tag{1.12}$$

This identifies the series construction, the nested-set construction, and the tree-boundary construction without assuming their equivalence.

Theorem 1.4 (Classical Cantor-set properties)**.** *The set $K_{1/3}$ is nonempty, compact, perfect, nowhere dense, uncountable, and totally disconnected. Its Lebesgue measure is zero.*

Proof. Compactness and nonemptiness were established above. Given $x = \pi(b)$, change one digit of b arbitrarily far out, leaving every other digit unchanged. The resulting point is distinct by Lemma 1.3 and tends to x by the tail bound. Thus every point is a limit point of other points of $K_{1/3}$; together with closedness this is perfectness.

If a connected subset $E \subset K_{1/3}$ contained distinct points x, y, choose n with $3^{-n} < |x - y|$. The two points belong to distinct level-n intervals. Those intervals have a gap between them. Intersecting E with the two open half-lines determined by a point of the gap separates E, a contradiction. Thus every connected subset is a singleton.

No nondegenerate open interval can lie in $K_{1/3}$: it would be a connected subset with more than one point. Since $K_{1/3}$ is closed, this proves nowhere density. Injectivity of π and the diagonal argument for binary sequences prove uncountability. Finally, monotonicity of Lebesgue measure and (1.7) give

$$0 \le \lambda(K_{1/3}) \le (2/3)^n \quad \text{for every } n,$$

and the right-hand side tends to zero. □

The zero-measure conclusion is specific to the classical ratios. It cannot be appended to the intrinsic topological definition of a Cantor set. Section 1.2.4 constructs a positive-measure example with exactly the same binary branching pattern.

1.2.2 Hausdorff error: what the finite picture certifies

For nonempty compact sets $E, F \subset \mathbb{R}$, write

$$d_H(E, F) = \max\left\{ \sup_{x \in E} \operatorname{dist}(x, F), \sup_{y \in F} \operatorname{dist}(y, E) \right\}.$$

The empty set is deliberately excluded: the displayed formula is not an ordinary finite-valued metric on a hyperspace containing it.

Proposition 1.5 (Three different approximation errors)**.** *Let E_n be the set of left endpoints of level-n cylinders, and F_n the set of both endpoints. Then*

$$d_H(K_n, K_{1/3}) = \frac{3^{-n}}{6}, \qquad d_H(E_n, K_{1/3}) = 3^{-n}, \qquad d_H(F_n, K_{1/3}) = 3^{-(n+1)}. \tag{1.13}$$

Proof. Inside a level-n cylinder of length $h = 3^{-n}$, the largest subsequently removed gap has length $h/3$. Its endpoints belong to $K_{1/3}$, and its midpoint is at distance $h/6$ from $K_{1/3}$. Every other removed gap is smaller. Since $K_{1/3} \subset K_n$, the first equality follows.

Every point of $K_{1/3}$ in a cylinder is at distance at most h from its left endpoint. The point 1 attains distance h from E_n, giving the second equality. For both endpoints, rescale a cylinder to $[0, 1]$. Points of $K_{1/3}$ lie in $[0, 1/3] \cup [2/3, 1]$, so their distance to $\{0, 1\}$ is at most $1/3$. It equals $1/3$ at $1/3$ and $2/3$. Endpoints of other level-n cylinders cannot be nearer than the endpoints of the cylinder containing the point. Rescale back to obtain the third equality. □

For a required enclosure error $\varepsilon > 0$, it suffices to select $n \ge \max\{0, \lceil \log_3(1/(6\varepsilon)) \rceil\}$ when using K_n. In exact code, this depth is more safely selected by multiplying an integer power of three until the rational inequality holds, instead of by rounding a floating logarithm near an integer threshold.

An operator proof of existence. On nonempty compact subsets of $[0, 1]$, define $\mathcal{H}(E) = S_0(E) \cup S_1(E)$. Matching each branch with itself shows

$$d_H(\mathcal{H}(E), \mathcal{H}(F)) \le \tfrac{1}{3} d_H(E, F).$$

To justify this estimate, for every $x \in E$ choose points of F as close as desired to the distance infimum; their images under the same S_b are one third as far away. Repeat in the opposite direction.

Completeness of the compact hyperspace then permits the contraction theorem. It gives a unique nonempty compact fixed set. Since the set just constructed satisfies $K_{1/3} = S_0(K_{1/3}) \cup S_1(K_{1/3})$, the fixed set is $K_{1/3}$. This is the present case of the invariant-set framework developed by Hutchinson [1, §§3–4]; the direct nested-interval proof above does not require importing that general theorem.

1.2.3 Hausdorff dimension, with both inequalities proved

Let

$$s = \frac{\log 2}{\log 3}. \tag{1.14}$$

Define a probability measure μ on $K_{1/3}$ by assigning mass 2^{-n} to each depth-n cylinder slice $K_{1/3} \cap I_w$. One construction is the pushforward of the fair product measure on Σ under π. Consistency of the masses follows because every cylinder has two children of equal mass. No singleton has positive mass: its mass is at most 2^{-n} for all n.

Theorem 1.6. *The Hausdorff dimension of $K_{1/3}$ is s.*

Proof. For $t > s$, the level-n cover has total t-cost

$$\sum_{|w|=n} (\operatorname{diam} I_w)^t = 2^n 3^{-nt} = (2\,3^{-t})^n \longrightarrow 0.$$

Thus $\mathcal{H}^t(K_{1/3}) = 0$ and $\dim_H K_{1/3} \le s$.

For the reverse inequality, take an interval J of length $0 < \ell \le 1$ and choose $n \ge 0$ so that $3^{-(n+1)} < \ell \le 3^{-n}$. Distinct level-$n$ cylinders are separated by at least 3^{-n}. Therefore J meets at most two of them, with endpoint coincidences harmless. It follows that

$$\mu(J) \le 2 \cdot 2^{-n} = 2(3^{-n})^s \le 2 \cdot 3^s \ell^s = 4\ell^s.$$

An arbitrary nonempty bounded subset of $\mathbb{R}$ is contained in a closed interval of the same diameter. Given a countable cover $\{U_i\}$ of $K_{1/3}$ by sufficiently small sets, replace its nonempty members by closed intervals J_i with the same diameters. Countable subadditivity on these measurable intervals gives

$$1 \le \sum_i \mu(J_i) \le 4 \sum_i (\operatorname{diam} U_i)^s.$$

Singleton cover elements have zero measure and cause no exception. Hence $\mathcal{H}^s(K_{1/3}) \ge 1/4$ under the diameter-cost convention, and $\dim_H K_{1/3} \ge s$. The constant $1/4$ is only a sufficient lower bound; no sharp value of $\mathcal{H}^s(K_{1/3})$ is asserted here. □

The equation $2(1/3)^s = 1$ suggests the exponent but does not prove the lower bound. The separation and mass estimate are the load-bearing steps. Without a separation hypothesis, a similarity-dimension calculation can overcount overlapping pieces.

Computational interpretation. At resolution $\varepsilon \asymp 3^{-n}$, a complete cylinder representation uses $2^n \asymp \varepsilon^{-s}$ leaves. This is an output-size estimate, not an automatic running-time guarantee for every operation on the set. It explains why geometry-aware hierarchical methods can require fewer cells than a uniform interval grid of size ε^{-1}, while still having exponential cost in the address depth.

1.2.4 General interval trees and the geometry/combinatorics split

Here is an explicit version of the setup behind source Exercise 1.4. Let T be a rooted, locally finite, pruned tree. "Pruned" means that every vertex extends to an infinite branch. Attach a nonempty compact interval I_v to every vertex. Children lie inside the parent, distinct same-level intervals are disjoint, and along every infinite branch the interval diameters tend to zero. No assumption of equal lengths is made.

The branch space $[T]$ is compact: successive subsequence extraction fixes one vertex at each level, and the diagonal sequence converges in the prefix topology. Each branch determines a unique point of $K = \bigcap_n \bigcup_{|v|=n} I_v$. Distinct branches determine distinct points, because at their first split they lie in disjoint compact intervals. Conversely, the unique interval containing any $x \in K$ at each level gives a compatible branch.

Lemma 1.7 (Branchwise shrinking implies uniform mesh shrinking)**.** *Under the preceding hypotheses,* $\max_{|v|=n} \operatorname{diam} I_v \longrightarrow 0$.

Proof. Fix $\varepsilon > 0$. Every branch has a prefix v with $\operatorname{diam} I_v < \varepsilon$. Its cylinder is an open neighborhood in $[T]$. These neighborhoods cover $[T]$, and compactness selects finitely many. Let N be the largest selected depth. Every level-N vertex lies on a branch, and that branch has one of the selected prefixes. Its interval is therefore contained in an interval of diameter less than ε. The same is true at later levels. □

Proposition 1.8 (The missing perfectness criterion)**.** *The limit K is a nonempty compact totally disconnected set. It is perfect if and only if no branch of $[T]$ is isolated. Equivalently, along every branch and after every finite prefix, a later vertex has at least two children with infinite continuations. Requiring every vertex to have at least two children is a sufficient, but stronger, condition.*

Proof. The coding just constructed is continuous: a sufficiently long common prefix traps both image points in an arbitrarily small interval. It is a continuous bijection from compact $[T]$ to Hausdorff K, hence a homeomorphism. A connected subset of $[T]$ cannot meet two disjoint prefix cylinders; refining the depth proves total disconnectedness. An isolated branch is exactly a branch whose sufficiently long prefix has a unique infinite continuation. A singleton open cylinder maps to an isolated point of K, and the converse follows from the homeomorphism. □

Validity analysis. The source's stated shrinking condition alone does not imply perfectness. A tree with one vertex at each level and intervals $I_n = [0, 3^{-n}]$ satisfies nesting and shrinking, yet its limit is $\{0\}$. All levels of this example are nonempty; the failure is not explained away by a missing nonemptiness assumption. The essential absent ingredient is persistent branching.

A precise comparison convention. For the quantitative statements below, "C-bounded geometry" means that for every nondegenerate parent I of length L, each child and every open gap between consecutive children has length at least L/C. The reverse comparison is automatic because these pieces lie in I. This fixes a specific realization of the source's comparability language; it is not an assertion that a hierarchy is canonical. The source explicitly allows adjustments that preserve the slices of the Cantor set.

If a parent has $m \geq 2$ children, it has $m - 1$ intervening gaps. Adding their lower length bounds yields

$$\frac{(2m-1)L}{C} \leq L, \qquad m \leq \left\lfloor \frac{C+1}{2} \right\rfloor. \tag{1.15}$$

Thus bounded geometry implies bounded combinatorics. The classical construction has $m = 2$ and $C = 3$, with equality. Furthermore, at least one gap removes a fraction $1/C$ from each parent, so

$$\lambda(K_n) \le (1 - 1/C)^n \lambda(I_\varnothing)$$

whenever the uniform comparison and the two-child condition hold at every level. Such a limit has zero Lebesgue measure.

Example 1.9 (Binary combinatorics without bounded geometry)**.** At level $n \ge 1$, retain two end children, each of relative length $r_n = 1/(n+3)$. The tree is always binary, every branch shrinks, and the limit is perfect. But the ratio of a parent to either child is $n + 3$, which is unbounded. Hence no fixed C works.

Example 1.10 (A positive-measure Cantor set)**.** At level $n \ge 1$, remove the open middle fraction $\delta_n = 2^{-(n+1)}$ from every interval. Each remaining child has relative length $(1 - \delta_n)/2 < 1/2$, so all branch diameters tend to zero and persistent binary splitting gives perfectness. The retained length is

$$\lambda(K_n) = \prod_{j=1}^{n} (1 - \delta_j) \ge 1 - \sum_{j=1}^{n} \delta_j \ge \tfrac{1}{2}.$$

The product inequality follows inductively from $(1 - a)(1 - b) \ge 1 - a - b$ for $a, b \in [0, 1]$. Continuity of measure from above gives $\lambda(K) \ge 1/2$. Its gap ratios δ_n tend to zero, so bounded geometry fails even though the child ratios remain near $1/2$.

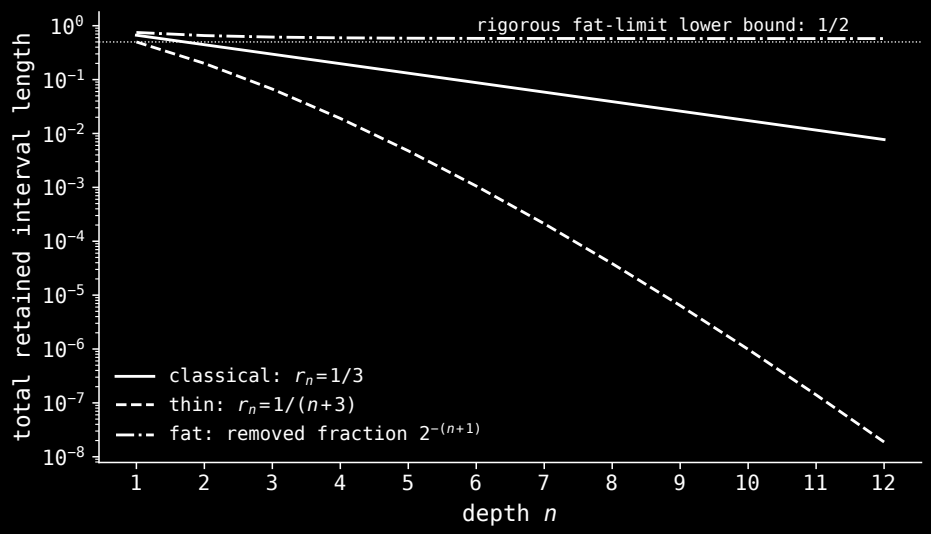


Figure 1.3: Identical full binary combinatorics, radically different retained lengths. The classical and thin constructions have vanishing length, whereas the fat construction has limit length at least $1/2$. The plotted values are finite products; the limit claims follow from the proofs, not from extrapolating a plot.

Effectivity qualification. The abstract compactness proof supplies no input format or complexity bound. If all finite levels and their rational endpoints can actually be enumerated, a search for the first level with mesh below a prescribed rational tolerance terminates by Lemma 1.7. Thus an effective finite-level presentation does yield a computable mesh modulus by exhaustive search. What is missing without further quantitative hypotheses is an efficient a priori rate, not necessarily computability.

1.3 Functional analysis: approximation spaces, measures, and operators

Additional development. This subsection is not part of the source's discussion on pp. 13–15. It turns the same hierarchy into an approximation and operator framework. The point is not to replace the elementary construction with a more difficult proof, but to extract numerical schemes with explicit error bounds and clear assumptions.

1.3.1 Cylinder functions as finite-dimensional approximation spaces

Write $C_w = K_{1/3} \cap I_w$. A cylinder slice is closed in $K_{1/3}$; its complement is the finite union of the other level-n slices and is also closed. Consequently C_w is both open and closed, and its indicator $\mathbf{1}_{C_w}$ is continuous on $K_{1/3}$.

Let $C(K_{1/3})$ denote the complex Banach space of continuous functions on $K_{1/3}$ with the supremum norm. Define

$$V_n = \operatorname{span}\{\mathbf{1}_{C_w} : |w| = n\}.$$

Then $V_n \cong \mathbb{C}^{2^n}$ and $V_n \subset V_{n+1}$. Under these coordinate identifications, the refinement map duplicates each coordinate:

$$(a_w)_{|w|=n} \longmapsto (a_w, a_w)_{|w|=n}.$$

This is an exact algebra homomorphism for pointwise multiplication and an isometry for the supremum norm.

Proposition 1.11 (Constructive density)**.** *The union $\bigcup_n V_n$ is uniformly dense in $C(K_{1/3})$. More precisely, if $f \in C(K_{1/3})$ has modulus of continuity*

$$\omega_f(\delta) = \sup\{|f(x) - f(y)| : x, y \in K_{1/3},\ |x - y| \le \delta\},$$

then choosing any $x_w \in C_w$ and setting $R_n f = \sum_{|w|=n} f(x_w)\mathbf{1}_{C_w}$ gives

$$\|f - R_n f\|_\infty \le \omega_f(3^{-n}) \longrightarrow 0. \tag{1.16}$$

Proof. If $x \in C_w$, then both x and x_w lie in an interval of diameter 3^{-n}. Apply the definition of ω_f. Uniform continuity follows from compactness of $K_{1/3}$. □

For an α-Hölder function, $0 < \alpha \le 1$, with seminorm

$$[f]_\alpha = \sup_{x \ne y} \frac{|f(x) - f(y)|}{|x - y|^\alpha},$$

the error is at most $[f]_\alpha 3^{-\alpha n}$. This is a usable approximation theorem: it specifies the subspace, samples, norm, and rate. A statement that cylinder functions are merely "dense" would not provide a stopping condition.

1.3.2 A probability measure distinct from ambient length

The cylinder measure μ introduced in Section 1.2.3 satisfies

$$\mu(C_w) = 2^{-|w|}, \qquad \int f \, d\mu = \frac{1}{2}\int f \circ S_0 \, d\mu + \frac{1}{2}\int f \circ S_1 \, d\mu. \tag{1.17}$$

The second equality follows by conditioning on the first bit of a fair binary sequence. It is valid first for cylinder indicators, then for bounded Borel functions by standard approximation. In

particular, $\mu(K_{1/3}) = 1$ while $\lambda(K_{1/3}) = 0$: μ is singular with respect to Lebesgue measure. Any numerical integral must state which of these measures it uses.

Let $X = \pi(B_1, B_2, \ldots)$ with independent fair bits. The tail after its first bit has the same distribution as X, independently of B_1, so $X \overset{d}{=} (2B_1 + X')/3$. If $m_j = \int x^j \, d\mu$, then $m_0 = 1$ and

$$(2 \cdot 3^j - 2)m_j = \sum_{r=0}^{j-1} \binom{j}{r} 2^{j-r} m_r \qquad (j \geq 1). \tag{1.18}$$

This follows by expanding $2 \cdot 3^j m_j = m_j + \int (x+2)^j \, d\mu$ and moving the two copies of m_j to the left. Thus

$$m_1 = \frac{1}{2}, \qquad m_2 = \frac{3}{8}, \qquad \mathrm{Var}_\mu(X) = \frac{1}{8}. \tag{1.19}$$

The recurrence reduces every polynomial integral to exact rational arithmetic. It is not an empirical fit to a histogram.

1.3.3 Conditional expectation and a complete Haar basis

In the Hilbert space $L^2(K_{1/3}, \mu)$, let E_n denote conditional expectation on the depth-n cylinder partition. Its explicit formula is

$$(\mathsf{E}_n f)|_{C_w} = 2^n \int_{C_w} f \, d\mu. \tag{1.20}$$

It is the orthogonal projection onto V_n. Indeed, subtracting the displayed average makes the integral of $f - \mathsf{E}_n f$ vanish on each cylinder, hence makes it orthogonal to every element of V_n. Therefore $\mathsf{E}_n^2 = \mathsf{E}_n$, $\|\mathsf{E}_n\|_{2\to 2} = 1$, and $\mathsf{E}_n \mathsf{E}_m = E_{\min\{m,n\}}$. For continuous f, averaging the oscillation within a cylinder gives $\|f - \mathsf{E}_n f\|_\infty \leq \omega_f(3^{-n})$.

For each finite word w, define

$$h_w = 2^{|w|/2}\big(\mathbf{1}_{C_{w0}} - \mathbf{1}_{C_{w1}}\big). \tag{1.21}$$

Then $\int h_w \, d\mu = 0$ and $\|h_w\|_2^2 = 2^{|w|}(2^{-|w|-1} + 2^{-|w|-1}) = 1$. If two supports are disjoint, the functions are orthogonal. Otherwise the shorter-address function is constant on the support of the longer-address function, whose mean is zero. It follows that

$$\{1\} \cup \{h_w : w \in \{0,1\}^*\}$$

is an orthonormal family. The functions with $|w| < n$, together with 1, number 2^n and belong to V_n, so they form a basis of V_n.

For completeness in L^2, note that finite unions of cylinders form an algebra generating the Borel sets. Sets approximable in measure by this algebra form a sigma-algebra: complements preserve approximation, and a countable union is first truncated to a finite union using the finiteness of μ, then approximated. Thus all Borel indicators can be approximated in L^2 by cylinder indicators. Simple functions and truncation then prove L^2-density. The orthonormal family is therefore complete, not merely a finite-stage pattern.

Worked multiresolution calculation. For the observable $f(x) = x$ and a word w of length k, the means on its children differ by $-(2/3)3^{-k}$. Multiplying this difference by their mass 2^{-k-1} and the normalization $2^{k/2}$ gives

$$\langle x, h_w \rangle = -2^{-k/2} 3^{-(k+1)}.$$

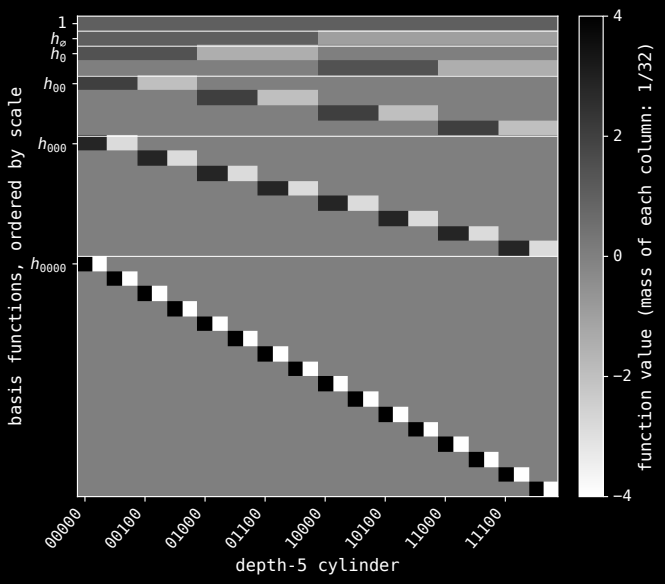


Figure 1.4: The complete depth-five Haar transform, evaluated on its 32 cylinder atoms. Rows are 1 and the h_w in increasing scale order. Every column has probability mass $1/32$, so the exact orthogonality identity is $HH^{\mathsf{T}}/32 = I$, not $HH^{\mathsf{T}} = I$. The shades encode signed function values; amplitudes are specified by (1.21). The image is a rendering, while the orthogonality proof is independent of its pixels.

There are 2^k words at that scale, so the total squared coefficient energy is $3^{-2(k+1)}$. Summing gives

$$\|x - \tfrac{1}{2}\|_2^2 = \sum_{k=0}^{\infty} 3^{-2(k+1)} = \frac{1}{8}, \qquad \|x - \mathsf{E}_n x\|_2^2 = \frac{3^{-2n}}{8}.$$

This reproduces the independently derived variance and provides an exact compression-error certificate. It also gives a sensitive regression test for an incorrectly normalized Haar implementation.

1.3.4 Transfer and Koopman operators: where contraction is valid

The map on the Cantor set corresponding to deletion of the first bit is

$$T(x) = \begin{cases} 3x, & x \in K_{1/3} \cap [0, 1/3], \\ 3x - 2, & x \in K_{1/3} \cap [2/3, 1]. \end{cases} \tag{1.22}$$

The two branches are separated in its domain, so $T : K_{1/3} \to K_{1/3}$ is continuous and $T\pi = \pi\sigma$, where σ is the left shift. The central open gap is not part of this domain.

Define

$$Uf = f \circ T, \qquad Pf = \tfrac{1}{2}(f \circ S_0 + f \circ S_1).$$

Both preserve constants and positivity. Shift invariance of the fair product measure gives $\int Uf\, d\mu = \int f\, d\mu$ and $\|Uf\|_2 = \|f\|_2$. With $\langle f, g\rangle = \int f\overline{g}\, d\mu$, splitting into the first-bit cylinders yields

$$\langle Uf, g\rangle = \langle f, Pg\rangle, \qquad P = U^*, \qquad PU = I. \tag{1.23}$$

Here U is an isometry and P is a coisometry. They are not mutual inverses: for example $Ph_\varnothing = 0$, although $h_\varnothing \neq 0$. The operator UP is the projection onto functions of the tail, not the finite-prefix projection in (1.20).

On the supremum norm, $\|Pf\|_\infty \leq \|f\|_\infty$, and equality holds for $f = 1$. Thus P is not a strict contraction on all of $C(K_{1/3})$. On the Hölder seminorm, however,

$$[Pf]_\alpha \leq 3^{-\alpha}[f]_\alpha. \tag{1.24}$$

This follows by comparing each pair $f(S_b x), f(S_b y)$ and using $|S_b x - S_b y| = |x-y|/3$.

Proposition 1.12 (Certified mixing of observables)**.** *For every α-Hölder f and every $n \geq 0$,*

$$\left\| P^n f - \int f \, d\mu \right\|_\infty \leq [f]_\alpha 3^{-\alpha n}. \tag{1.25}$$

For a merely continuous f, the right-hand side can be replaced by $\omega_f(3^{-n})$.

Proof. Expand the iterate and iterate (1.17):

$$P^n f(x) = 2^{-n} \sum_{|w|=n} f(S_w x), \qquad \int f \, d\mu = 2^{-n} \sum_{|w|=n} \int f(S_w y) \, d\mu(y).$$

Subtract, apply the triangle inequality, and use $|S_w x - S_w y| \leq 3^{-n}$. The weights sum to one. □

Computational interpretation. The operator provides a deterministic dynamic-programming scheme for averaging a continuous observable over a symbolic execution space. Its correct rate is a bound in a specified regularity class. A software claim of "geometric convergence" without identifying that class is too weak: continuous functions need not share one uniform Hölder constant or exponent.

1.3.5 Recovering the space from its function algebra

The preceding algebraic and functional perspectives fit together more strongly than a loose analogy. With pointwise multiplication and complex conjugation, $C(K_{1/3})$ is the norm completion of the increasing finite algebras V_n. Its idempotents are precisely the indicators of clopen sets: $f^2 = f$ forces every value to be 0 or 1, and continuity makes the two level sets clopen.

A nonzero complex-linear multiplicative map $\chi : C(K_{1/3}) \to \mathbb{C}$ is evaluation at a unique point of $K_{1/3}$. To see this, $\chi(1) = 1$ and the mutually orthogonal depth-n cylinder idempotents sum to 1, so exactly one has image 1. Compatibility under refinement selects an infinite branch b. On every V_n, χ equals evaluation at $\pi(b)$.

For the extension to the completion, boundedness need not be assumed silently. If $\chi(f) \notin f(K_{1/3})$, then $f - \chi(f)1$ has a continuous pointwise inverse. Multiplicativity would give $0 = 1$, a contradiction. Therefore $|\chi(f)| \leq \|f\|_\infty$. Density now shows $\chi(f) = f(\pi(b))$ for every continuous f. Uniqueness follows from cylinder indicators separating points.

This elementary spectrum argument is useful in computer science because it identifies an infinite state by its consistent finite Boolean observations. A character choosing incompatible cylinder idempotents is not a valid state; the refinement equations are its consistency contract.

1.4 Topological rederivation and source Exercise 1.2

1.4.1 Cantor space as an inverse limit

Let $D_n = \{0,1\}^n$ with the discrete topology and let $p_n : D_{n+1} \to D_n$ delete the last digit. The inverse limit is

$$\varprojlim (D_n, p_n) = \{(w_n)_{n \geq 0} : p_n(w_{n+1}) = w_n\}.$$

Every compatible family consists of all the prefixes of a unique infinite word. Thus $\Sigma \cong \varprojlim D_n$. Compactness follows either from the product theorem for the finite discrete factors or directly from the diagonal argument. Its prefix cylinders form a countable clopen basis.

The coding π is continuous by (1.9), injective by Lemma 1.3, and onto $K_{1/3}$ by (1.12). A continuous bijection from compact to Hausdorff is a homeomorphism: the image of a closed set is compact, hence closed, so the inverse is continuous. We have therefore proved

$$K_{1/3} \cong \Sigma \cong \varprojlim D_n. \tag{1.26}$$

The spaces D_n are discrete, not unions of intervals. An inverse-limit argument about D_n must not be replaced by an unjustified assertion that each K_n is already homeomorphic to a finite discrete set.

1.4.2 Why the classification statement needs its hypotheses

The source describes an intrinsic Cantor set as a perfect totally disconnected set and then asks that all Cantor sets be shown homeomorphic. For a correct abstract classification, this synopsis uses:

$$X \neq \varnothing, \quad X \text{ compact metrizable}, \quad X \text{ has no isolated points and is totally disconnected.} \tag{1.27}$$

The book's surrounding metrizability convention does not itself supply compactness or nonemptiness for every possible set. Without compactness, $\bigcup_{m\in\mathbb{Z}}(2m + K_{1/3})$ is a perfect totally disconnected closed subset of $\mathbb{R}$ but cannot be homeomorphic to compact $K_{1/3}$. The empty set also requires exclusion if "perfect" is allowed to hold vacuously.

Lemma 1.13 (Clopen bases from compact total disconnectedness)**.** *A compact metrizable totally disconnected space has a basis of clopen sets.*

Proof. Fix $x \in X$, and let Q_x be the intersection of all clopen neighborhoods of x. We first prove Q_x is connected. Otherwise split it into two nonempty compact pieces and choose disjoint ambient open neighborhoods U, V covering those pieces, with $x \in U$. The compact set $X \setminus (U \cup V)$ is disjoint from Q_x. For each of its points, some clopen neighborhood of x excludes that point. Finitely many suffice; their intersection W is clopen, contains Q_x, and lies in $U \cup V$. Now $W \cap U$ is clopen in X: within W, its complement is $W \cap V$, and W itself is clopen. It contains x but excludes the nonempty piece $Q_x \cap V$, contradicting the definition of Q_x.

Every connected set containing x lies in all its clopen neighborhoods, so the connected component of x lies in Q_x. Conversely, the connected set Q_x lies in that component. Total disconnectedness therefore gives $Q_x = \{x\}$. For any open neighborhood O of x, use compactness of $X \setminus O$ again to intersect finitely many clopen neighborhoods excluding it. The resulting clopen neighborhood lies in O. □

Theorem 1.14 (Corrected form of source Exercise 1.2)**.** *Any two spaces satisfying* (1.27) *are homeomorphic. In particular, each is homeomorphic to* $K_{1/3}$*.*

Proof. Let the spaces be X, Y. We construct finite clopen partitions $\mathcal{P}_n$ of X and $\mathcal{Q}_n$ of Y, with a bijection between their cells compatible under refinement. Start with the single matched pair (X, Y).

Suppose U, V are matched nonempty clopen cells. By Lemma 1.13 and compactness, partition U into finitely many nonempty clopen cells of diameter less than 2^{-n}; a finite clopen cover can be made disjoint by taking successive differences. Do the same for V. Their numbers of cells need not agree. Every nonempty clopen cell has at least two points, because the original space has no isolated points. It can be split into two nonempty clopen cells using the clopen basis. Repeatedly split cells in the smaller partition until both partitions have the same number of cells. Splitting preserves the diameter bound. Choose a bijection between these refined cells and repeat for every matched pair.

A point $x \in X$ determines nested cells U_n, and the matched cells $V_n \subset Y$ are nested, nonempty, compact, and of diameters tending to zero. Their intersection is one point $h(x)$. Reversing the roles constructs its inverse, so h is a bijection. If x' lies in the same sufficiently fine cell as x, the images lie in a cell of arbitrarily small diameter. Thus h is continuous, and the same argument proves continuity of its inverse. □

This proof preserves a matched refinement tree. It does not produce a canonical homeomorphism, a computable one from an unspecified presentation, or a bi-Lipschitz one. Those are genuinely additional requirements.

1.4.3 Ambient equivalence: complete solution of source Exercise 1.5

Suppose $K, \widetilde{K} \subset \mathbb{R}$ arise from the same *ordered* pruned interval tree, with the shrinking and perfectness conditions of Proposition 1.8. The source explicitly orders siblings; this datum is needed for the following ambient construction.

Let $\pi_T, \widetilde{\pi}_T$ be the two branch codings. Define $h_0 = \widetilde{\pi}_T \circ \pi_T^{-1} : K \to \widetilde{K}$. Both codings are homeomorphisms. At the first level where two addresses differ, their order is the order of their sibling intervals, and it is the same in both realizations. Consequently h_0 is increasing.

Write $a = \min K$, $b = \max K$, $\widetilde{a} = h_0(a)$, and $\widetilde{b} = h_0(b)$. Every bounded complementary gap (u, v) of K corresponds to the gap $(h_0(u), h_0(v))$ of $\widetilde{K}$. Indeed, a point of $\widetilde{K}$ strictly between these endpoints would have a preimage in K strictly between u, v. Define $h = h_0$ on K, and on each such gap set

$$h(x) = h_0(u) + \frac{x - u}{v - u}\big(h_0(v) - h_0(u)\big). \tag{1.28}$$

On the two exterior rays use $h(x) = x - a + \widetilde{a}$ for $x < a$ and $h(x) = x - b + \widetilde{b}$ for $x > b$.

The resulting map is strictly increasing and onto $\mathbb{R}$: it maps the Cantor set, each complementary gap, and the exterior rays onto the corresponding pieces. A strictly increasing surjection $\mathbb{R} \to \mathbb{R}$ has no jump discontinuity, since a jump omits a nonempty interval from its range. It is therefore continuous, and its inverse is continuous by the same argument. This gives the required ambient homeomorphism and makes its restriction induce the identity on the common tree.

Validity analysis. An arbitrary unordered-tree automorphism need not extend to an ambient homeomorphism of the line. For instance, interchange the children below the left root child but leave the root children and the right subtree fixed. On the classical set this reverses the order of 0 and 1/3 but leaves 1 to their right. The resulting map is neither increasing nor decreasing. Every homeomorphism of $\mathbb{R}$ is monotone, so no such ambient extension exists. Preserving a tree as an abstract graph is weaker than preserving its planar order.

1.5 Algebraic topology: invariants that distinguish the right things

Additional development. The following calculations use standard singular and Čech constructions; Hatcher provides background for singular homology and homotopy [2]. The calculations are supplied here because a generic claim that a Cantor set is "topologically trivial" would hide several different and incompatible meanings of triviality.

1.5.1 Homotopy and singular homology

Every continuous map from a connected space into $K_{1/3}$ is constant. In particular, every path in $K_{1/3}$ is constant, and for every basepoint $x \in K_{1/3}$,

$$\pi_j(K_{1/3}, x) = 0 \qquad (j \geq 1).$$

Nevertheless $K_{1/3}$ is not contractible. A contraction $H(x,t)$ would give, for each x, a path from x to the contraction point. Such a path must be constant, which cannot hold for all distinct x.

Every singular simplex $\Delta^j \to K_{1/3}$ is constant, so the singular chain complex splits as a direct sum, indexed by points of $K_{1/3}$, of the singular chain complex of a point. For a constant j-simplex the boundary coefficient is $\sum_{i=0}^{j}(-1)^i$, which is 0 for odd j and 1 for even j. Therefore, with integer coefficients,

$$H_0^{\mathrm{sing}}(K_{1/3};\mathbb{Z}) = \bigoplus_{x \in K_{1/3}} \mathbb{Z}, \qquad H_j^{\mathrm{sing}}(K_{1/3};\mathbb{Z}) = 0 \quad (j \geq 1). \tag{1.29}$$

The degree-zero group has uncountable rank. Vanishing higher groups do not mean that there is one connected component.

For singular cohomology, applying $\mathrm{Hom}(-,\mathbb{Z})$ to the same chain complex gives

$$H^0_{\mathrm{sing}}(K_{1/3};\mathbb{Z}) = \prod_{x \in K_{1/3}} \mathbb{Z}, \qquad H^j_{\mathrm{sing}}(K_{1/3};\mathbb{Z}) = 0 \quad (j \geq 1).$$

Degree-zero singular cocycles are arbitrary functions on path components, not necessarily continuous functions on $K_{1/3}$.

1.5.2 Clopen nerves and Čech cohomology

Every open cover of compact $K_{1/3}$ has a finite clopen refinement, and that refinement can be chosen to be a partition by sufficiently fine cylinder slices. Its nerve is a finite discrete set. Such partitions are cofinal among open covers, so their higher cohomology groups vanish and

$$\check{H}^j(K_{1/3};\mathbb{Z}) = 0 \quad (j \geq 1), \qquad \check{H}^0(K_{1/3};\mathbb{Z}) = \varinjlim \left(\mathbb{Z}^{2^n},\ (a_w) \mapsto (a_w, a_w)\right). \tag{1.30}$$

This direct limit is naturally $C(K_{1/3},\mathbb{Z})$, the group of continuous integer-valued functions. Indeed, every such function is locally constant. Compactness reduces its local cylinder neighborhoods to a finite cover, which is refined by one common cylinder level. The function therefore lies in a finite stage of the displayed direct limit.

The distinction from singular degree-zero cohomology is substantial: $C(K_{1/3},\mathbb{Z})$ is a countable union of the countable sets $\mathbb{Z}^{2^n}$, whereas $\prod_{x \in K_{1/3}} \mathbb{Z}$ contains all, including discontinuous, assignments of integers to points. The topology of locally constant observations is captured by the former; arbitrary pointwise labels are allowed by the latter. Neither should be silently substituted for the other.

Computational interpretation. A data structure based on finite partitions naturally realizes the refinement system in (1.30). A connected-component routine applied to a finite sampled point cloud computes a different object, at a specified scale. Neither calculation establishes the topology of an unresolved continuum without a connecting approximation theorem. In this example, the exact gap sizes provide that connection explicitly in Section 1.9.

1.6 Computer-science application I: exact representations and decisions

1.6.1 Three input types, three different computational questions

An infinite binary address, an exact rational number, and a numerical approximation to an arbitrary real are different inputs. They must not share one unqualified "is in the Cantor set" API.

An address b represents the point $\pi(b) \in K_{1/3}$ by construction. Reading n bits returns the enclosure (1.9). This is a real-number representation with an explicit precision modulus. An exact rational $x = p/q$ permits a terminating membership algorithm. A general real supplied by progressively accurate approximations does not admit a total membership decider; the obstruction is proved below.

1.6.2 A terminating exact-rational membership algorithm

For $x \in \mathbb{Q}$, first reject $x \notin [0,1]$. Repeatedly apply the following rules to the current exact rational y:

$$\begin{array}{c|l} y = 0 \text{ or } y = 1 & \text{accept,} \\ 1/3 < y < 2/3 & \text{reject,} \\ y \le 1/3 & y \leftarrow 3y, \\ y \ge 2/3 & y \leftarrow 3y - 2. \end{array} \tag{1.31}$$

If an already visited value recurs before rejection, accept. The closed endpoint comparisons are intentional.

Proposition 1.15 (Correctness and termination on rational inputs)**.** *The algorithm decides membership in $K_{1/3}$ for every exact rational input. With an initial denominator $q > 0$, it visits at most $q + 1$ distinct in-range states before terminating or detecting a cycle.*

Proof. Within either retained branch, $y \in K_{1/3}$ is equivalent to the updated value belonging to $K_{1/3}$, by the coding identity $T\pi = \pi\sigma$. The endpoints $0, 1$ are in $K_{1/3}$, and the open central gap is disjoint from it. Thus all finite accept/reject conclusions are sound.

For cycle acceptance, repeating a finite orbit that avoids the gap produces an infinite valid sequence of branch choices. Reversing any finite prefix of those choices puts the initial point in the corresponding cylinder. It therefore belongs to every K_n and hence to $K_{1/3}$.

To prove termination, keep the original denominator q without reduction. Every update changes an integer numerator p to either $3p$ or $3p - 2q$. As long as the algorithm has not rejected, an in-range numerator belongs to $\{0, 1, \ldots, q\}$. A longer nonterminating run repeats a state. □

For example,

$$\frac{1}{3} \longmapsto 1 \quad \text{accepts,} \qquad \frac{1}{4} \longmapsto \frac{3}{4} \longmapsto \frac{1}{4} \quad \text{accepts,} \qquad \frac{1}{2} \quad \text{rejects immediately.}$$

The second example is the periodic address $(01)^\infty$, whose value is $A(01)/(3^2 - 1) = 2/8 = 1/4$.

The implementation returns a finite *witness*: the exact orbit, termination reason, decision, and any repeated-state index. A separate checker replays the transitions without calling the search routine. The package also contains an independent long-division implementation; it agrees on all the tested rational inputs. This comparison is a finite regression check, not a substitute for Proposition 1.15.

Complexity, with the numeric denominator exposed. The above bound is $O(q)$ transitions and $O(q)$ stored states, not polynomial in the bit length $\log q$. Straightforward fixed-denominator integer updates take $O(\log q)$ bit work each, giving an $O(q \log q)$ arithmetic bound. This is a pseudopolynomial reference algorithm, not a claim of optimal rational-membership complexity. Floyd's cycle detection can reduce state storage, but an independently replayable certificate still needs an orbit or another succinct witness.

For a single address of length n, Horner encoding has n updates, each on an integer with $O(n)$ bits. Schoolbook accounting gives $O(n^2)$ bit work and $O(n)$ working bits. Whole-stage enumeration has an exponential output size. The delivered Python code emphasizes inspectability and exactness instead of optimal base-conversion or tree-enumeration kernels.

1.6.3 Why arbitrary-real membership cannot be decided

Suppose a real is given by a total procedure returning a rational r_n with $|r_n - x| \le 2^{-n}$. Assume, for contradiction, that a single algorithm decides $x \in K_{1/3}$ for every such valid procedure.

For a program e, define a real

$$x_e = \begin{cases} \frac{1}{2}\, 3^{-t}, & \text{if } e \text{ first halts at step } t \ge 1, \\ 0, & \text{if it never halts.} \end{cases}$$

This real has a uniformly computable fast approximation procedure. To answer query n, simulate e for $n + 1$ steps. If it halts, output the corresponding rational; otherwise output 0. Any later halting contribution is less than $3^{-(n+2)}/2 < 2^{-n}$, so the approximation guarantee is valid regardless of whether e halts.

If e halts at time t, then

$$3^{-(t+1)} < x_e < 2 \cdot 3^{-(t+1)},$$

which is the open middle gap in the cylinder with prefix 0^t. Hence $x_e \notin K_{1/3}$. If e never halts, then $x_e = 0 \in K_{1/3}$. A membership decider would therefore decide halting and nonhalting, a contradiction.

The sound replacement: certified nonmembership or unresolved. The complement of $K_{1/3}$ is the union of the exterior rays and the enumerable open middle-third gaps. Given a real approximation interval, certify nonmembership only when the entire closed enclosure lies inside one of those open sets. For $x \notin K_{1/3}$, this eventually succeeds because its distance to the closed set $K_{1/3}$ is positive. For $x \in K_{1/3}$, it may run forever. A fixed-depth procedure should return *unresolved*, not *member*, when it has not found a separating gap.

This is a useful interface for verified numerics: negative certificates are stable under an explicitly bounded input error, whereas proximity to a fractal set is not the same as membership in it.

1.7 Computer-science application II: automata, safety, and entropy

1.7.1 A worked finite-state constraint on infinite addresses

Consider binary execution traces in which two 1 bits are never adjacent. An automaton needs two nonrejecting states: q_0, meaning that the previous bit was 0 or no bit has yet been read, and q_1, meaning that it was 1. A third state is rejecting and absorbing.

Let Σ_A be the infinite accepted traces and $K_A = \pi(\Sigma_A)$. The set of rejected traces is a union of cylinders exhibiting the finite bad block, so Σ_A is closed. It is compact, and the coding maps

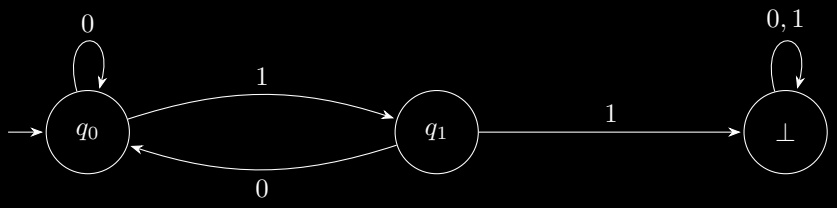


Figure 1.5: A safety monitor for the forbidden block 11. An infinite trace is accepted precisely when it never reaches $\perp$; this is a safety condition, not a finite-word acceptance convention left implicit.

it homeomorphically to K_A. No accepted trace is isolated: after any prefix, one can append a 0 and then choose either 0 or 1 before continuing with zeros. Thus this particular constrained state space is again Cantor. An arbitrary safety automaton need not have this property; one allowing only 0^∞ produces a singleton.

If a_n is the number of accepted length-n words, splitting according to the last bit gives

$$a_0 = 1, \quad a_1 = 2, \quad a_n = a_{n-1} + a_{n-2} \quad (n \geq 2).$$

Therefore $a_n = F_{n+2}$, where $F_0 = 0, F_1 = 1$. Equivalently the transition matrix is

$$A = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix},$$

whose eigenvalues solve $\lambda^2 - \lambda - 1 = 0$. Writing $\varphi = (1 + \sqrt{5})/2$ gives the word-growth rate

$$\lim_{n\to\infty} \frac{1}{n} \log a_n = \log \varphi. \tag{1.32}$$

This calculation quantifies the combinatorial reduction from 2^n to F_{n+2} feasible prefixes.

1.7.2 A geometric complexity exponent derived from the automaton

Proposition 1.16. *The constrained set has* $\dim_H K_A = \log \varphi / \log 3$.

Proof. Let $t = \log \varphi / \log 3$. Since $a_n \leq C\varphi^n$ for a fixed constant C, the cylinder cover has vanishing u-cost for every $u > t$: $a_n 3^{-nu} \leq C(\varphi 3^{-u})^n \to 0$.

For the lower bound use the positive right eigenvector $v = (\varphi, 1)^{\mathsf{T}}$ of A. Starting at the virtual state 0, assign transition probabilities $p_{ij} = A_{ij} v_j / (\varphi v_i)$. Their row sums are one because $Av = \varphi v$. The mass of an admissible word $b_1 \cdots b_n$ telescopes to

$$\nu(C_{b_1 \cdots b_n}) = \frac{v_{b_n}}{\varphi^n v_0} \leq \varphi^{-n}.$$

These consistent masses define a probability measure on the accepted traces and hence on K_A. An interval of length $3^{-(n+1)} < \ell \leq 3^{-n}$ meets at most two level-n cylinders, so $\nu(J) \leq 2\varphi^{-n} \leq 2\varphi \ell^t$. The same cover argument as in Theorem 1.6 proves $\dim_H K_A \geq t$. □

Validity analysis. The measure ν is not the fair Cantor measure conditioned on K_A by ordinary division. In fact $\mu(K_A) = \lim_n F_{n+2}/2^n = 0$. The transition construction specifies a new probability law. Its initial virtual state also need not have the stationary state distribution. These points matter when using constrained traces to calculate probabilities, not just dimensions or feasible-state counts.

1.7.3 A sound abstraction of the shift

For a nonempty word $w = b_1 \cdots b_n$, deletion of the first bit gives the exact identity

$$T(C_w) = C_{b_2 \cdots b_n} = C_{b_2 \cdots b_n 0} \cup C_{b_2 \cdots b_n 1}.$$

At fixed depth n, an abstract transition therefore shifts the address and appends either possible next bit. A safety monitor can be combined with this transition system by a product construction. Under a regular constraint, only monitor-compatible successors are retained.

The soundness statement must name its concrete domain: it describes points already represented by valid infinite addresses in $K_{1/3}$. A bare real interval I_w contains points that eventually fall into a removed gap. Treating every point of that interval as a valid symbolic execution would discard a possible escape. An interval implementation must retain an escape or unresolved state unless membership has separately been established. This is the distinction between a sound abstraction and a convenient but unjustified discretization.

1.8 Computer-science application III: certified integration and optimization

1.8.1 A quadrature rule with a computable stopping condition

For a function $f : K_{1/3} \to \mathbb{C}$, select the left endpoint $\ell_w = A(w)/3^n \in K_{1/3}$ and put

$$Q_n(f) = 2^{-n} \sum_{|w|=n} f(\ell_w). \tag{1.33}$$

Integrating the cylinderwise approximation in Proposition 1.11 gives

$$\left| \int f \, d\mu - Q_n(f) \right| \le \omega_f(3^{-n}) \le L 3^{-\alpha n} \tag{1.34}$$

when $[f]_\alpha \le L$. To reach error $\varepsilon < L$, take $n = \lceil \alpha^{-1} \log_3(L/\varepsilon) \rceil$. The number of evaluations is at most $2(L/\varepsilon)^{s/\alpha}$, with $s = \log 2 / \log 3$. This is a cost for this rule under a supplied regularity bound, not a universal lower bound or an optimality assertion.

For a real monotone function defined on $[0, 1]$, the rule admits a direct interval certificate. If f is increasing, then

$$Q_n^-(f) := 2^{-n} \sum_w f(\ell_w) \le \int f \, d\mu \le 2^{-n} \sum_w f(r_w) =: Q_n^+(f). \tag{1.35}$$

For decreasing f, exchange the endpoints. This inequality is obtained by integrating the pointwise bounds on each cylinder; it does not replace μ by normalized Lebesgue measure on K_n.

1.8.2 Exact test case: the second moment

Set $h = 3^{-n}$. Averaging cylinder left endpoints gives $2^{-n} \sum_w \ell_w = (1 - h)/2$. Conditional on a cylinder, the random point is $\ell_w + hX'$ with $\mathbb{E}X' = 1/2$ and $\mathbb{E}(X')^2 = 3/8$. Expansion yields the exact identities

$$\frac{3}{8} - Q_n^-(x^2) = \frac{h}{2} - \frac{h^2}{8}, \qquad Q_n^+(x^2) - Q_n^-(x^2) = h, \tag{1.36}$$

$$\frac{3}{8} - Q_n^{\mathrm{mid}}(x^2) = \frac{h^2}{8}, \qquad Q_n^{\mathrm{mid}}(x^2) = 2^{-n} \sum_w (\ell_w + h/2)^2. \tag{1.37}$$

The midpoint improvement is not free: the midpoint of a cylinder is outside $K_{1/3}$. It is a legitimate evaluation here because x^2 is defined on the whole interval. For a function supplied only as an oracle on $K_{1/3}$, midpoint sampling may be an invalid query.

n	2^n	$Q_n^-(x^2)$	$Q_n^+(x^2)$	midpoint error
0	1	0	1	$\frac{1}{8}$
1	2	$\frac{2}{9}$	$\frac{5}{9}$	$\frac{1}{72}$
2	4	$\frac{26}{81}$	$\frac{35}{81}$	$\frac{1}{648}$
3	8	$\frac{260}{729}$	$\frac{287}{729}$	$\frac{1}{5832}$
4	16	$\frac{2420}{6561}$	$\frac{2501}{6561}$	$\frac{1}{52488}$
5	32	$\frac{22022}{59049}$	$\frac{22265}{59049}$	$\frac{1}{472392}$

Table 1.1: Exact rational integration results generated by the accompanying code. The target value is the independently derived moment 3/8.

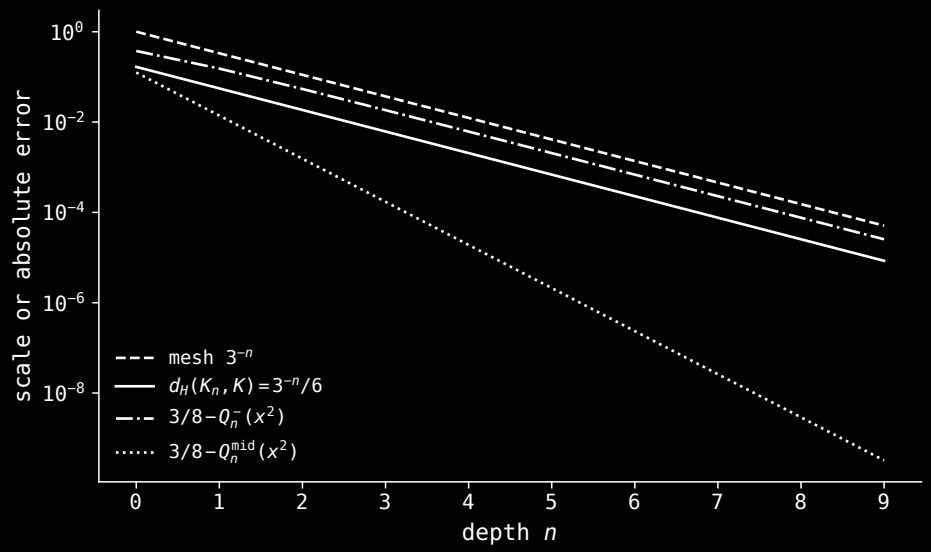


Figure 1.6: Four genuinely different quantities: cylinder mesh, set-approximation error, the exact left-rule error for x^2, and the exact midpoint error for x^2. The second-order midpoint slope comes from (1.37), not from a fitted regression. A shared depth does not make these error notions interchangeable.

1.8.3 A nonpolynomial observable with outward-rounded bounds

For $f(x) = 1/(1+x)$, the function is decreasing on $[0,1]$ and $|f'(x)| \le 1$. Exact rational endpoint evaluations therefore produce

$$2^{-n} \sum_w \frac{1}{1+r_w} \le J := \int \frac{1}{1+x}\, d\mu(x) \le 2^{-n} \sum_w \frac{1}{1+\ell_w}.$$

The interval width is at most 3^{-n}. No closed-form value of J is assumed. The following displayed decimal bounds are rounded outward; the unrounded rational bounds are retained in `results/convergence.csv`.

1.8.4 A complete nonconvex optimization certificate

Consider the optimization problem

$$\min_{x \in K_{1/3}} \left(x - \frac{1}{2}\right)^2. \tag{1.38}$$

The convex relaxation $x \in [0,1]$ returns the incorrect feasible answer $x = 1/2$ and the lower bound 0. At the first Cantor refinement, $K_{1/3} \subset [0,1/3] \cup [2/3,1]$, every feasible point has

n	nodes	lower bound	upper bound
0	1	0.5000000000	1.0000000000
2	4	0.6781250000	0.7368983958
4	16	0.7036084884	0.7101453050
6	64	0.7065058700	0.7072321919
8	256	0.7068285826	0.7069092851
9	512	0.7068554823	0.7068823832

Table 1.2: A certified enclosure of a nonpolynomial expectation. Lower decimals are rounded down and upper decimals up, so the printed interval remains valid.

$|x - 1/2| \geq 1/6$. Both $1/3$ and $2/3$ belong to $K_{1/3}$, so

$$\min_{x \in K_{1/3}} (x - 1/2)^2 = \frac{1}{36}, \qquad \operatorname{argmin} = \{1/3, 2/3\}. \tag{1.39}$$

This is a complete lower-bound and feasible-witness certificate, not an optimizer trajectory that happened to converge near a minimizer.

More generally, if f is α-Hölder with constant L and $U_n = \min_{|w|=n} f(\ell_w)$, compactness gives an attained minimum m and

$$U_n - L3^{-\alpha n} \leq m \leq U_n.$$

The upper bound uses feasible sample points. For the lower bound, take a minimizer, locate its cylinder, and compare it with that cylinder's left endpoint. An adaptive branch-and-bound routine can prune a cylinder only when a sound lower bound there exceeds its current feasible upper bound. A midpoint outside $K_{1/3}$ is not a feasible incumbent, even when it is a useful evaluation for a smooth extension of the objective.

1.9 Computer-science application IV: topology of finite observations

1.9.1 Exact gap statistics and persistent components

For a compact $E \subset \mathbb{R}$ and $r \geq 0$, define its closed offset $E_r = \{x : \operatorname{dist}(x, E) \leq r\}$. Two neighboring offset components merge when the gap between the original pieces has length at most $2r$. In the classical Cantor construction, the gaps removed at generation $j \geq 1$ have length 3^{-j} and multiplicity 2^{j-1}. Therefore, for every $r > 0$,

$$\#\pi_0((K_{1/3})_r) = 1 + \sum_{\substack{j \geq 1 \\ 3^{-j} > 2r}} 2^{j-1}. \tag{1.40}$$

The sum is finite. The strict inequality is necessary because a closed gap is filled when its two offset fronts meet. For $r = 0$, there are uncountably many components; formula (1.40) is not a finite algorithm at that endpoint.

For the endpoint sample F_n, sort the 2^{n+1} sample points. Their adjacent spacings comprise the original gaps of generations 1 through n, plus 2^n cylinder-length spacings of size 3^{-n}. The total count is

$$\sum_{j=1}^{n} 2^{j-1} + 2^n = 2^{n+1} - 1.$$

Thus its degree-zero offset barcode, over any field, consists of one essential component and one finite bar dying at half each adjacent spacing. For example, F_4 has one death at $1/6$, two at $1/18$,

four at 1/54, and 24 at 1/162. This finite barcode is exactly computable from rational numbers, without estimating distances from image pixels.

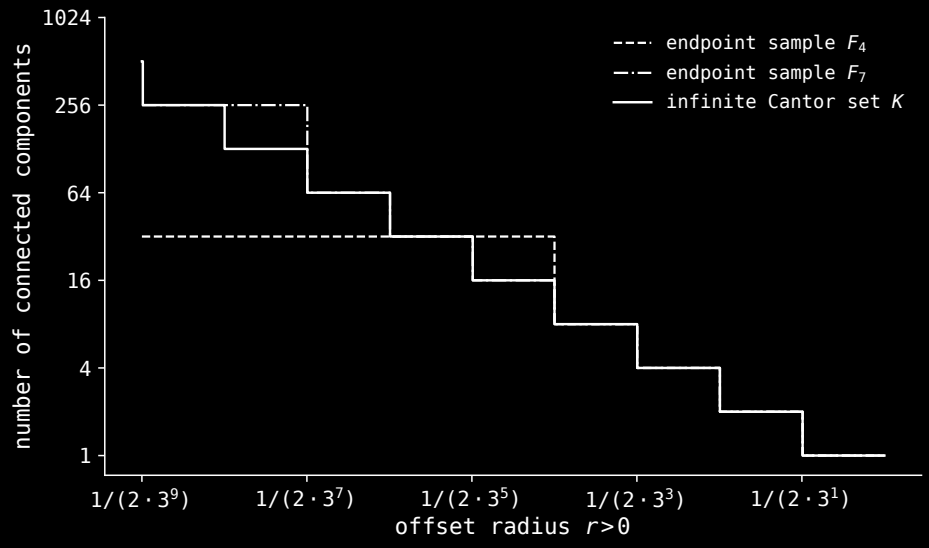


Figure 1.7: Positive-scale component counts for the infinite set and two finite endpoint samples. The step radii are prescribed by exact rational gaps. A finite sample can lose unresolved small-scale components and can also introduce extra components inside a sampled cylinder. Hausdorff proximity controls the scale of the discrepancy, not equality of component counts at every fixed radius.

1.9.2 The approximation theorem connecting data to topology

If $d_H(E, F) \le \delta$, then the triangle inequality gives

$$E_r \subset F_{r+\delta}, \qquad F_r \subset E_{r+\delta} \quad \text{for every } r \ge 0. \tag{1.41}$$

Indeed, a point within r of E lies within $r + \delta$ of F by first moving to a nearest point of E and then within δ of F. Applying homology to these inclusions supplies the usual scale-shift comparison of persistent components. Here $\delta = d_H(F_n, K_{1/3}) = 3^{-(n+1)}$ is explicit.

The infinite-set discussion is deliberately phrased at positive scales. The uncountable rank at radius zero is not replaced by an assertion that an ordinary finite barcode completely represents $H_0(K_{1/3})$. Nor do these one-dimensional offsets acquire genuine positive-degree cycles: they are unions of intervals. A software pipeline reporting a robust loop in this particular one-dimensional offset model has changed the model or introduced an error; it has not found hidden genus in the Cantor set.

1.10 A bridge back to quadratic dynamics: a certified real repeller

The source's broader subject is the quadratic family $f_c(z) = z^2 + c$ [3, Preface]. The following is an additional worked bridge, not a claim that the construction occurs in source §1.1.1 or a replacement for the later Fatou–Julia theory.

Take the real map $f(x) = x^2 - 3$ and the rational interval $I = [-B, B]$ with $B = 231/100$. Its inverse branches on I are

$$g_0(y) = -\sqrt{y + 3}, \qquad g_1(y) = \sqrt{y + 3}.$$

They are real-valued because $3 - B = 69/100 > 0$. Moreover, $B^2 > 3 + B$ implies $g_i(I) \subset I$, and

$$\sqrt{3 - B} > \frac{83}{100}, \qquad |g_i'(y)| \le \frac{1}{2\sqrt{3 - B}} < \frac{50}{83} < 1.$$

Their images are disjoint, separated by a gap containing $[-83/100, 83/100]$.

Define

$$K_f = \bigcap_{n\geq 0} f^{-n}(I),$$

where preimages are taken in $\mathbb{R}$. Because $f^{-1}(I) \subset I$, these are nested compact sets. Every inverse cylinder $g_{b_1} \circ \cdots \circ g_{b_n}(I)$ has diameter at most

$$2B\left(\frac{50}{83}\right)^n. \tag{1.42}$$

The two inverse branches supply persistent splitting, so the interval-tree argument makes K_f a Cantor set coded by binary sequences. On K_f, $|f'(x)| = 2|x| > 166/100 > 1$, which explains the term *expanding real repeller*.

This application uses the same computational pattern as the classical construction but changes its arithmetic: the inverse maps are nonlinear. The delivered `code/repeller.py` encloses square roots on a rational grid using integer square roots. For $u \geq 0$ and grid denominator D, put

$$k = \left\lfloor \sqrt{\lfloor uD^2 \rfloor} \right\rfloor.$$

Then $k/D \leq \sqrt{u} < (k+1)/D$, unless the lower endpoint is exact, in which case equality can be detected by squaring. Squaring the two rational bounds independently verifies each enclosure. The negative branch reverses their order.

Using $D = 10^{10}$, the finite cylinder for the word 00110101 is enclosed by

$$[-0.8741166948,\ -0.8733722911].$$

These decimal endpoints are exact grid rationals, and the rounding directions are outward. If a normalized inverse-step rounding error is at most D^{-1}, its propagated endpoint error is bounded by

$$\frac{1}{D}\sum_{j=0}^{n-1}(50/83)^j \leq \frac{1}{D(1-50/83)}.$$

One must account for this arithmetic enclosure error separately from the geometric cylinder diameter in (1.42).

Validity analysis. The negative inverse branch reverses orientation. As a result, ordinary lexicographic order of the fixed branch labels is not the spatial order at every node: the local ordering must track orientation parity. For example, $g_0 g_1(I)$ lies to the left of $g_0 g_0(I)$. Blindly reusing the increasing middle-third address order would be incorrect. Also, this section proves a real invariant-repeller statement; it does not claim that the accompanying finite calculation identifies an entire complex Julia set or settles a parameter-plane problem.

Chapter 2

Proper holomorphic maps, degree and lifting

2.1 Local analytic models and proper maps

Mathematical reference. Exercises 3.1–3.3; manuscript p. 106.

2.1.1 The local power model: solution of Exercise 3.1

Theorem 2.1 (M01: Local holomorphic normal form)**.** *Let $f : X \to Y$ be nonconstant and holomorphic. At any $a \in X$ there are holomorphic coordinates u at a and v at $f(a)$ for which $v \circ f \circ u^{-1}(\zeta) = \zeta^e$, for a unique integer $e \geq 1$.*

Proof. Choose preliminary coordinates with $a = 0 = f(a)$. The Taylor series has a first nonzero coefficient: $f(z) = z^e g(z)$, where g is holomorphic and $g(0) \neq 0$. Shrink to a disk on which g has no zero. A logarithm exists there, so $q(z) = \exp(\log g(z)/e)$ is a nonvanishing holomorphic eth root. Set $u(z) = zq(z)$. Then $u'(0) = q(0) \neq 0$, hence u is a coordinate on a smaller disk, and $f(z) = u(z)^e$. The exponent is invariant under invertible coordinate changes because they have nonzero first derivative. This is exactly the orientation-preserving topological local model required in the source. □

The local index is e; the derivative has a zero of order $e-1$. The distinction matters already for $f(z) = z^7$: one critical location contributes six to a ramification sum and seven to the multiplicity of the fibre over zero.

Algebraic-geometric derivation. In completed local rings, the map is $\mathbb{C}[[v]] \to \mathbb{C}[[u]]$, $v \mapsto u^e$. The local fibre is $\mathbb{C}[[u]]/(u^e)$, of complex vector-space length e. Its underlying set has one point, but its nilpotents retain the multiplicity lost by set-theoretic counting. A basis is $1, u, \ldots, u^{e-1}$. This is a genuine local algebraic model; it does not assert that every open Riemann surface is a projective algebraic curve.

Algebraic-topological derivation. On a small punctured neighbourhood, the map is an e-sheeted covering. Its map on $\pi_1 \cong \mathbb{Z}$ is multiplication by e. A positive generator in the source traces the target circle e times. Analytic orientation makes the sign positive. For an orientation-reversing local model $\bar{z}^e$, the same unsigned fibre count would conceal index $-e$.

Functional-analytical interpretation. For a germ $h(u) = \sum_{j\geq 0} c_j u^j$, average over the e local inverse branches:

$$\mathcal{T}h(v) = \frac{1}{e} \sum_{u^e=v} h(u) = \sum_{k\geq 0} c_{ek} v^k.$$

The roots-of-unity sum removes all powers not divisible by e. Individual branches are multivalued; their symmetric average is holomorphic through the critical value. This local cancellation is the first bridge to stable observable computation. The series identity holds on a sufficiently small disk of convergence. It does not justify averaging divergent series.

Worked problem: A critical point that is not a singular surface

For $f(z) = z^3(1+z)$ near zero, take $u = z(1+z)^{1/3}$ with the root equal to one at zero. The source surface is smooth, while the map has local degree three. The ring fibre has length three and the differential vanishes to order two. None of these facts says that the graph $w = z^3(1+z)$ is singular: its equation has derivative -1 with respect to w. Criticality of a projection and singularity of its total space are different claims.

2.1.2 Properness turns local multiplicities into a global degree

Theorem 2.2 (M02: Finite, constant weighted fibre cardinality)**.** *For a nonconstant proper holomorphic map $f : X \to Y$, the sum $\sum_{a\in f^{-1}(b)} e_a$ is finite, locally constant in b, and hence independent of b. The map is surjective.*

Proof. A fibre is discrete by M01 and compact by properness, so it is finite. Take disjoint coordinate neighbourhoods of its points. A sufficiently small neighbourhood of b has no other preimages: otherwise points outside these neighbourhoods with images converging to b would have a convergent subsequence by properness, producing an omitted point of the fibre. Within the chosen neighbourhoods, M01 supplies exactly e_a preimages with multiplicity. Thus the weighted count is locally constant. The image is open by M01 and closed by properness; connectedness of the target makes it all of Y. □

The proof uses a *compact discrete fibre*, not the false assertion that every discrete subset of a compact space is finite. The set $\{1/n : n \geq 1\}$ is a counterexample to the latter; it is not closed in $[0, 1]$. A fibre is closed because it is the inverse image of a closed point.

Applicability. The map $e^z : \mathbb{C} \to \mathbb{C}^*$ is locally biholomorphic but not proper: a point has infinitely many preimages. The inclusion $\mathbb{D} \hookrightarrow \mathbb{C}$ is injective and holomorphic but is neither proper nor surjective. A finite plot of an image establishes neither property.

2.1.3 Ideal ends: a hypothesis-aware solution of Exercise 3.2

The source's full ideal compactification is introduced in Theorem 2.59 for a finitely generated Fuchsian group, hence a finite-type hyperbolic surface. Cusps receive points and noncuspidal tame ends receive circles. Exercise 3.2 suppresses this inherited restriction in its wording. Here it is kept explicit; arbitrary infinite-type surfaces are not declared to have the same compactification. Source Lemma 5.7 revisits the tame-end statement on p. 118. [3, pp. 80–81,118]

Theorem 2.3 (M03: Extension over finite-type ideal ends)**.** *Let $f : X \to Y$ be a finite-degree proper holomorphic map between finite-type hyperbolic surfaces, with their standard ideal point/circle completions. It extends continuously to those completions, holomorphically through filled cusp points and locally as an unbranched covering on ideal boundary circles.*

Proof. First isolate every end by a peripheral collar. Properness sends sufficiently far-out components into an end collar of the target and gives a proper restriction on those components. The topological finite-type hypothesis rules out unbounded accumulation of independent ramification handles toward an end; equivalently one may first use the finite-type Riemann–Hurwitz argument M08 on an exhaustion to see that only finitely many ramification contributions occur. Remove those finitely many critical values by shortening the collars. Each remaining end map is an unbranched finite covering of annuli.

The classification of annular coverings now gives logarithmic-strip coordinates in which the map is a finite power covering, up to conformal changes of the collar coordinates. Covering moduli are scaled by the degree; in particular a puncture end corresponds to a puncture end. At a puncture its local expression is $u \mapsto u^e$ times a nonvanishing unit and it extends by removable singularities. At a noncuspidal circular end, the strip/power model extends continuously to the boundary line/circle and has nonzero tangential derivative in suitable conformal charts. The finitely many extensions agree with f on their common interiors, so they patch uniquely. Properness on the compact completions and the local models give the asserted branched covering. □

Dependency audit. This proof uses finite-type topology and annular uniformization, not merely the local inverse theorem. The exhaustion argument behind the finiteness assertion can be made explicit: choose compact target subsurfaces with boundary avoiding critical values, include all source components over them, and apply the lifted-cell count. As genus and the number of ends of the whole source are finite, the Euler defect cannot increase without bound; each new ramification contribution is nonnegative. This is a topological argument before it is an analytic extension argument.

Why a broader claim is unsafe. A finite-degree map can have infinitely many critical points on an infinite-type source. Locally defined double covers $w^2 = h(z)$, with a holomorphic h having infinitely many simple zeros accumulating only at a boundary point of $\mathbb{D}$, illustrate the obstruction: each zero is a branch value of the proper degree-two projection to $\mathbb{D}$. Over a compact subset the projection is proper because h is bounded there, but critical values need not be finite globally. For example a suitable nonzero constant multiple of $\sin(1/(1-z))$ has such a zero sequence in $\mathbb{D}$. This does not contradict M03, because the source is not in its finite-type class. It prevents silently applying the compact boundary-circle picture to every finite-degree map.

Three perspectives. Algebraically, filled punctures restore valuations at the missing places of a finite function-field extension. Topologically, the covering of a peripheral circle determines the local degree at the filled end. Analytically, boundedness in a cusp coordinate permits removal of a singularity. The circle and point completions encode different data; replacing both by a point is not the same compactification.

2.1.4 Finite Blaschke products: Exercise 3.3(i)

Theorem 2.4 (M04: Classification of proper disk maps)**.** *A proper holomorphic map $B : \mathbb{D} \to \mathbb{D}$ of degree d has the form*

$$B(z) = \lambda \prod_{j=1}^{d} \frac{z - a_j}{1 - \overline{a_j} z}, \qquad |\lambda| = 1, \quad a_j \in \mathbb{D},$$

where zeros are repeated with their multiplicities. Conversely every such product is a proper degree-d disk map.

Proof. The zero fibre has exactly d points with multiplicity by M02. Divide B by the displayed product without λ. Removable singularities at the zeros give a nonvanishing holomorphic quotient q. Properness implies $|B(z)| \to 1$ uniformly as $|z| \to 1$: failure would give a sequence escaping every compact subset whose images remain in a fixed compact subdisk. Each elementary factor also has boundary modulus one and no pole on $\overline{\mathbb{D}}$. Therefore $|q| \to 1$ at the boundary. Apply the maximum principle on concentric disks to q and $1/q$ to obtain $|q| \leq 1$ and $|q| \geq 1$. A holomorphic function with constant modulus is constant, giving $q = \lambda$.

For the converse, the identity

$$|1 - \bar{a}z|^2 - |z - a|^2 = (1 - |a|^2)(1 - |z|^2)$$

places each factor inside the disk. The product extends across the unit circle with modulus one there. Compact inverse images stay away from the boundary, hence the map is proper. Its zero multiplicity is exactly d, establishing its degree. □

Recorded source repair. The source product runs from $k = 0$ to d, giving $d + 1$ factors, while the map is said to have degree d. The corrected bounds are $1, \ldots, d$ (or $0, \ldots, d - 1$). The degree-one test $a_j = 0$ detects the discrepancy immediately. This is a visible indexing issue, not an assertion that the classification itself is wrong.

Algebraic geometry. The product is a rational map $\mathbb{P}^1 \to \mathbb{P}^1$. A zero a_j in the disk has its reflected pole at $1/\bar{a}_j$ outside, with infinity included when $a_j = 0$. The divisors of zeros and poles each have degree d. On homogenizing numerator and denominator, common factors must be cancelled before declaring the projective degree; for $|a_j| < 1$ the reflected locations cannot cancel the disk zeros. At $|a| = 1$ cancellation does occur and the factor becomes constant after simplification. Degree is not continuous across that excluded boundary of parameter space.

Algebraic topology. On $\mathbb{T}$, the restriction has degree d. If $B(e^{it}) = e^{i\beta(t)}$ with a continuous argument lift, then

$$\beta'(t) = \sum_{j=1}^{d} \frac{1 - |a_j|^2}{|e^{it} - a_j|^2} > 0, \qquad \beta(t + 2\pi) - \beta(t) = 2\pi d.$$

Differentiate the logarithm of each factor, or verify directly by multiplying numerator and denominator by conjugates. Positivity means that the boundary has no critical point. The integral is a global winding budget; local phase can concentrate sharply near a zero approaching the circle without changing that budget.

A quantitative consequence. If $|a_j| \leq \rho < 1$ for all j, then

$$d\frac{1 - \rho}{1 + \rho} \leq \beta'(t) \leq d\frac{1 + \rho}{1 - \rho}.$$

The bounds follow from $1 - |a| \leq |e^{it} - a| \leq 1 + |a|$. They certify a boundary conditioning envelope, not an equally strong Euclidean inverse-derivative bound at interior branch points.

2.1.5 Half-plane rational maps: Exercise 3.3(ii)

Theorem 2.5 (M05: Positive-residue representation)**.** *A proper holomorphic map $F : \mathbb{H} \to \mathbb{H}$ of degree d with $F(\infty) = \infty$ has the rational representation*

$$F(z) = \lambda_0 z + a_0 - \sum_{j=1}^{d-1} \frac{\lambda_j}{z - a_j}, \quad \lambda_j > 0, \quad a_0, a_j \in \mathbb{R},$$

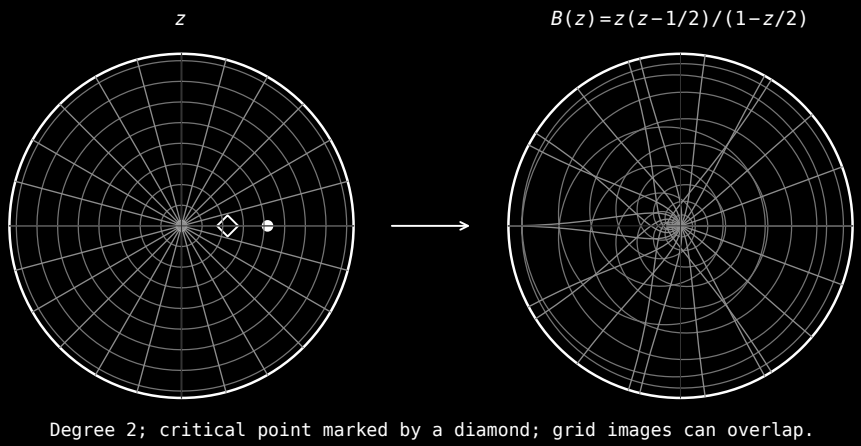


Figure 2.1: An original image of a polar grid under a degree-two Blaschke map. The plot records a many-to-one image, not a globally injective conformal coordinate change. The critical point is computed independently of the pixels.

where the finite poles a_j are distinct. Conversely this representation defines such a map.

Proof. Conjugating by a Cayley map reduces to M04, so F is rational and extends meromorphically across the real circle. Reflection makes it real on the real line away from poles. A pole on the real boundary must be simple: a pole of order $m > 1$ makes the leading imaginary part change sign as a small upper semicircle sweeps through angles $0 < \theta < \pi$. For a simple pole, the residue is negative, since $\operatorname{Im}[-\lambda/(z-a)] = \lambda \operatorname{Im} z/|z-a|^2$. The same argument at infinity shows that its pole is simple and that its linear coefficient is positive. Subtracting its principal parts leaves a real constant. There are d poles on the sphere with multiplicity, one at infinity, hence $d-1$ finite poles.

Conversely, for $z = x + iy$, $y > 0$,

$$\operatorname{Im} F(z) = y\left(\lambda_0 + \sum_{j=1}^{d-1} \frac{\lambda_j}{(x-a_j)^2 + y^2}\right) > 0.$$

The boundary values lie on the extended real line, the degree is d, and compact inverse images cannot approach a real pole, infinity, or a nonpole real point while their images remain in a compact subset of $\mathbb{H}$. Thus F is proper. □

A repeated pole location must be combined into a single residue. Counting two entries at the same location as two poles incorrectly inflates the degree. Positivity of imaginary part is not, by itself, properness: $F(z) = z + i$ maps $\mathbb{H}$ into $\mathbb{H}$ but is not onto it. The EE volume distinguishes lossless rational reactance models of this form from general dissipative positive-real models.

2.1.6 Functional analysis added to 3.1: the analytic cost of a zero

Definition 2.6. The Hardy space $H^2(\mathbb{D})$ consists of $f(z) = \sum_{n\geq 0} c_n z^n$ with $\sum |c_n|^2 < \infty$, with inner product linear in the first variable. Boundary norms are interpreted by radial limits, equivalently by the coefficient norm. Write $M_B f = Bf$ and $K_B = H^2 \ominus BH^2$.

Theorem 2.7 (M06: Inner multiplication, defect dimension and index)**.** *For a finite Blaschke product of degree d, M_B is an isometry, its range is closed,* $\dim K_B = d$*, and its Fredholm index is $-d$.*

Proof. Boundary modulus one gives $\|Bf\|_2 = \|f\|_2$, first for polynomials and then by completion. An isometry has closed range and zero kernel. For an elementary factor $b_a = (z-a)/(1-\bar{a}z)$, its range consists exactly of the functions vanishing at a. One inclusion is immediate. For the other, divide a function with a zero at a by b_a: the apparent pole is removable and its radial boundary norm remains finite. Evaluation at a is a continuous nonzero functional by Cauchy–Schwarz,

$$|f(a)| \le \frac{\|f\|_2}{\sqrt{1-|a|^2}}.$$

Thus K_{b_a} is one dimensional, spanned by $k_a(z) = (1-\bar{a}z)^{-1}$. Finally

$$K_{B_1 B_2} = K_{B_1} \oplus B_1 K_{B_2}$$

is an orthogonal decomposition: decompose $H^2 = K_{B_1} \oplus B_1 H^2$ and then the second summand. Induction gives dimension d. Kernel dimension zero minus cokernel dimension d gives the index. □

This is an operator-theoretic re-expression of the same degree, but its proof also specifies the space and its norm. On a finite truncation, boundary effects can destroy an apparent isometry. Replacing the infinite shift by an $N \times N$ matrix creates a nontrivial kernel; its ordinary finite-dimensional index is zero. An array of singular values is therefore not a proof of the Fredholm index above.

For $B(0) = 0$, composition $C_B f = f \circ B$ is an isometry on H^2: the powers $1, B, B^2, \dots$ are orthonormal, since the inner product of two distinct powers reduces by boundary modulus one to the mean of a nonconstant power of B, namely its value at zero. Without $B(0) = 0$, orthogonality fails: $\langle B, 1\rangle = B(0)$. Multiplication and composition operators are not interchangeable.

Consolidated adversarial questions for 3.1. Can one delete properness? No, by the exponential and disk-inclusion examples. Can one replace multiplicity by a count of dots? No, at $z^e = 0$. Can a disk zero lie on the boundary without changing the degree? No, cancellation occurs. Does a proper map have to be univalent? No, a finite Blaschke product of degree greater than one supplies a counterexample. Does an isometry of one function space certify arbitrary sampled data processing? No, the norm and truncation contract have changed.

2.2 Riemann–Hurwitz from cells, divisors and energy

Mathematical reference. The unnumbered formula, Remarks 3.4–3.5 and Corollary 3.6; pp. 106–107.

2.2.1 Cellular proof with every defect accounted for

Theorem 2.8 (M07: Riemann–Hurwitz for closed surfaces)**.** *Let $f : X \to Y$ be a degree-d branched covering between closed connected oriented surfaces. Then*

$$\chi(X) = d\chi(Y) - R_f, \qquad R_f = \sum_{x \in X} (e_x - 1).$$

For closed surfaces of genera g_X, g_Y this is

$$2g_X - 2 = d(2g_Y - 2) + R_f.$$

Proof. The branch set is finite. Triangulate the target with every branch value as a vertex. On the interiors of edges and faces the covering is unbranched, giving dE_Y edges and dF_Y faces.

For a vertex y, let the distinct points in its fibre have local degrees $e_1, \dots, e_r$. The fibre degree identity says $\sum e_i = d$, so the number r of lifted vertices is $d - \sum(e_i - 1)$. Summing over target vertices gives $V_X = dV_Y - R_f$. Substitute these three identities into $V_X - E_X + F_X$. Replacing Euler characteristics by $2 - 2g$ proves the genus form. □

Why the lifted cells are legitimate. A sufficiently small triangle adjacent to a branch value lifts in sectors in the local power chart. Its lifted open faces are disks, and their closures glue along lifted edges. The resulting decomposition may be treated as a CW decomposition; one can subdivide further to obtain an ordinary triangulation. It is not necessary to pretend that a branch fibre contains d distinct vertices.

Permutation version. Label a regular fibre by $\{1, \dots, d\}$. Monodromy around a branch value is a permutation σ. A cycle of length e is a local degree-e point. If $c(\sigma)$ is the number of cycles, including fixed points, its total defect is $d - c(\sigma)$. For a sphere target with branch permutations $\sigma_1, \dots, \sigma_r$,

$$2 - 2g_X = 2d - \sum_{j=1}^{r} (d - c(\sigma_j)).$$

The product must be the identity with a consistent loop convention, and the generated group must act transitively for the source to be connected. An integer genus is only one acceptance condition; arbitrary permutations satisfying the Euler equation need not form the specified covering.

2.2.2 Algebraic-geometric derivation through the canonical divisor

Let X, Y be compact Riemann surfaces, viewed equivalently as smooth projective complex curves when the algebraic formulation is used. Choose a nonzero meromorphic differential ω on Y. Near x, choose coordinates with $v = u^e$ and write $\omega = v^m a(v)\, dv$, $a(0) \neq 0$. Pullback gives

$$f^*\omega = e\, u^{em+e-1} a(u^e)\, du.$$

Consequently

$$\operatorname{ord}_x(f^*\omega) = e_x \operatorname{ord}_{f(x)} \omega + (e_x - 1), \qquad K_X = f^* K_Y + \mathcal{R}_f,$$

where $\mathcal{R}_f = \sum(e_x - 1)[x]$ is the ramification divisor. Taking degrees, using $\deg f^* D = d \deg D$ and $\deg K_X = 2g_X - 2$, yields M07.

This derivation reveals information absent from the single integer R_f: the location and multiplicity of every ramification point. It also exposes a field hypothesis. Over characteristic zero, differentiating u^e produces a nonzero coefficient e. In positive characteristic dividing e, the derivative may vanish identically and the tame formula must be replaced by a separability/different calculation. The map $u \mapsto u^p$ in characteristic p is not a valid counterexample to the complex theorem; it is outside its field assumptions.

Independence accounting. The cell argument and divisor argument reorganize the same geometric fact. The second imports the canonical-degree theorem, so it is not counted as a second axiom-free proof from identical primitives. The perspective ledger records its distinct algebraic mechanism and its additional prerequisite instead of inflating a "number of proofs" metric.

2.2.3 Punctures, boundary components and the finite-type restriction

Theorem 2.9 (M08: Finite-type and bordered versions)**.** *For a proper finite-degree branched covering between connected oriented finite-type surfaces, with ordinary covering behaviour at genuine boundary circles, the same Euler-defect formula holds, with ramification counted at interior points of the surfaces under consideration.*

Proof. For punctured surfaces, fill each topological end by a point to obtain $\widehat{X}, \widehat{Y}$. Properness extends the map topologically at ends. If n_X, n_Y are their numbers of ends and R_{end} is the total ramification at filled end points, then the fibre count over those points is

$$n_X = dn_Y - R_{\text{end}}.$$

Apply the closed formula to the filled surfaces:

$$\chi(\widehat{X}) = d\chi(\widehat{Y}) - R_{\text{int}} - R_{\text{end}}.$$

Subtract n_X on the left and dn_Y on the right. The end contributions cancel, giving $\chi(X) = d\chi(Y) - R_{\text{int}}$.

For compact bordered surfaces, triangulate with boundary edges and branch values among the vertices. Boundary fibres are unbranched, so those vertices have d lifts. All edge and face counts scale by d, while the interior vertex defect remains R_{int}. Alternatively double the surfaces and the map across their boundaries; Euler characteristic and interior ramification both double, then divide the closed-surface identity by two. The doubling argument requires the stated boundary extension, not arbitrary maps of closed sets. □

Remarks 3.4 and 3.5. Ideal triangulations encode punctures by omitted vertices; including their correction recovers exactly the subtraction above. Removing the boundary of a compact surface leaves a homotopy-equivalent interior, so Euler characteristic does not change. These statements are correct controls once the conventions are fixed. The ideal *conformal* completion, in contrast, may retain noncuspidal boundary circles; it should not be confused with one point per topological end.

2.2.4 The disk criterion: solution and qualification of Corollary 3.6

Theorem 2.10 (M09: A branch budget characterizes a disk source)**.** *Under M08, with disk target, the connected source is a topological disk exactly when $R_f = d - 1$.*

Proof. The formula gives $\chi(X) = d - R_f$. A connected oriented finite-type surface with a nonempty end/boundary set has $\chi(X) = 2 - 2g - b$, where $b \geq 1$ counts ends or boundary components in the corresponding compact-core model. Therefore $\chi(X) = 1$ is equivalent to $g = 0, b = 1$. This is a disk, open or closed according to the chosen convention. Conversely a disk has Euler characteristic one and hence $R_f = d - 1$. □

Connectedness is essential. A disjoint union of a disk and an annulus also has total Euler characteristic one, but is not a disk. Orientability and finite-type hypotheses are also doing real work. The short source proof should be read within that ambient category, not as a classification of every space with Euler characteristic one.

Worked problem: Power maps on three different domains

On $\mathbb{P}^1$, $z \mapsto z^d$ has critical points zero and infinity, each of local degree d. Thus $R_f = 2d - 2$ and $2 = 2d - (2d - 2)$. On $\mathbb{C}$, infinity is deleted and $R_f = d - 1$, so $1 = d - (d - 1)$. On $\mathbb{C}^*$ both critical locations are deleted and $R_f = 0$, so $0 = d \cdot 0$. The polynomial did not change; the source and target surfaces did. A computer routine which always inserts or always omits infinity will fail at least one of these tests.

2.2.5 A smooth double cover: genus is determined after normalization

Consider the compact smooth model of

$$y^2 = (x^2 - 1)(x^2 - k^2), \qquad 0 < k < 1.$$

The projection to the x-sphere has degree two. Its four finite branch values $\pm 1, \pm k$ are simple, each contributing one. In the coordinate $t = 1/x$, $v = y/x^2$, the equation becomes

$$v^2 = (1 - t^2)(1 - k^2 t^2).$$

At $t = 0$ there are two distinct unramified points $v = \pm 1$, so infinity contributes nothing. Hence $2 - 2g = 4 - 4 = 0$ and $g = 1$.

The affine equation is not the complete compact surface. One must use a smooth compactification/normalization, not count an apparent singular point of an unsuitable plane homogenization as additional ramification. When $k \to 1$, branch points collide and the smooth-family hypotheses fail; the normalization of the limiting singular curve can have smaller genus. "Genus is topological" does not mean that genus stays constant through singular degeneration.

The differential dx/y is holomorphic on this genus-one curve. Near a simple branch point, use $x - a = u^2$ and $y = u\,q(u)$: then $dx/y = 2\,du/q(u)$ has no pole. Near infinity, $dx/y = -dt/v$ is also regular. This local computation is more informative than merely attaching the word elliptic to a quartic polynomial.

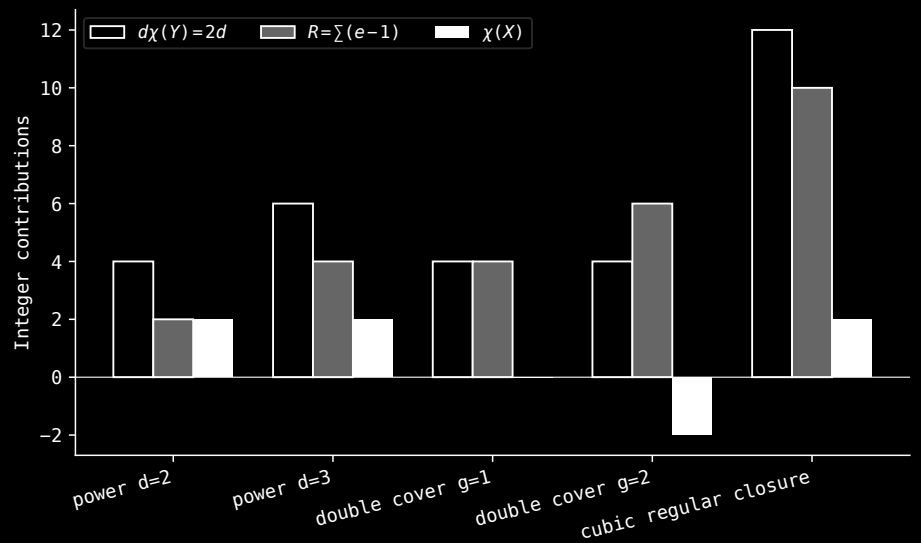


Figure 2.2: Branch-value cycle defects for five cover examples. The explicit table and the Euler identities, instead of the graph, establish the genus. Infinity is included in the compact-sphere examples.

2.2.6 Functional analysis: pullback of energy and holomorphic forms

Theorem 2.11 (M10: Degree is the energy multiplier)**.** *For a proper holomorphic degree-d map of planar domains, and a continuously differentiable real-valued u with finite Dirichlet energy on the target,*

$$\int_X |\nabla(u \circ f)|^2 \, dA = d \int_Y |\nabla u|^2 \, dA.$$

An analogous identity holds for square-integrable holomorphic one-forms with the intrinsic norm $\|\omega\|^2 = (i/2) \int \omega \wedge \bar{\omega}$.

Proof. Away from critical points, the real derivative of a holomorphic map is a rotation followed by scaling by $|f'|$. Therefore $|\nabla(u \circ f)|^2 = |f'|^2 |\nabla u|^2 \circ f$. The change-of-variables formula on the inverse branches sums d copies of the target integral. Critical points and values are discrete and have area zero, so they do not change the integral. Exhaustion and monotone convergence treat finite-energy functions without compact support. For forms, $f^*(\omega \wedge \bar{\omega}) = f^*\omega \wedge \overline{f^*\omega}$, and the same degree-counting integral argument applies. □

Thus $d^{-1/2} f^*$ is an isometry on the corresponding spaces of one-forms, and on Dirichlet spaces modulo constants. It is *not* an isometry of every unweighted function L^2 norm: ordinary area pullback introduces the Jacobian in a different place. It also does not automatically preserve physical conductivity or elapsed diffusion time. Those require the model and metric transformations specified in the EE volume.

An instructive consequence is an obstruction: if ω is a nonzero holomorphic one-form on a compact target, its pullback cannot vanish identically. Hence pullback injects the space of holomorphic differentials, consistent with the genus inequality implied by M07. This is an operator statement with a nonzero norm, not a pictorial inference about holes.

2.3 the topological argument principle

Mathematical reference. Exercise 3.7, Propositions 3.8–3.9 and Exercise 3.10; pp. 107–108.

2.3.1 Winding as a homological invariant: Exercise 3.7

Theorem 2.12 (M11: The punctured plane has one integral winding coordinate)**.** *Radial retraction onto a circle gives*

$$H_1(\mathbb{C} \setminus \{b\}; \mathbb{Z}) \cong H_1(S^1; \mathbb{Z}) \cong \mathbb{Z}.$$

The isomorphism sends a positive once-around loop to 1 *and agrees with the radian-normalized argument integral on smooth loops.*

Proof. Translate b to zero. The homotopy

$$H(z, s) = \big((1 - s) + s/|z|\big) z$$

never meets zero and retracts $\mathbb{C}^*$ onto the unit circle. The circle has a CW decomposition with one vertex and one oriented edge; both adjacent endpoints are that vertex, so the cellular boundary of the edge is zero. Thus its first homology is $\mathbb{Z}$. For a smooth circle parametrization, the integral is 1. The form $d\arg$ is closed on $\mathbb{C}^*$ and its integral is invariant under homotopy; subdividing loops into charts with continuous arguments identifies the integral with the change of a global argument lift divided by 2π. The generator normalization therefore fixes the same homomorphism on every cycle. Continuous loops are treated by uniform polygonal approximation in a neighbourhood avoiding zero and homotopy to that approximation. □

There is no global real-valued continuous argument on $\mathbb{C}^*$. If one existed, its increment around every closed loop would be zero, contradicting the generator calculation. A local branch of logarithm is compatible with this theorem; a globally single-valued logarithm on the punctured plane is not.

Algebraic geometry. For a meromorphic function, the same integer is the degree of its zero divisor minus its pole divisor inside a contour. This algebraic description requires meromorphy; it is not a replacement for signed topological index for arbitrary continuous maps. Multiplicity is again length in a local ring.

Functional analysis. The contour functional $f \mapsto (2\pi i)^{-1} \int f'/f$ is locally constant on the open set of nonvanishing boundary functions for which the expression is defined. A perturbation smaller than the boundary clearance stays in the same connected component of the nonvanishing-function space. This is the functional viewpoint behind a robust root-count certificate. It does not say the nonlinear map $f \mapsto f'/f$ is defined where f vanishes.

2.3.2 Nonzero boundary degree forces a value: Proposition 3.8

Theorem 2.13 (M12: Value-existence obstruction)**.** *Let D be a bounded Jordan disk and $f : \overline{D} \to \mathbb{C}$ continuous with $b \notin f(\partial D)$. If $\mathrm{wind}_b(f|_{\partial D}) \neq 0$, then $b \in f(D)$.*

Proof. Suppose b is absent from $f(D)$ as well. Then f maps the whole closed disk into $\mathbb{C} \setminus \{b\}$. Contract its boundary in the source disk and compose the contraction with f. This null-homotopes the boundary image in the punctured plane, forcing winding zero by M11, a contradiction. □

The map need not be analytic, differentiable or injective. Its boundary image need not be simple. Conversely winding zero does not prove that b is absent: local indices of opposite signs can cancel.

Worked problem: A zero with no locally surjective map

For $f(x + iy) = x^2 + iy$ and target $b = 0$, the origin is an isolated preimage. On a small circle the image has local index zero; the image lies in the right half-plane and does not contain any disk about zero. For a target $b > 0$ close to zero, there are two preimages $\pm\sqrt{b}$, with real Jacobian determinants $\pm 2\sqrt{b}$. Their indices add to zero. The two preimages are not detected by a nonzero-degree existence test because their signed contributions cancel.

2.3.3 The local-to-global sum: Proposition 3.9

Theorem 2.14 (M13: Boundary winding equals the sum of local indices)**.** *Under the hypotheses of M12, assume $f^{-1}(b)$ is discrete in D. Then the fibre is finite and*

$$\mathrm{wind}_b(f \circ \Gamma) = \sum_{a \in f^{-1}(b)} \mathrm{ind}_a(f),$$

where Γ is a positively oriented parametrization of ∂D.

Proof. The fibre is closed in $\overline{D}$ and avoids its boundary, so it is compact; discreteness makes it finite. Choose disjoint closed small disks about its points, contained in D and containing no other preimage of b. Let their positively oriented boundaries be γ_j. In the region obtained by deleting their interiors, the oriented boundary is $\Gamma - \sum_j \gamma_j$. Triangulate this region; cancellation

of interior edges shows that this boundary is a singular boundary. Thus $[\Gamma] = \sum_j [\gamma_j]$ in the first homology of $\overline{D} \setminus f^{-1}(b)$. Apply the map induced by f and then the winding isomorphism M11. The small-loop winding is the local index by definition. □

The source's short phrase about a "2-cell" is made precise here by a triangulated 2-chain with holes. The compactness claim is retained as a positive control after its closedness argument is supplied. Neither clarification changes the theorem.

For holomorphic f, M01 makes every local index a positive zero multiplicity, hence M13 becomes a zero-counting theorem. For meromorphic f, remove its poles as additional holes; a pole of order m contributes $-m$. This gives the analytic argument principle

$$\frac{1}{2\pi i} \int_\Gamma \frac{f'(z)}{f(z)}\, dz = N_Z - N_P.$$

The equation identifies an integer but does not certify a numerical quadrature output. A finite sample sum can miss a loop between samples or a zero on the path.

2.3.4 Quantitative local surjectivity: Exercise 3.10

Theorem 2.15 (M14: A boundary-clearance radius)**.** *Let a be an isolated preimage of $b = f(a)$ with nonzero local index. For every sufficiently small radius r with $\overline{D_r(a)} \subset D$ and no other preimage of b in that disk, set*

$$\delta = \min_{|z-a|=r} |f(z) - b| > 0.$$

Then $D_\delta(b) \subset f(D_r(a))$.

Proof. For $|b' - b| < \delta$, the straight homotopy between the two translated boundary loops $f(z) - b$ and $f(z) - b'$ stays nonzero, since its perturbation has magnitude smaller than the minimum boundary magnitude. Their winding numbers agree and are nonzero. M12 gives a preimage of b' inside the disk. Given a larger prescribed neighbourhood of a, choose such a sufficiently small r inside it. □

This proof corrects the repeated b in the source hint's comparison by explicitly introducing b'. For $f(z) = z^e$, the exact radius is $\delta = r^e$. Near a critical point a large input neighbourhood can therefore map to a much smaller guaranteed output neighbourhood. Local surjectivity does not imply a well-conditioned single-valued inverse.

A useful perturbation theorem. If a continuous perturbation g satisfies $\sup_{\partial D} |f - g| < \min_{\partial D} |f - b|$, the boundary counts for $f - b$ and $g - b$ agree. In the holomorphic case this is the homotopy form of Rouché's theorem. The strict inequality is important: at equality a zero can lie on the boundary. The CSE certificate checker produces witnesses for this type of nonvanishing homotopy on every subdivided contour segment.

2.3.5 Functional analysis: winding as a Fredholm index

Theorem 2.16 (M28: Polynomial winding and a Hardy-space index)**.** *Let p be a nonzero polynomial with no zero on $\mathbb{T}$, and let N count its zeros in $\mathbb{D}$ with multiplicity. Multiplication $M_p : H^2 \to H^2$ is Fredholm with index $-N$, and*

$$-\operatorname{ind}(M_p) = \operatorname{wind}_0(p|_{\mathbb{T}}).$$

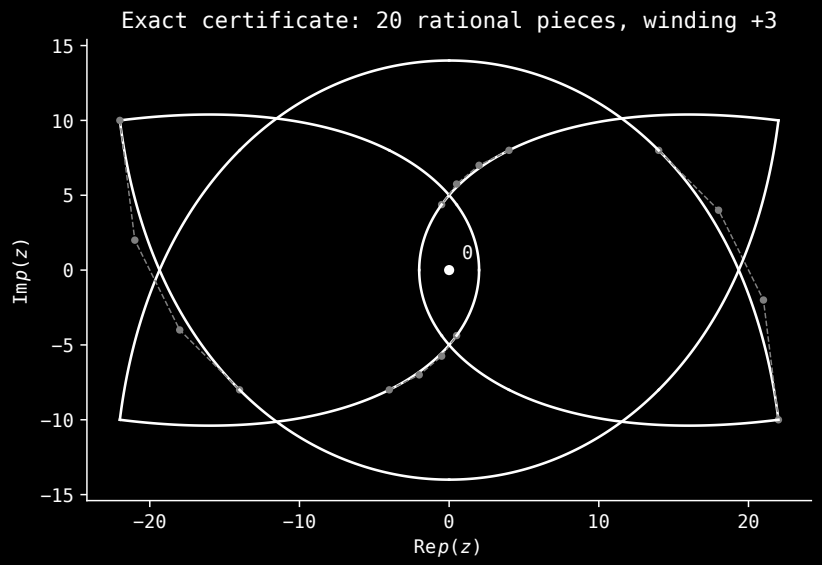


Figure 2.3: Polynomial images of a rational square contour, with certified subdivision vertices. The exact checker encloses the entire intervening polynomial arcs in open coordinate half-planes; the visible polygon alone is not its proof.

Proof. Factor the interior zeros into a finite Blaschke product B, repeated with multiplicity, and write $p = BG$. The factors $1 - \bar{a}z$ introduced in G, together with the remaining polynomial factors, have no zero on the closed disk. Thus G and $1/G$ are bounded holomorphic functions on the disk, and M_G is a bounded invertible operator. Consequently M_p has the same range codimension as M_B, namely N by M06, and has zero kernel. Its index is $-N$. The analytic argument principle, already established from M13, identifies N with the boundary winding. □

This is a genuinely functional-analytic rederivation of the same integer through the codimension of a closed operator range. It does not assert the general Toeplitz index theorem for arbitrary continuous symbols. The finite Blaschke factorization supplies the exact special case required here. A finite truncated multiplication matrix is not substituted for the infinite operator: truncation can change its kernel and range at the last rows.

2.4 an explicitly supplied degree subsection

Mathematical reference. The heading "Degree of proper maps" appears without subsequent content on p. 108. Everything in this section is an addition, linked to source Proposition 1.104 and section 3.3.

Theorem 2.17 (M15: Three compatible notions of degree)**.** *For a finite proper holomorphic map of oriented surfaces, the weighted fibre cardinality, the multiplier on local top homology, and the boundary winding about a regular value in a suitable compact target disk agree. Degrees multiply under composition.*

Proof. The weighted count is constant by M02. For a small disk about a target point, excision identifies the relative top homology of its preimage with the sum of the local relative groups. Each local power map multiplies its orientation class by its exponent, giving the same sum. Taking boundaries of relative classes gives the sum of small-circle windings, equal to the target-boundary degree by M13. For composition, group the preimages by the intermediate point. M01 and the chain of local power coordinates give $e_x(f \circ g) = e_x(g)e_{g(x)}(f)$. Summing yields $(\deg g)(\deg f)$. □

For an arbitrary smooth real map, the regular-value formula is instead $\sum \operatorname{sign} \det Df$. Properness is required to make this finite and stable under variation of the value. In an infinite-dimensional

Banach-space setting, a superficially similar degree theory requires additional compactness or Fredholm hypotheses. The finite surface result is not a license to use the same count for arbitrary neural operators.

The algebraic norm and trace. For a finite separable function-field extension, trace and norm are sums and products over the local conjugate branches. They are single-valued meromorphic functions downstairs. In the simplest quadratic algebra $A = \mathbb{C}[w][z]/(z^2 - w)$, multiplication by z has matrix

$$M_z = \begin{pmatrix} 0 & w \\ 1 & 0 \end{pmatrix} \quad \text{in the basis } (1, z).$$

It satisfies $M_z^2 = wI$, $\operatorname{tr} M_z = 0$ and $\det M_z = -w$. The matrix and its traces are single-valued even though its individual eigenvalues $\pm\sqrt{w}$ exchange along a loop. This is the computational interface used by the CSE branches, not an inference that eigenvectors are nonsingular at $w = 0$.

An independent functional warning. A spectral projection P satisfying $P^2 = P$ need not identify the intended spectral subspace: zero and identity are also idempotents. A root-count or degree certificate therefore needs a domain and a rank, not merely an algebraic residual.

2.5 lifting is a compatibility condition

Mathematical reference. Lemma 3.11 and Exercises 3.12–3.13; pp. 108–109.

2.5.1 The subgroup criterion, including the base point

Theorem 2.18 (M16: Lifts are controlled by covering subgroups)**.** *For connected coverings $p : X \to Y$ and $\widetilde{p} : \widetilde{X} \to \widetilde{Y}$, and a homeomorphism $h : Y \to \widetilde{Y}$, choose $x_0 \in X$ and $\widetilde{x}_0$ over $h(p(x_0))$. A based lift of $h \circ p$ exists exactly when*

$$h_* p_* \pi_1(X, x_0) \subseteq \widetilde{p}_* \pi_1(\widetilde{X}, \widetilde{x}_0).$$

It is unique once its value at x_0 is fixed. If the subgroups are equal, the lift is a homeomorphism.

Proof. A lift takes every closed path at x_0 to a closed path at $\widetilde{x}_0$, proving necessity. For sufficiency, join x_0 to x by a path, lift its image under $h \circ p$ starting at $\widetilde{x}_0$, and define $H(x)$ as the endpoint. Two chosen paths differ by a based loop; the subgroup inclusion says its image lifts closed, so the endpoint is independent of the choice. Evenly covered neighbourhoods give continuity and the local lifting identity. A lift agreeing at one point agrees in an open and closed subset of the connected source, hence everywhere. Equality allows the same construction for the inverse map; uniqueness makes the two lifts inverse. □

Without base points, the relevant objects are conjugacy classes of subgroups. Ignoring that conjugacy is a common error in a program that relabels sheets during continuation. The isomorphism on a punctured disk is $n \mapsto n$ for an orientation-preserving homeomorphism and $n \mapsto -n$ for an orientation-reversing one. Either preserves the subgroup $d\mathbb{Z}$; it need not preserve each signed winding number.

2.5.2 Double covers and the general single-branch case

Theorem 2.19 (M17: The precise lift count in Lemma 3.11)**.** *Two connected degree-d branched covers of disks, each with a single branch point of local degree d, admit exactly d lifts of a base homeomorphism carrying the branch value to the branch value. The lift is fixed by one regular-fibre value. For $d = 2$ this is source Lemma 3.11. If the base map and covering maps are holomorphic, all lifts are holomorphic.*

Proof. Delete the branch point and value. Both maps become connected degree-d coverings of punctured disks. Their covering subgroups are $d\mathbb{Z} \subset \mathbb{Z}$. By M16, the base homeomorphism lifts after any of the d possible choices of the image of a regular base point; uniqueness makes these lifts distinct and exhaustive. In the local power model, a sequence tending to the puncture must lift to a sequence tending to the puncture, so adjoining the branch points extends the lift and its inverse continuously. In the holomorphic case, the lift is holomorphic away from the puncture. In a bounded local target coordinate it extends by removable singularities; the inverse does too. □

Correction to a shorthand in the source proof. The image subgroup consists of loops of winding *an integer multiple of two*, not only loops of winding exactly two. The set containing only winding two is not a subgroup. The source's intended equality of subgroups is preserved by this correction.

Worked problem: Closed loops and an observable choice

Let $w(t) = re^{2\pi imt}$ and $z(0) = r^{1/d} > 0$ under $z^d = w$. The unique continued branch is

$$z(t) = r^{1/d} e^{2\pi imt/d}.$$

It closes exactly when d divides m. The unordered fibre returns after every loop, while a selected sheet need not. The observable z^d is single valued and $\sum_{z^d=w} z = 0$ for $d > 1$, so some quantities can be computed without a sheet choice. A branch-sensitive output such as z cannot be replaced by its trace without changing the question.

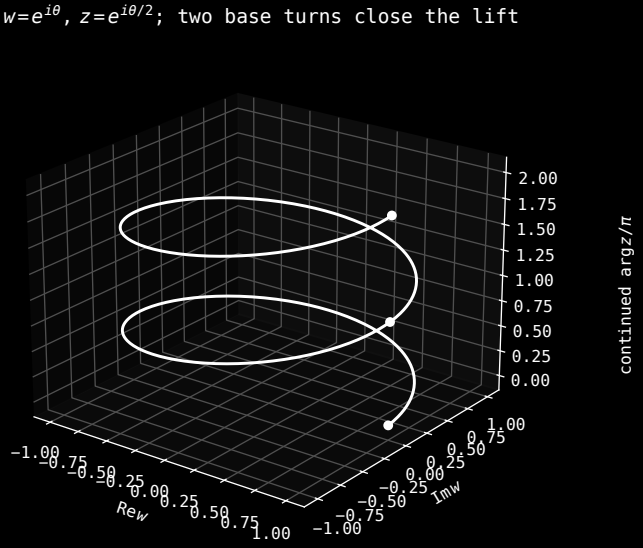


Figure 2.4: A continuous square-root lift over two turns of a parameter loop. Height is the real part of the chosen root, not an embedding of all four real coordinates of the algebraic graph. One base turn exchanges sheets; two close the lift.

2.5.3 Two or more branch values: solution of the analysis in Exercise 3.12

For several branch values, the number and order of local ramification indices no longer suffice. Delete all branch values and all their preimages. Choose a reference fibre and loops about the deleted values. The resulting representation

$$\rho : \pi_1(Y \setminus V, y_0) \to S_d$$

records how every loop permutes the sheets. A base homeomorphism lifts between two such covers if and only if the transported monodromy representations are conjugate by the permutation determined by the chosen fibre bijection. This is M16 expressed through permutation actions. It is a simultaneous conjugacy condition, not separate relabelling for each branch loop.

Theorem 2.20 (M18: Two branch points do not determine a cover)**.** *Even disk covers with two critical points and two critical values can have the same degree and source topology but fail to be equivalent over the identity of the base.*

Proof. Consider degree five. A transitive branch tuple

$$\sigma_1 = (1\,2\,3), \qquad \sigma_2 = (3\,4\,5)$$

has defects two and two. Its product is a five-cycle (with either consistent permutation convention, possibly inverted). The corresponding connected disk covering obtained by slit-and-sheet gluing has total defect four, hence source Euler characteristic $5 - 4 = 1$ and one boundary component. It has two critical points, of local degrees three and three.

A second tuple

$$\tau_1 = (1\,2\,3\,4), \qquad \tau_2 = (4\,5)$$

is also transitive, has a five-cycle product and defects three and one. It also gives a disk cover, with two critical points, now of local degrees four and two. A homeomorphism over the base must preserve local degree at the fibre of each branch value. Since the first branch permutations have different cycle types, no simultaneous conjugacy exists. The covers are not equivalent over the identity. □

The slit construction glues labelled sheets along cuts according to each permutation. The complex structure is pulled back away from branch points, and local power coordinates complete the branch points. Uniformization supplies disk realizations. The tuple calculation is the finite part of the construction; it does not independently prove uniformization.

If a particular exercise means two *critical values* instead of two critical points, the same warning applies. A permutation can have several nontrivial cycles above a single critical value. The certificate format consequently stores both cycle type and the actual permutations, together with the loop-order convention.

2.5.4 Real symmetry: Exercise 3.13 and a sharp negative control

The source writes $f(T \cap \mathbb{R})$ although f is defined on S. We use the well-typed condition

$$h\big(f(S \cap \mathbb{R})\big) = \widetilde{f}(\widetilde{S} \cap \mathbb{R})$$

inside the double-cover setting, together with real symmetry of the disks, branch points and maps.

Theorem 2.21 (M19: Both double-cover lifts preserve reflection under the side condition)**.** *With those assumptions, both lifts in M17 for $d = 2$ satisfy $H \circ \sigma = \sigma \circ H$, where $\sigma(z) = \bar{z}$ in the chosen symmetric coordinates.*

Proof. Near a real branch point, a symmetric holomorphic double cover has a real local coordinate form $\pm u^2$ on its real diameter. On the appropriate side of the branch value, a nearby real regular value has two real preimages. The side condition says that the corresponding target fibre also consists of two real points. Therefore either lift takes a selected real regular point to a real point. The maps H and $\sigma \circ H \circ \sigma$ are both lifts of the same base map, because all the input maps commute with reflection. They agree at that regular point. Uniqueness in M16 makes them equal everywhere. □

Worked problem: Why the side condition cannot be omitted

Take $f = \widetilde{f} : z \mapsto z^2$ on the unit disk, and $h(w) = -w$. All maps commute with conjugation, and the branch value remains zero. The lifts are $H(z) = iz$ and $H(z) = -iz$. Neither preserves the real line. Here $h([0, 1)) = (-1, 0]$ is not the real-image side $[0, 1)$ of the target cover. The extra condition excludes precisely this failure.

Functional-analytical addition: unitary sheet relabelling. Over a simply connected regular base region, a vector of fibre values belongs to $\mathbb{C}^d$ with normalized inner product $d^{-1} \sum u_j \bar{v}_j$. A sheet permutation is unitary. The mean projection is invariant under it, but a coordinate projection onto a selected sheet is not. Continuation therefore transports an orthonormal finite fibre space without requiring a global coordinate label. This is a useful finite-dimensional prototype of a local system; it is not a global trivialization of the branched surface.

2.5.5 Pullback algebras and an operator intertwining identity

The lift relation $\widetilde{f} \circ H = h \circ f$ gives a commuting diagram of local holomorphic rings by pullback. For compact surfaces it similarly gives compatible embeddings of meromorphic function fields. These maps preserve composition and local vanishing orders; the subgroup compatibility is the topological condition that makes the algebraic diagram globally realizable.

On bounded holomorphic function spaces, define $C_\phi g = g \circ \phi$. The same relation becomes

$$C_H C_{\widetilde{f}} = C_f C_h \quad \text{on } H^\infty(\widetilde{T}).$$

Every factor is a bounded operator of norm at most one for the specified supremum norms; the pullbacks by the onto conformal isomorphisms H and h are isometries. The identity follows by evaluation at every point. This is an operator consequence of a correctly typed lift, not a substitute for proving that the lift exists. It also shows exactly which data would be lost by composing symbols whose source and target surfaces do not match.

2.6 Galois coverings and controlled symmetrization

Mathematical reference. Exercise 3.14 and Proposition 3.15; pp. 109–110.

2.6.1 Holomorphic deck groups and quotients: Exercise 3.14

Theorem 2.22 (M20: Deck transformations are analytic, and analytic quotients have power charts)**.** *A deck transformation of a holomorphic branched cover is a holomorphic automorphism. Conversely an effective properly discontinuous holomorphic group action with finite stabilizers defines a Riemann-surface quotient with local branched-power projection. For a finite group, proper discontinuity is automatic.*

Proof. Away from branch points, write a deck transformation as an inverse branch of the covering composed with the covering; it is holomorphic. At a branch point it is a homeomorphism into a bounded coordinate neighbourhood and has a removable singularity. Apply the same argument to its inverse.

For the converse, consider the finite stabilizer at a point. Its derivative representation is faithful: a finite-order germ with derivative one is the identity, as follows either from its first nonlinear term under iteration or from analytic linearization by averaging. A finite subgroup of $\mathbb{C}^*$ is cyclic. If γ generates a stabilizer of order e with derivative ζ, the averaged coordinate

$$u(z) = \frac{1}{e} \sum_{k=0}^{e-1} \zeta^{-k} \gamma^k(z)$$

has nonzero derivative and satisfies $u(\gamma z) = \zeta u(z)$. The invariant coordinate is u^e. These local charts, together with ordinary quotient charts at free points, define the complex structure on the quotient and make the projection holomorphic. □

An arbitrary dense group of rotations is not an admissible action for this construction: its quotient need not be a surface. Likewise an antiholomorphic reflection produces a different quotient category with boundary. The word quotient does not erase the hypotheses required to obtain holomorphic charts.

2.6.2 The normal core: a full proof of Proposition 3.15

Theorem 2.23 (M21: Galois closure of degree at most factorial)**.** *Let $f : X \to Y$ be a connected finite-degree holomorphic branched covering of degree d. There is a connected Galois branched cover $g : Z \to Y$, ramified only over critical values of f, with $g = f \circ h$ and $\deg g \leq d!$.*

Proof. Let V be the critical-value set, $Y^* = Y \setminus V$ and $X^* = X \setminus f^{-1}(V)$. Properness and isolated critical points make V locally finite: infinitely many critical values in a compact target set would give accumulating critical points in a compact source inverse image, contradicting analyticity. Removing a discrete set does not disconnect a connected surface, so the unbranched covering remains connected.

Fix y_0 and label its fibre. Path lifting gives a transitive monodromy homomorphism $\rho : \pi_1(Y^*, y_0) \to S_d$. Let $N = \ker \rho$. It is a normal subgroup, and the connected covering corresponding to N is regular, with deck group $G = \rho(\pi_1(Y^*))$. Its degree is $|G| \leq |S_d| = d!$. The stabilizer subgroup of the selected sheet contains N, so the regular cover factors through X^*. Pull back the complex charts to make both maps holomorphic.

Near a deleted value, take a sufficiently small punctured disk containing no other deleted value. Every connected component of its inverse image in the finite covering is a finite covering of a punctured disk, hence has a power-map model. Adjoin the missing point in that model and extend the maps by the corresponding powers. Perform this operation in the locally finite family of such neighbourhoods. The extended surface is Hausdorff and second-countable, and the extended map is proper: a compact target set meets only finitely many branch charts, with finitely many sheets elsewhere. Deck transformations extend across the filled points by M20. No point outside V was modified, so no new branch value is created. □

The bound counts the order of the monodromy image, not the number of all permutations generated by a naive factorial enumeration. An implementation should compute the subgroup actually generated and retain the larger bound only as a worst case. The source proof refers to the final local map by f where the factor map is h; the typed construction above keeps the maps distinct.

Algebraic geometry. For compact curves, $\mathbb{C}(X)/\mathbb{C}(Y)$ is a finite separable extension. Its normal closure corresponds to the completed regular cover. The monodromy kernel and field-theoretic normal closure express the same symmetry from topological and algebraic sides. The extension is separable because the field has characteristic zero. The degree may be far below $d!$ if the original cover has special symmetry.

Algebraic topology. The normal core of a covering subgroup H is $\bigcap_\gamma \gamma H \gamma^{-1}$. Its quotient acts faithfully on the cosets of H, giving the same embedding into S_d. The full subgroup, not merely the abelianized winding data, is needed for noncommuting branch loops. Homology alone loses the ordering information in a nonabelian monodromy representation.

Functional analysis. On a regular fibre, averaging over the deck group gives a self-adjoint idempotent

$$\mathcal{A}v = \frac{1}{|G|} \sum_{g \in G} U_g v.$$

Its range is the invariant subspace. Idempotence follows because each group element occurs exactly $|G|$ times among the products; self-adjointness follows by replacing g with g^{-1}. A transitive action on d sheets makes the invariant vectors constant. This exact projection is not an inverse of any branch-sensitive coordinate map.

2.6.3 An explicit closure for a cubic: the abstract construction becomes rational

Theorem 2.24 (M22: A degree-six closure of $x^3 - 3x$)**.** *Set $x = t + t^{-1}$ and $w = t^3 + t^{-3}$ on the t-sphere. Then $x^3 - 3x = w$. The rational map $t \mapsto w$ has degree six and deck transformations generated by $t \mapsto \zeta t$ and $t \mapsto t^{-1}$, with $\zeta^3 = 1$. They form the dihedral group of order six, isomorphic to S_3.*

Proof. Expanding $(t+t^{-1})^3 - 3(t+t^{-1})$ cancels the mixed terms and gives $t^3 + t^{-3}$. The rational map is $(t^6+1)/t^3$, of degree six. Both displayed transformations preserve it. Their relations are $r^3 = s^2 = 1$ and $srs = r^{-1}$, giving six distinct transformations. A generic fibre has six points, so the group is the full deck group and the map is Galois. The degree-two map to x followed by the degree-three polynomial to w gives the factorization. □

The roots of the cubic over w are $\zeta^j t + \zeta^{-j} t^{-1}$, $j = 0, 1, 2$. The two finite critical values of $x^3 - 3x$ are ± 2; the third branch value on the compact sphere is infinity. In the regular six-sheet action, the finite local monodromies have three transpositions each, and the infinity monodromy has two three-cycles. The defects are $3, 3, 4$, totaling ten. Riemann–Hurwitz gives $2 = 12 - 10$, so the closure is genus zero, consistent with its explicit rational parameter t.

This worked closure joins all three mathematical perspectives: a rational parametrization and function-field normal closure, an S_3 monodromy representation, and a finite unitary averaging operator. Each answers a different question. Counting them as three unrelated examples would overstate the information content.

2.6.4 A compositional invariant of the closure

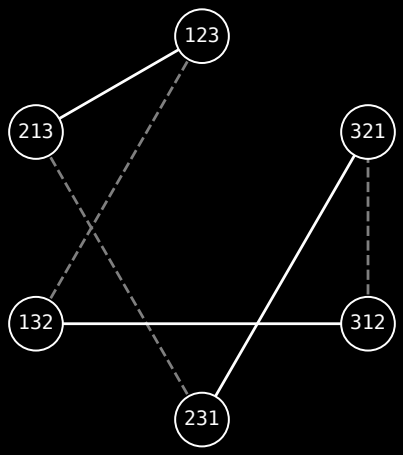
Theorem 2.25 (M23: Symmetric observables descend without choosing roots)**.** *For the cubic roots x_j of $x^3 - 3x - w = 0$, the power sums obey*

$$s_0 = 3, \quad s_1 = 0, \quad s_2 = 6, \qquad s_n = 3s_{n-2} + ws_{n-3} \quad (n \geq 3).$$

Thus every polynomial symmetric observable generated by these traces can be computed as a single-valued polynomial in w without numerical root tracking.

Proof. The initial values follow from Vieta's relations. Multiply $x_j^3 = 3x_j + w$ by x_j^{n-3} and sum over the three roots, including multiplicity. This gives the recurrence. At a critical value two roots coincide, but the polynomial identity remains valid, so the formula has no branch singularity there. □

The recurrence does not reconstruct a distinguished root. It changes representation for an invariant output while preserving that output exactly. The CSE volume compares this operation count with explicitly computing three roots, and the PG branches ask what conditioning survives after cancellation.



Solid: right multiplication by (12); dashed: by (23).

Figure 2.5: The six monodromy states of the cubic's regular closure, generated by two transpositions. Vertices are permutations; edge labels are generators. This is a finite group graph, not a spatial drawing of the branched surface.

2.6.5 A dependency correction: orbifold covering does not imply regularity

The note at the end of source 3.5 refers to Exercise 1.113 (p. 50), which states an equivalence with a *Galois* branched covering. The local subgroup-quotient model of an orbifold covering on p. 49 does not impose global normality. With that local definition, the Galois qualifier is too strong in the converse. This is a targeted dependency finding, not a review of all of source section 1.

Theorem 2.26 (M30: An explicit nonregular orbifold covering)**.** *The map $f(x) = x^3 - 3x$ restricts to an unbranched degree-three covering*

$$\mathbb{C} \setminus \{-2, -1, 1, 2\} \longrightarrow \mathbb{C} \setminus \{-2, 2\}.$$

Regarded as orbifolds with all local groups trivial, it is an orbifold covering but is not Galois.

Proof. The only finite critical points are ± 1, with critical values ∓ 2. The preimages of the two deleted values are exactly the four listed points, including multiplicities at ± 1. Restriction of a proper polynomial map to the full inverse image of the punctured target remains proper, and the derivative is nowhere zero there. The covering degree is three.

A holomorphic deck map extends across each removed finite point: from $f(H(x)) = f(x)$, its image is bounded near that puncture, because a polynomial tends to infinity when its argument tends to infinity. Removable singularities apply. Near infinity its image tends to infinity by the same identity. Applying the argument to its inverse shows that the extension is a sphere automorphism fixing infinity, hence $H(x) = ax + b$. Comparing the x^3, x^2 and x coefficients in $f(ax + b) = f(x)$ gives $a^3 = 1$, $b = 0$ and $a = 1$. The deck group is trivial, so it cannot act transitively on three-point fibres. The covering is not regular. □

The compact orbifold version assigns target weights two at ± 2 and three at infinity, and source weights two at the unramified preimages ± 2. All other source weights are one. The local rule $q_Y(f(x)) = e_x q_X(x)$ holds at every point, yet the underlying map is still nonregular. Its orbifold Euler characteristics are 1 and $1/3$, correctly related by the degree three. The correct general implication is "orbifold cover with compatible local weights," with regularity as an additional global group condition. Conversely, a Galois branched cover does supply an appropriate regular orbifold cover. This distinction prevents an incorrect normality assumption from propagating into the closure algorithm.

2.6.6 Functional analysis of a finite deck group

Theorem 2.27 (M29: Deck averaging is an orthogonal projection)**.** *Let a finite deck group Γ act on a measured surface X with a Γ-invariant finite measure μ. On $L^2(X, \mu)$ put $U_\gamma f = f \circ \gamma^{-1}$ and*

$$P_\Gamma = \frac{1}{|\Gamma|} \sum_{\gamma \in \Gamma} U_\gamma.$$

Then P_Γ is the orthogonal projection onto the invariant functions.

Proof. Measure invariance makes every U_γ unitary and gives $U_\gamma^* = U_{\gamma^{-1}}$. Inversion permutes the finite group, so $P_\Gamma^* = P_\Gamma$. Each group element occurs exactly $|\Gamma|$ times among the products $\gamma\eta$, giving $P_\Gamma^2 = P_\Gamma$. Its image is invariant because multiplication by any fixed group element permutes the summands; an invariant function is fixed by the average. A self-adjoint idempotent is an orthogonal projection. □

An invariant measure can be constructed by averaging the translates of any finite starting measure; ordinary area need not already be invariant under a conformal deck action. On a finite permutation fibre, the projection is onto functions constant on group orbits. Its trace is the number of orbits, also equal to the average number of fixed points of a group element. This supplies a finite check of information lost in passing to the quotient, with the measure and the representation explicitly fixed.

2.7 telescoping with explicit composition domains

Mathematical reference. The Telescoping Lemma and Corollary 3.16, p. 110.

2.7.1 Why a composition must carry its target restriction

The source writes $\phi_i : U_i \to \mathbb{C}$ with $\phi_i(U_i) \supset U_{i+1}$ and describes the domain of the full composition as mapped onto U_l. An unrestricted composition does not by itself impose the last target condition. For $l = 1$, $U_0 = U_1 = \mathbb{D}$ and $\phi_0(z) = 2z$, the original map is defined on all of $\mathbb{D}$ and maps it onto $2\mathbb{D}$, not $\mathbb{D}$. The intended domain is the pullback of the prescribed terminal disk. That restriction is made part of the definition below.

Theorem 2.28 (M24: Typed telescoping lemma)**.** *Let $U_0, \dots, U_l$ be open topological disks and let $\phi_i : U_i \to \mathbb{C}$ be univalent holomorphic maps with $U_{i+1} \subset \phi_i(U_i)$. Define*

$$F_0 = \mathrm{id}, \quad E_0 = U_0, \qquad E_{j+1} = \{z \in E_j : F_j(z) \in \phi_j^{-1}(U_{j+1})\}, \quad F_{j+1} = \phi_j \circ F_j|_{E_{j+1}}.$$

Then $F_j : E_j \to U_j$ is a biholomorphism for every j, and $E_{j+1} \subset E_j$. In particular $F_l : E_l \to U_l$ is the intended composition.

Proof. Induct on j. The claim is immediate at zero. Since ϕ_j is univalent and its image contains U_{j+1}, its inverse maps U_{j+1} biholomorphically onto a disk inside U_j. The inverse of F_j maps that disk biholomorphically onto E_{j+1}. Their inverses compose to give F_{j+1}. The inclusion follows directly from the definition. □

This definition checks every intermediate domain. A final residual $F_l(z) \in U_l$ is insufficient if an earlier stage crossed a pole or the chosen branch cut. In software, these are preconditions on a chain of partial maps.

2.7.2 Compact nesting and a repaired Corollary 3.16

Theorem 2.29 (M25: The correct pullback nest)**.** *If $U_{j+1} \Subset \phi_j(U_j)$ at every stage, then $E_{j+1} \Subset E_j$. With $D^k = E_{l-k}$ one has*

$$E_l = D^0 \Subset D^1 \Subset \cdots \Subset D^l = U_0, \qquad F_{l-k} : D^k \longrightarrow U_{l-k}$$

biholomorphic for $0 \le k \le l$.

Proof. The continuous inverse ϕ_j^{-1} sends the compact closure of U_{j+1}, contained in its domain, to a compact subset of U_j. The continuous inverse F_j^{-1} then sends it to a compact subset of E_j. This contains the closure of E_{j+1} and proves relative compactness. Reversing the index gives the displayed nest and its correctly matched targets. □

The source uses n in the full composition where l is the declared chain length, and its corollary pairs reverse-indexed domains with forward-indexed maps to a common target U_l. The identity above is an explicit repair. We do not silently copy those indices into the branches or animate an inconsistent nest.

Worked problem: An exact nested family

Take all $U_j = \mathbb{D}$ and all $\phi_j(z) = 2z$. Then $E_j = 2^{-j}\mathbb{D}$ and $F_j(z) = 2^j z$. The nested domains are real disks in the same initial coordinate chart. For $l = 3$, $D^0 = \frac{1}{8}\mathbb{D}$, $D^1 = \frac{1}{4}\mathbb{D}$, $D^2 = \frac{1}{2}\mathbb{D}$, $D^3 = \mathbb{D}$. The map from D^1 to $\mathbb{D}$ is F_2, not F_0. A single radius calculation catches the indexing problem.

2.7.3 Algebraic geometry of composition: ramification is not merely added

Theorem 2.30 (M26: Ramification divisor under composition)**.** *For finite holomorphic maps $g : X \to Y$ and $f : Y \to Z$,*

$$\mathcal{R}_{f\circ g} = \mathcal{R}_g + g^*\mathcal{R}_f.$$

For compact covers this gives $R_{f\circ g} = R_g + (\deg g)R_f$.

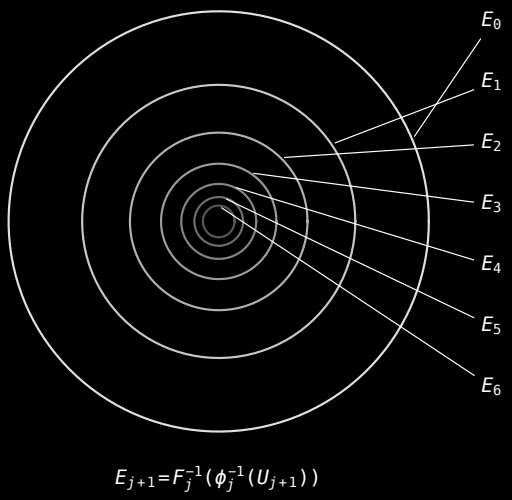


Figure 2.6: Seven target-aware pullback domains for $\phi_j(z) = (20/13)z$ and $U_j = \mathbb{D}$, with radii $(13/20)^j$. The nested sets share the initial coordinate plane. This is a second exact affine example; the preceding $2z$ example instead has radii 2^{-j}.

Proof. At x, let the two local degrees be e_g, e_f. Their product is the local degree of the composition. The coefficient identity

$$e_g e_f - 1 = (e_g - 1) + e_g(e_f - 1)$$

is exactly the local divisor identity; g^* multiplies a point-divisor coefficient by e_g. Sum and use the degree of pullback for the compact total. □

For a quadratic polynomial $q_c(z) = z^2 + c$, its nth iterate has degree 2^n. As a polynomial on $\mathbb{C}$, the total finite critical multiplicity is $2^n - 1$; on the sphere infinity contributes a further $2^n - 1$. The chain rule

$$(q_c^n)'(z) = 2^n \prod_{j=0}^{n-1} q_c^j(z)$$

locates the critical points with multiplicities, even when different precritical levels collide. At $c = 0$ every factor vanishes only at zero, yet the total multiplicity is still $2^n - 1$. Counting distinct preimages instead of multiplicities fails at precisely this important dynamical example.

On univalent restricted branches, every local degree is one, so the ramification divisor vanishes. The global iterate can be highly branched while its selected puzzle-piece inverse branch is univalent. This is the conceptual link between source section 3 and later pullback/puzzle arguments; no later puzzle theorem is claimed proved by this example.

2.7.4 Functional analysis of telescoping: an error theorem

Theorem 2.31 (M27: A compositional error budget with domain guards)**.** *Let maps T_i be L_i-Lipschitz on validated domains. Suppose exact and approximate states remain in those domains and*

$$x_{i+1} = T_i(x_i), \qquad \|\widehat{x}_{i+1} - T_i(\widehat{x}_i)\| \le \eta_i.$$

Then

$$e_l \le e_0 \prod_{i=0}^{l-1} L_i + \sum_{j=0}^{l-1} \eta_j \prod_{i=j+1}^{l-1} L_i, \qquad e_i = \|\widehat{x}_i - x_i\|.$$

Proof. The triangle inequality gives $e_{i+1} \le L_i e_i + \eta_i$. Substitute this inequality successively; induction yields the stated product-and-sum formula, with empty products equal to one. □

This estimate is meaningful only if the enclosures also prove domain membership. A small propagated bound computed after a wrong square-root branch has been selected certifies the wrong trajectory. Inverse derivatives can diverge near critical values, so the L_i must be established on whole enclosures, not evaluated at their centres alone.

If all $L_i \le q < 1$ and all $\eta_i \le \eta$, then $e_l \le q^l e_0 + \eta(1 - q^l)/(1 - q)$. This is a contraction statement for a declared metric. Holomorphic maps of the disk are nonexpansive in the hyperbolic metric, but can expand Euclidean distances. A norm cannot be changed halfway through the chain without conversion constants.

Topological addition. A chain of validated univalent branches transports the fundamental-group data of each allowed region without new local branching. If a target region begins to surround a critical value, that chain may no longer admit the chosen inverse branch. A domain guard is therefore not merely a floating-point safety check; it can be a topological prerequisite.

Chapter 3

Multiplicity, conditioning and compositional estimates

3.0.1 Q01 / source 3.1: normal forms separate topology from conditioning

Common problem. Given a local holomorphic map with a zero of order e, determine the number of nearby sheets and the response of a chosen sheet to a change in target. A conventional polynomial calculation expands and solves for each nearby root. The proposed decomposition first identifies $f(z) = z^e g(z)$ with $g(0) \neq 0$, then solves the coordinate problem $u^e = w$, where $u = zg(z)^{1/e}$. It retains a separate bound for the inverse coordinate change.

Theorem 3.1 (Q01: The local sensitivity exponent is not the local degree)**.** *For a chosen local branch of $u^e = w$ with $w \neq 0$,*

$$\left|\frac{du}{dw}\right| = \frac{1}{e}|w|^{1/e-1}.$$

If $z = \psi(u)$ and $|\psi'(u)| \leq L$ on the branch region, then $|dz/dw| \leq L|w|^{1/e-1}/e$. At the branch point the inverse map has a Hölder, not a uniform Lipschitz, scale.

Proof. Differentiate $u^e = w$ at $u \neq 0$ to obtain $eu^{e-1}u' = 1$. Taking moduli and using $|u| = |w|^{1/e}$ gives the first formula. The chain rule gives the second. At $w = 0$, the magnitude relation $|u| = |w|^{1/e}$ is exact. For $e > 1$, a finite constant times $|w|$ cannot bound $|w|^{1/e}$ in any punctured neighborhood of zero. □

Measured quantity versus theorem. The exponent is a proved statement; no timing study is needed to establish it. The graph shows the conditioning law for three degrees. A claim that normalizing the coordinate removes the physical sensitivity would be false: ψ changes a prefactor, not the ramification exponent. The algebraic-geometric quotient $\mathcal{O}_a/(f - f(a))$ preserves dimension e; the topological winding preserves e; the functional-analytic derivative norm can nevertheless diverge. Those are complementary observations, not three competing answers to one question.

When the baseline remains preferable. If the task is one explicit quadratic root far from the discriminant, the quadratic formula with a suitable stable evaluation is simpler than constructing a local coordinate. The decomposition becomes valuable for a family, a collision, an error specification, or the need to retain multiplicity after distinct values merge.

Exercise 3.1 (Q01a)**.** For $u^4 = w$, a residual report is $|u^4| \leq 10^{-12}$. What coordinate-radius certificate follows at the multiple zero? Is a radius 10^{-12} justified?

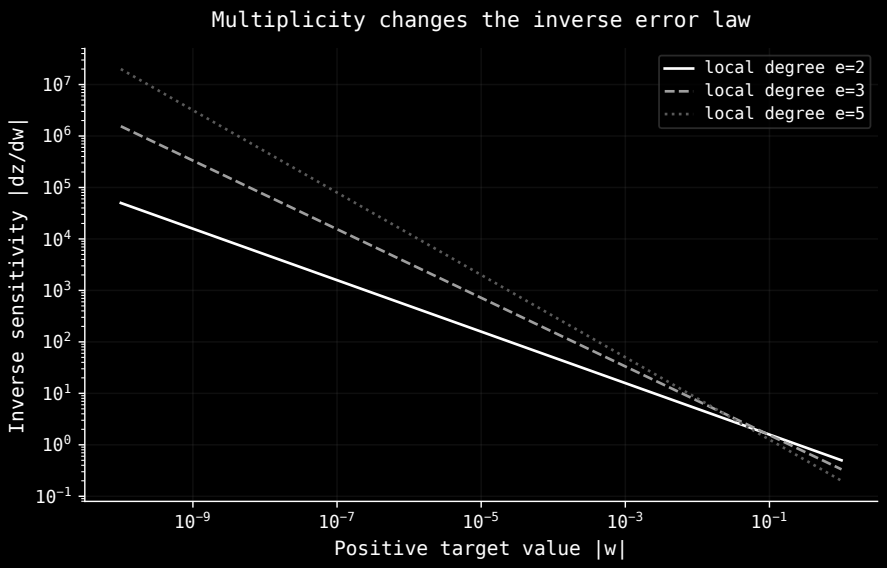


Figure 3.1: The sheet count is constant while inverse sensitivity is not. Each curve is derived from Q01; no fitted exponent is used.

Solution. The exact relation gives $|u| \le 10^{-3}$. A radius 10^{-12} is not justified. A small scalar residual becomes a much weaker coordinate guarantee when the derivative vanishes. If $z = \psi(u)$ with $\psi(0) = 0$ and $|\psi'| \le L$, the final radius is at most $L10^{-3}$, provided the entire connecting segment lies in the controlled coordinate patch.

3.0.2 Q02 / source 3.1: proper-map structure versus boundary sampling

Suppose the input is a factored disk map $B(z) = \lambda \prod_{j=1}^{d} b_{a_j}(z)$ with exact coefficients, $|\lambda| = 1$, and $|a_j| < 1$. Numerical sampling of $|B(e^{it})|$ is a natural diagnostic. The structural certificate is the identity

$$|1 - \bar{a}z|^2 - |z - a|^2 = (1 - |a|^2)(1 - |z|^2).$$

The certificate replaces an unbounded collection of boundary tests by a factor-wise polynomial identity and parameter inequalities. Its advantage is logical coverage of the whole disk, not an assertion that it evaluates B faster at a requested point. Evaluating a product at N points still uses $O(Nd)$ factor work unless a separate algorithm is introduced.

Degree-loss adversary. As $a \to 1$ along the real line, b_a tends to -1 on compact subsets of the disk. For every $a < 1$ the map has degree one; at the cancelled boundary parameter the limit is constant. A figure that animates $a = 1$ as one more ordinary proper map is wrong. In the functional-analytic representation the normalized defect vector

$$k_a(z) = \frac{\sqrt{1 - a^2}}{1 - az}$$

has H^2 norm one for every $a < 1$ but converges weakly to zero. A moving rank-one projection can converge strongly to zero while its operator norm remains one. The source's finite-degree classification therefore needs its parameter domain preserved in any numerical continuation.

Computable margin. The condition $|a| \le \rho < 1$ gives $|1 - \bar{a}z| \ge 1 - \rho$ on $|z| \le 1$. This is a denominator exclusion that a solver can use before division. The corresponding phase-derivative bounds are developed in the EE volume; they are not counted again as new core results here.

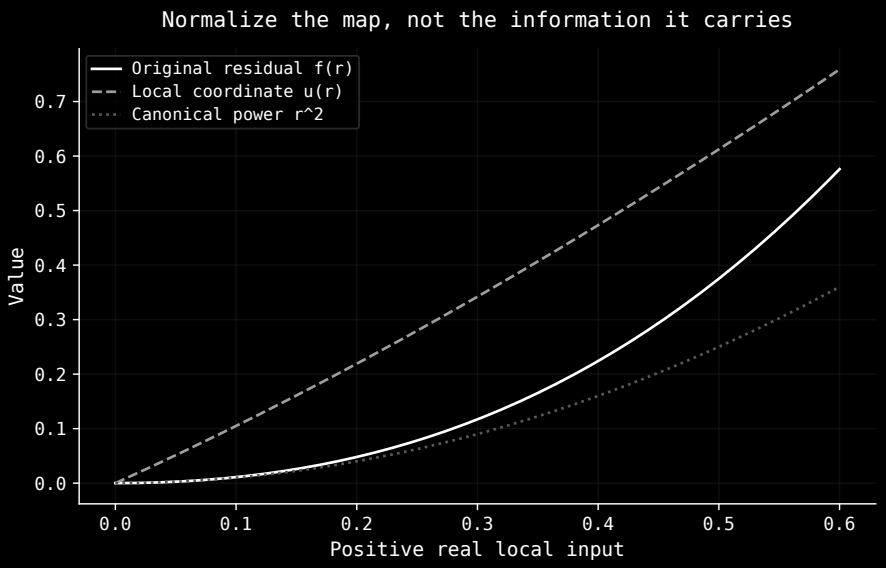


Figure 3.2: A real slice illustrates the exact local coordinate $u = z\sqrt{1+z}$ for $f = z^2(1+z)$. The complex inverse-function theorem, not the apparent monotonicity of this slice, supplies local invertibility.

3.0.3 Q03 / source 3.2: the branch ledger is a compressed Euler calculation

Same output. Both methods seek the Euler characteristic of a specified connected smooth compact covering surface. The cellular baseline builds a compatible triangulation and counts lifted vertices, edges and faces. The Riemann–Hurwitz representation stores d, $\chi(T)$ and local ramification contributions. If there are B listed branch points, summing their $e_a - 1$ contributions takes $B - 1$ additions after the data are known, independently of the number of unramified triangles. This is an arithmetic saving in a representation with pre-certified local models, not a cost estimate for discovering all branch points.

Theorem 3.2 (Q03: A completeness condition for the branch ledger)**.** *For a connected compact holomorphic cover with all local degrees verified, a ramification list is sufficient for Riemann–Hurwitz only when it contains every point with $e_a > 1$, including points in the infinity charts. Omitting one point of local degree e changes the computed Euler characteristic by $e - 1$.*

Proof. The formula M07 subtracts exactly $\sum_a (e_a - 1)$. Omitting one summand increases the calculated value by that summand. The chart at infinity is part of the compact surface and has no exemption. The local normal-form proof applies there after a coordinate change just as it does at a finite point. □

A second independent question. For a proposed passport, $g = 1 - d + R/2$ being a nonnegative integer is only a necessary arithmetic test. It does not enforce the product-one relation among branch permutations or transitivity. For example, $d = 4$ with two branch permutations both equal to (12) gives product one and $R = 2$, but the generated action fixes sheets 3 and 4. It does not describe the claimed connected four-sheet cover; the arithmetic would already give $g = -2$. Even when genus arithmetic passes, transitivity and local passports must still be checked. The CSE branch performs those discrete checks without confusing them with analytic realizability.

Exercise 3.2 (Q03a)**.** For $y^2 = P(x)$ with squarefree P of degree seven, account for the infinity chart and compute the genus. How many independent holomorphic differentials should a complete basis contain?

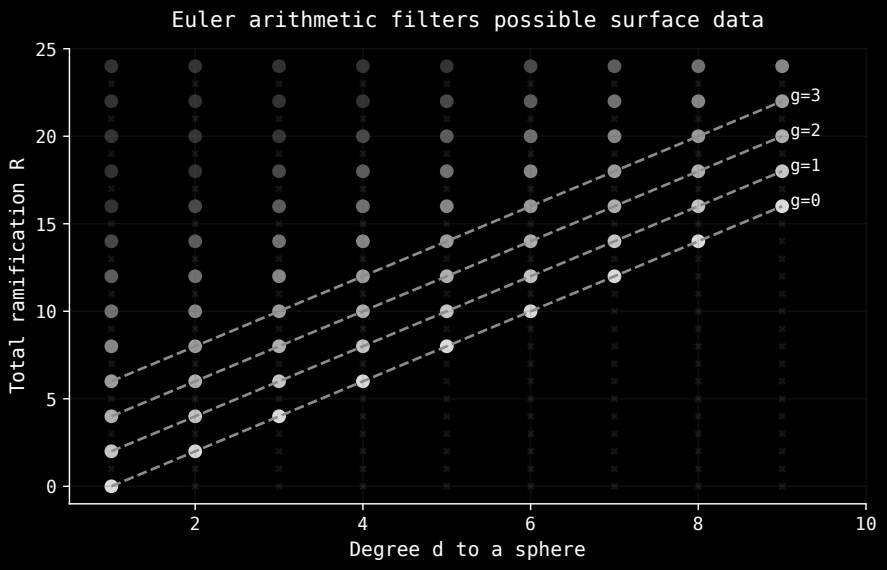


Figure 3.3: Necessary genus arithmetic for maps to a sphere. The grid is a filter on possible data, not a realizability classification.

Solution. Seven simple finite branches are present. With $x = t^{-2}$, the local square-root description gives one further simple branch at infinity. Thus $R = 8$, $d = 2$ and $2-2g = 4-8 = -4$, so $g = 3$. The differentials dx/y, $x\,dx/y$, $x^2\,dx/y$ are the three candidates; their infinity orders are $4, 2, 0$, respectively. Their independence and absence of other poles finish the usual basis argument with the dimension theorem from the graduate branch.

3.0.4 Q04 / source 3.2: topology survives before geometry does

A discriminant can certify that a squarefree polynomial family has no collision on a parameter region. It does not give a sharp local condition number automatically. For a monic degree-d polynomial with roots in $|z| \le R$ and $M = d(d-1)/2$ unordered pairs, the elementary bound is

$$s_{\min} \ge \frac{\sqrt{|\operatorname{Disc} p|}}{(2R)^{M-1}}.$$

Indeed, choose a closest pair and bound the other $M-1$ distances by $2R$ in the product formula for the discriminant. This globally valid estimate may be very pessimistic. The graph uses roots $0, \varepsilon, 1, 2$ to show its relationship to the actual minimum distance. A sharper local separation certificate can be preferable even though both guarantee the same topological type. A degeneration can preserve a simple arithmetic description while creating a large integral. The model

$$I(\varepsilon) = \int_0^1 \frac{dx}{\sqrt{x^2+\varepsilon^2}} = \operatorname{arsinh}(1/\varepsilon) = \log\frac{1+\sqrt{1+\varepsilon^2}}{\varepsilon}$$

satisfies $I - \log(2/\varepsilon) = O(\varepsilon^2)$. The equality follows by $x = \varepsilon \sinh t$. The asymptotic follows by expanding the smooth square root at zero. This is a transparent local model for a logarithmic neck integral; reducing a global period to this model requires its own coordinate and remainder estimates.

3.0.5 Q05 / source 3.3: turn degree stability into a usable error radius

The traditional argument-principle proof gives an integer after verifying nonvanishing on a contour. A plotted phase often hides the nonvanishing obligation. The proposed workflow exports a lower bound δ for the entire contour and only then allows a perturbation certificate.

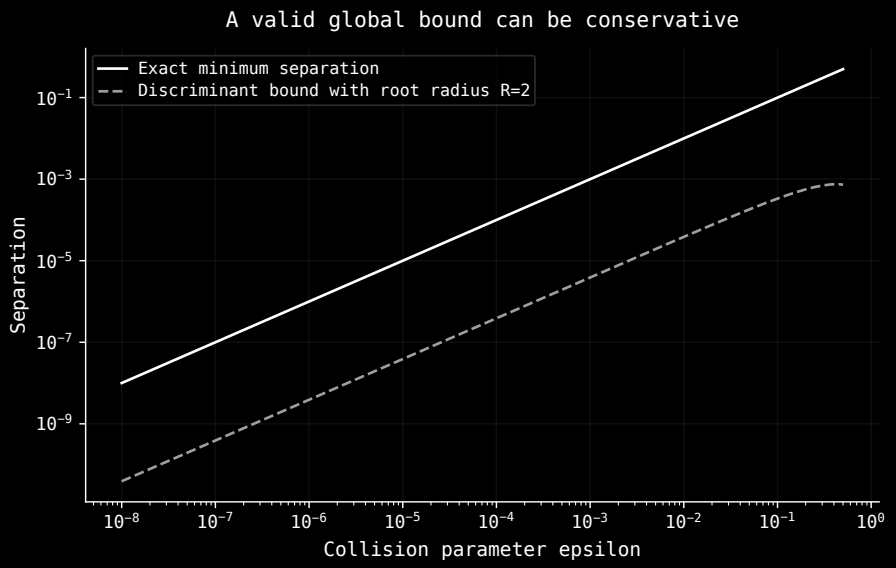


Figure 3.4: Validity and sharpness are distinct metrics. The conservative discriminant bound and exact minimum separation are shown on the same declared family.

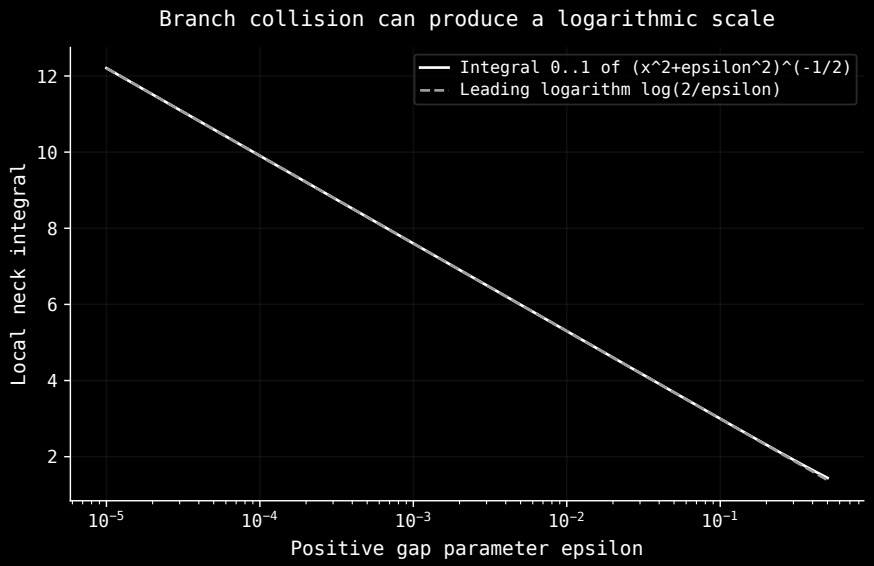


Figure 3.5: The exact local neck integral and its leading logarithm. Increasing geometric resolution without acknowledging the logarithmic scale can misinterpret slow convergence as numerical failure.

Theorem 3.3 (Q05: Certified degree under a boundary perturbation)**.** *Let f, g be continuous on a Jordan boundary Γ, with $\min_\Gamma |f-b| \geq \delta > 0$ and $\max_\Gamma |g-f| < \delta$. Then $\mathrm{wind}_b(g|_\Gamma) = \mathrm{wind}_b(f|_\Gamma)$. If both extend holomorphically inside, their b-fibre multiplicity counts agree.*

Proof. For $0 \leq t \leq 1$, set $h_t = f + t(g - f)$. The reverse triangle inequality gives $|h_t - b| \geq \delta - t\|g - f\|_\infty > 0$. This is a homotopy of loops in the punctured plane; homotopy invariance proves the winding equality. M13 and positivity of holomorphic local indices identify that winding with the fibre count. □

Information gained. The bound turns a qualitative topological result into a permissible implementation tolerance. It also states exactly when no such common tolerance can be asserted: a family of curves may have constant winding while its clearance tends to zero. A single count does not bound the robustness radius.

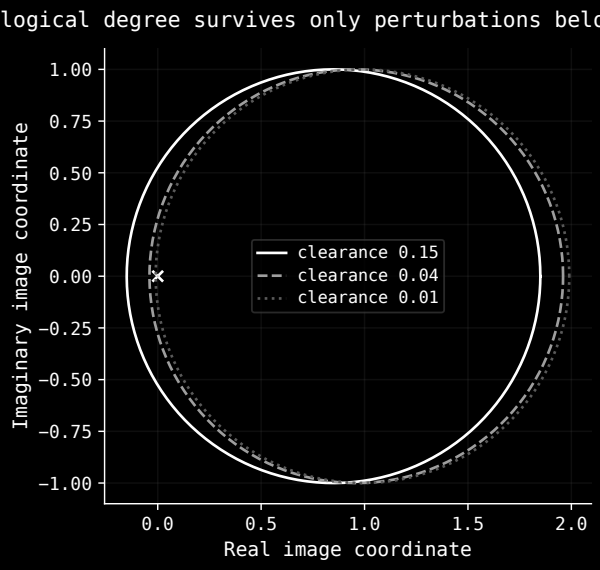


Figure 3.6: Identical winding and decreasing robustness. The exclusion radius, not the integer alone, controls whether perturbations preserve the conclusion.

Exercise 3.3 (Q05a)**.** A contour certificate proves $|p| \geq 3/200$. A coefficient perturbation gives $|q - p| \leq 1/100$ on that contour. What is the remaining clearance along the straight homotopy? What happens when the error bound is $3/200$ instead?

Solution. The clearance is at least $3/200 - 1/100 = 1/200 > 0$, so winding is preserved. Equality at $3/200$ gives only a nonnegative lower bound and does not exclude a zero. For $p \equiv 3/200$ and $q \equiv 0$, the endpoint loop actually hits zero. Strictness is mathematically necessary.

3.0.6 Q06 / source 3.3.1: retain the algebra when values collide

The supplied heading for degree is empty, so this is a corpus extension. In $A_\lambda = \mathbb{C}[x]/(x^2 - \lambda)$, the basis $(1, x)$ remains two-dimensional at $\lambda = 0$. Multiplication by x is

$$M_\lambda = \begin{pmatrix} 0 & \lambda \\ 1 & 0 \end{pmatrix}, \qquad M_\lambda^2 = \lambda I.$$

For $\lambda \neq 0$ a diagonalization exposes the two roots. At zero the same two-dimensional algebra becomes nonreduced, and diagonalization no longer exists. The invariant calculation $F(M) = c_0(\lambda)I + c_1(\lambda)M$ can remain well-defined. This is the algebraic-geometric reason to retain an algebra object instead of deduplicating a list of coordinates.

Functional interpretation. For an entire F, the scalar coefficient functions follow by separating the power series into even and odd parts. At a multiple eigenvalue, derivatives encode the nilpotent direction. The complexity advantage is conditional on a low-degree polynomial relation that is actually known. Computing such a relation for a large arbitrary matrix can cost more than a robust general matrix-function method; the experiments include that general baseline instead of a deliberately fragile alternative alone.

3.0.7 Q07 / source 3.4: continuation is a path contract

A principal square-root routine is not an incorrect algorithm. It solves a pointwise convention on a cut plane. A continued branch solves a different problem: given a path $w(t) \neq 0$ and

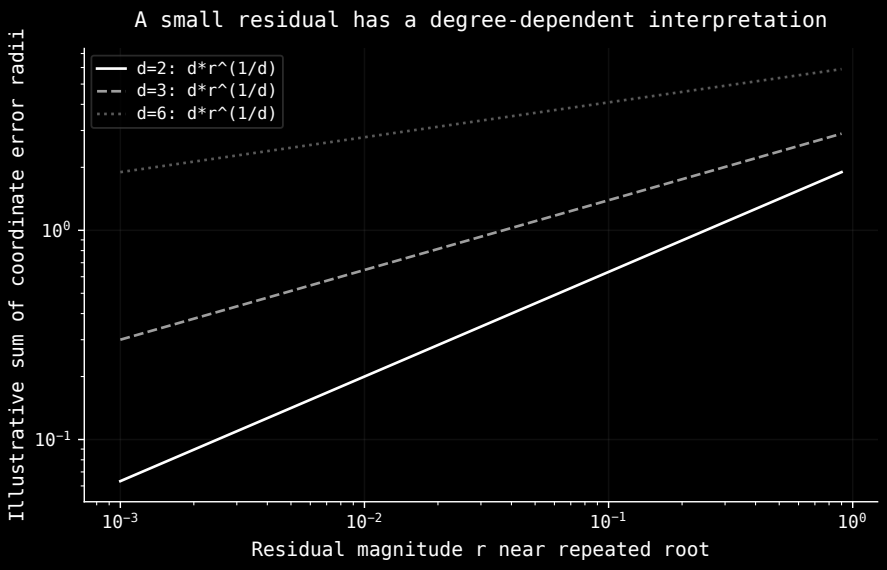


Figure 3.7: Residual exponents depend on multiplicity. A small residual cannot be interpreted by the same forward-error formula at simple and repeated fibres.

$z(0)^2 = w(0)$, find the continuous lift. The fair comparison gives both methods the continued-lift task and reports where pointwise evaluation ceases to meet it.

Theorem 3.4 (Q07: A sharp matching condition on a circular square-root path)**.** *For consecutive base angles differing by $0 \le \Delta\theta < \pi$, selecting the nearer of the two exact square roots to the previous lifted root reproduces the continuous lift of $w = re^{i\theta}$. At $\Delta\theta = \pi$ the two distances are equal, so nearest matching alone is ambiguous.*

Proof. Rotate and scale the previous root to $\sqrt{r}$. The continued candidate is $\sqrt{r}e^{i\Delta\theta/2}$ and the other is its negative. Their distances to $\sqrt{r}$ are $2\sqrt{r}\sin(\Delta\theta/4)$ and $2\sqrt{r}\cos(\Delta\theta/4)$. On the stated interval the first is strictly smaller; at the endpoint they are equal. Apply this argument successively. □

With computed roots, subtract the root-enclosure errors from the distance gap before deciding. The theorem about exact candidates is not itself a floating-point proof. Away from a circular path, root-tube bounds, path increments and separation margins replace this special formula. The monodromy remains a global obstruction: after one full base loop the selected square root changes sign, even if every local step was certified.

3.0.8 Q08 / source 3.5: the normal closure is sized by its actual group

The $d!$ theorem is a worst-case bound, not an instruction to enumerate every permutation. For a transitive permutation group G of a d-element fibre, the regular closure has degree $|G|$. Enumerating a subgroup generated by k explicit permutations can use a frontier and hash set. With permutations represented by arrays, each product costs $O(d)$ elementary entries, so the group-enumeration work is $O(kd|G|)$ up to dictionary and representation costs. The input generators, not merely their number, determine $|G|$.

Algebraic-topological check. A cover is represented by a monodromy homomorphism, not a bare group order. Changes of loops may conjugate generators or perform branch moves. The physical branch labels, basepoint and multiplication convention must be retained. Two groups of the same order do not imply identical covering data. On the other hand, computing an invariant

of one orbit may need only d states, not $|G|$ full orderings. That smaller computation is a genuine saving only if the requested observable factors through the orbit projection.

Operator check. Averaging over a finite deck group gives a projection only with the specified action and invariant measure. A uniform average over arbitrary numerical root labels does not automatically implement that operator. The mathematical quotient, topological quotient and conditional expectation agree only after these relations are proved.

3.0.9 Q09 / source 3.6: correct domains precede numerical efficiency

The telescoping correction is not a cosmetic reindexing. A formula can be algebraically meaningful at an input that fails an intermediate or terminal domain requirement. The correct partial program exports both its value and a valid-domain witness. Under the Lipschitz and local-error hypotheses of M27,

$$e_n \le \Big(\prod_{j=0}^{n-1} q_j\Big)e_0 + \sum_{k=0}^{n-1}\Big(\prod_{j=k+1}^{n-1} q_j\Big)\eta_k.$$

The baseline that uses one worst-case $q = \max q_j$ can overestimate every weight. The proposed bound keeps the ordered products and permits stage-dependent precision. The proof is an induction on $e_{j+1} \le q_j e_j + \eta_j$; the ordered positions cannot be discarded.

A negative control. Error averaging is invalid here: one large late error is not diluted by many small early errors. Nor can a final residual certify that every partial map was evaluated legally. The CSE parent gives explicit domain-rejection logic, while the mathematics research branch solves a precise integer allocation problem. The latter obtains an optimum under declared bit costs; it is not presented as a theorem that the whole written corpus is globally optimal.

3.0.10 Cross-perspective audit and a useful stopping rule

For every source unit, the ledger records a result, its assumptions, the requested perspective, an explanatory or computational use, a failure boundary, and where verification stops. A perspective is omitted from a particular *calculation* when it adds no information, while its honest role in the source unit remains documented. The quadratic formula, a three-line Routh check, or a direct local coordinate can remain the final recommended solution. Greater mathematical coverage does not require choosing the most expensive technique at every opportunity.

Chapter 4

Normal families, conformal mapping and distortion

4.1 Reference §4.1: Little Montel Theorem

4.1.1 The local estimate and its scale

Theorem 4.1 (M01: Cauchy derivative estimate)**.** *Let f be holomorphic on a neighbourhood of $\overline{D(z_0, R)}$, and suppose $|f| \le M$ on its boundary. For every integer $j \ge 0$,*

$$|f^{(j)}(z_0)| \le \frac{j!M}{R^j}.$$

If f is holomorphic and bounded by M on an open domain U, then

$$|f'(z)| \le \frac{M}{\operatorname{dist}(z, \partial U)}, \qquad z \in U,$$

with the usual interpretation for $U = \mathbb{C}$.

Proof. Cauchy's formula gives

$$f^{(j)}(z_0) = \frac{j!}{2\pi i} \int_{|\zeta - z_0| = R} \frac{f(\zeta)}{(\zeta - z_0)^{j+1}}\, d\zeta.$$

The contour has length $2\pi R$. Taking absolute values produces $j!M/R^j$. In an open domain one first uses any $R < \operatorname{dist}(z, \partial U)$ and then lets R increase to the distance. For an entire bounded function the same estimate holds for every R, so the derivative vanishes. □

The contour length is essential. The printed estimate with a squared distance in the denominator has the wrong scale for a first derivative. On $U = D_2$, take $f(z) = z$. The proposed right side at the origin would be $2/2^2 = 1/2$, while $|f'(0)| = 1$. The scale-correct estimate is attained in this example.

4.1.2 A source-ordered proof of normality

Theorem 4.2 (M02: locally bounded holomorphic families)**.** *Let $U \subset \mathbb{C}$ be a domain and $\mathcal{F} \subset \operatorname{Hol}(U)$. Suppose that for each compact $K \subset U$ there exists $M_K < \infty$ such that*

$$\sup_{f \in \mathcal{F}} \sup_{z \in K} |f(z)| \le M_K.$$

Every sequence in $\mathcal{F}$ has a subsequence converging locally uniformly to a holomorphic function on U.

Proof. Choose a compact exhaustion $K_m \subset \operatorname{int} K_{m+1}$ with $\bigcup_m K_m = U$. Such an exhaustion may be obtained from bounded distance-to-complement sets, with a subsequence or a finite union of disks used to secure the interior inclusion. Each K_m has a neighbourhood whose closure lies in U. The local boundedness hypothesis and M01 give a common derivative bound in this neighbourhood. On sufficiently small convex disks, the fundamental theorem of calculus along line segments gives a common Lipschitz bound. A finite cover supplies equicontinuity on K_m.

The restrictions are equibounded and equicontinuous on K_1; Arzelà–Ascoli yields a uniformly convergent subsequence there. Repeat on $K_2, K_3, \ldots$ and take the diagonal subsequence. The limit is continuous and the local limits agree on intersections. For a closed disk compactly contained in U, uniform convergence on its boundary allows passage to the limit in Cauchy's formula. The resulting formula makes the limit holomorphic on its interior. Such disks cover U. □

This proves source Theorem 4.1 and its locally bounded version. The proof uses second countability through the countable exhaustion. An individual bound depending on f is insufficient; the family $f_n(z) = nz$ is a counterexample on every neighbourhood of zero.

4.1.3 Finite jets and the compact restriction mechanism

For $f(z) = \sum_{j\geq 0} a_j z^j$ on D_R, its order-N jet is the residue class

$$J_N f = \sum_{j=0}^{N} a_j z^j \quad \text{in} \quad \mathbb{C}[z]/(z^{N+1}).$$

The quotient is a finite local algebra. Its maximal ideal is generated by z, and its nilpotence records truncation. The ring structure preserves products modulo order $N+1$. Analytic convergence requires an additional norm estimate.

Proposition 4.3 (M03: quantitative finite-jet approximation)**.** *If $|f| \leq M$ on D_R and $0 < r < R$, then*

$$\|f - J_N f\|_{\overline{D_r}} \leq M \frac{(r/R)^{N+1}}{1 - r/R}.$$

Consequently restriction from the bounded unit ball of $H^\infty(D_R)$ to $C(\overline{D_r})$ has finite-dimensional approximations with geometric uniform error.

Proof. For any $R' < R$, Cauchy's estimate gives $|a_j| \leq M(R')^{-j}$. Let $R' \uparrow R$ and sum the geometric tail on $|z| \leq r$. The coefficient vectors of $J_N f$ lie in a bounded subset of a finite-dimensional space and hence have compact closure. Combining finite coefficient nets with the uniform remainder proves total boundedness of the restricted family. □

The algebraic coordinates alone do not prove M02: a compatible infinite jet can have radius of convergence zero. The analytic estimate is the additional structure required for compactness.

4.1.4 Topological and functional formulation

The space $\mathrm{Hol}(U)$ is a closed subspace of $C(U)$ for compact-open convergence. With seminorms $p_m(f) = \sup_{K_m} |f|$, it is a Fréchet space. Completeness follows by uniform convergence on each K_m and the Cauchy-formula argument. Local boundedness implies relative compactness through M02. Thus holomorphic function spaces possess a compactness property that is absent from a general infinite-dimensional Banach space.

For comparison, the unit vectors e_n in ℓ^2 satisfy $\|e_n - e_m\|_2 = \sqrt{2}$ for $m \neq n$. They have no norm-convergent subsequence. In a holomorphic family, restriction to a smaller domain suppresses high-order coefficients. This suppression will yield an exact approximation-width result in M35.

Exercise 4.1 (The derivative scale)**.** A holomorphic function is bounded by 3 on $D(2, 5)$. Give certified bounds for its first and third derivatives at 2, and determine the first-derivative bound after the variable is rescaled by a factor of 10.

Solution. The bounds are $3/5$ and $3!\,3/5^3 = 18/125$. For $g(w) = f(2 + 10w)$, the domain is $D_{1/2}$ and $|g'(0)| \leq 3/(1/2) = 6$. The chain rule gives the same result, $|g'(0)| = 10|f'(2)| \leq 6$. □

Exercise 4.2 (A certified jet depth)**.** For $M = 2$, $r/R = 1/2$, find the smallest $N \geq 0$ for which M03 guarantees error at most 2^{-12}.

Solution. The bound is 2^{1-N}. Therefore $N \geq 13$, and $N = 13$ is the least certified degree for this bound. The answer concerns the class-wide guarantee; individual functions can need fewer terms. □

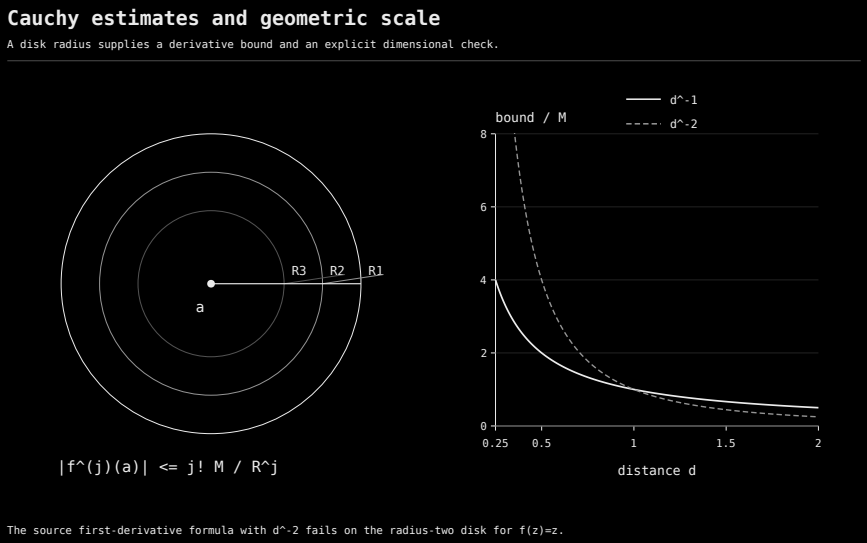


Figure 4.1: The first-derivative estimate changes as M/R. The marked linear example on D_2 exposes the inconsistent squared-distance estimate.

4.2 Reference §4.2: the Riemann Mapping Theorem

4.2.1 The topological obstruction and bounded reduction

A conformal equivalence preserves the fundamental group. The punctured plane and a genuine annulus cannot be conformally equivalent to $\mathbb{D}$: their fundamental groups are isomorphic to $\mathbb{Z}$. The plane has trivial fundamental group, but an equivalence $\mathbb{C} \to \mathbb{D}$ would be a bounded nonconstant entire function. These examples explain both hypotheses in the source statement.

Lemma 4.4 (M04: bounded conformal reduction)**.** *A simply connected domain $\Omega \subset \widehat{\mathbb{C}}$ whose complement contains at least two points admits a conformal embedding into a bounded planar domain.*

Proof. Choose two distinct complementary points and apply a Möbius transformation sending them to 0 and ∞. The transformed domain lies in $\mathbb{C}^*$. Simple connectivity supplies a holomorphic

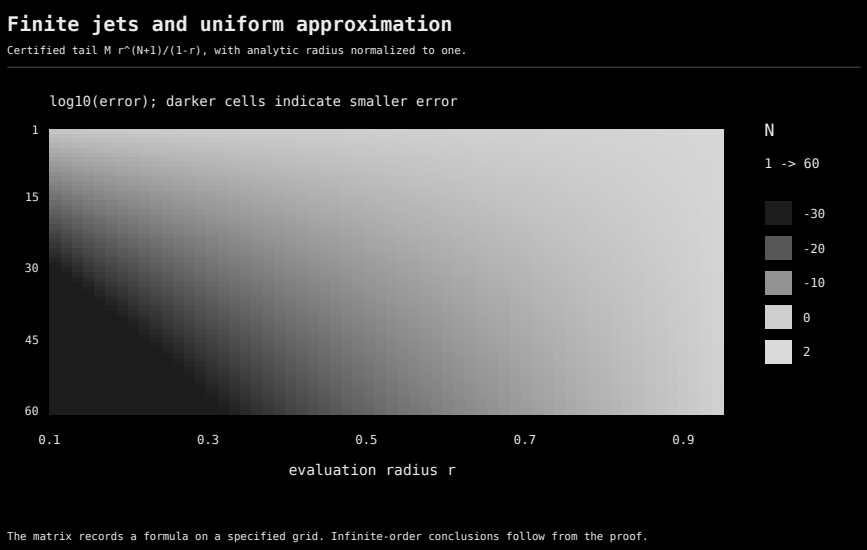


Figure 4.2: Geometric tail envelopes for bounded analytic families. Each curve specifies the ratio r/R and the truncation order. The curves are evaluations of M03.

logarithm, hence a single-valued square root Q. This square root is injective because $Q(z)^2 = z$. Its image V is disjoint from $-V$: a point common to them would give two opposite values of the same chosen square root at one nonzero point.

Choose $b \in -V$ and a disk $D(b, \delta) \subset -V$. The map $w \mapsto 1/(w - b)$ is injective on V and has modulus at most $1/\delta$ there. Composition provides the required bounded embedding. An affine transformation then embeds its image into $\mathbb{D}$ and can send any specified interior point to zero. □

In algebraic terms, Q selects one sheet of the correspondence $q^2 = z$. The fundamental-group hypothesis trivializes the monodromy of that correspondence on the domain. Neither the algebraic equation nor a pointwise principal square root supplies the global sheet by itself.

4.2.2 Compactness and attainment of the extremum

Lemma 4.5 (M05: limits of univalent functions)**.** *Let f_n be univalent on a domain Ω and converge locally uniformly, in the Euclidean metric, to f. Then f is univalent or constant.*

Proof. The limit is holomorphic by M02's closure argument. Suppose it is nonconstant and $f(z_1) = f(z_2)$ at distinct points. Choose disjoint closed disks centred at z_1, z_2, compactly contained in Ω, with $f - f(z_2)$ nonzero on the first boundary. The functions $f_n(z) - f_n(z_2)$ converge uniformly on that boundary to $f(z) - f(z_2)$. Rouché's theorem supplies a zero in the first disk for all sufficiently large n. The point z_2 is a second zero outside that disk. This contradicts injectivity of f_n. □

Theorem 4.6 (M06: normalized Riemann mapping)**.** *Let $\Omega \subsetneq \mathbb{C}$ be simply connected and let $a \in \Omega$. There is a conformal equivalence $f : \Omega \to \mathbb{D}$ such that $f(a) = 0$ and $f'(a) > 0$. These two normalization conditions determine f uniquely.*

Proof. By M04 the family $\mathcal{E}$ of conformal embeddings $\Omega \to \mathbb{D}$ taking a to zero is nonempty. Put $s = \sup_{f \in \mathcal{E}} |f'(a)|$. Cauchy's estimate on a fixed disk about a shows $s < \infty$. An embedding has nonzero derivative, so $s > 0$.

Choose $f_n \in \mathcal{E}$ with $|f_n'(a)| \to s$. M02 gives a locally uniformly convergent subsequence. Cauchy's formula also gives derivative convergence. Thus $f(a) = 0$ and $|f'(a)| = s > 0$. M05

implies injectivity. The limit takes values in the closed unit disk, and the maximum principle, together with nonconstancy, places its image in the open disk.

It remains to show surjectivity. Suppose $b \in \mathbb{D} \setminus f(\Omega)$. Then $b \neq 0$. Define

$$m_b(w) = \frac{w - b}{1 - \bar{b}w}, \qquad q^2 = -b,$$

and choose on Ω the square root Q of $m_b \circ f$ satisfying $Q(a) = q$. The function under the square root is zero-free; simple connectivity supplies the branch. Moreover $|Q| < 1$ and Q is injective. Therefore $F = m_q \circ Q$ lies in $\mathcal{E}$. Its derivative satisfies

$$\frac{|F'(a)|}{|f'(a)|} = \frac{1 - |b|^2}{2\sqrt{|b|}(1 - |b|)} = \frac{1 + |b|}{2\sqrt{|b|}} > 1,$$

since $(1 - \sqrt{|b|})^2 > 0$. This contradicts extremality. Hence f is onto. Multiplication by a unimodular constant makes its derivative positive.

For uniqueness, two normalized equivalences differ by a disk automorphism fixing zero. Schwarz's lemma applied to that automorphism and its inverse makes it a rotation. The positive derivative condition sets the rotation equal to the identity. □

The sphere version in the source follows by the initial Möbius reduction. The inverse map $\phi : \mathbb{D} \to \Omega$ is determined by $\phi(0) = a$ and the argument of $\phi'(0)$. The number

$$\operatorname{crad}(\Omega, a) = |\phi'(0)| = |f'(a)|^{-1}$$

is independent of the remaining rotation.

4.2.3 Source Exercise 4.2: the degree-two map

Proposition 4.7 (M07: an explicit covering with prescribed critical value)**.** *Let $b \in \mathbb{D} \setminus \{0\}$, choose $q \in \mathbb{D}$ with $q^2 = -b$, and set*

$$B_b(w) = m_b^{-1}\Big(\big(m_q^{-1}(w)\big)^2\Big).$$

Then $B_b : \mathbb{D} \to \mathbb{D}$ is a proper map of degree two, $B_b(0) = 0$, its unique critical point is $-q$, and its critical value is b. Also

$$|B_b'(0)| = \frac{2\sqrt{|b|}}{1 + |b|} < 1.$$

Proof. Disk automorphisms preserve properness and degree, and $w \mapsto w^2$ has degree two. Since $m_q^{-1}(0) = q$ and $q^2 = m_b(0)$, the value at zero is zero. The critical point is characterized by $m_q^{-1}(w) = 0$, hence $w = m_q(0) = -q$. Its image is $m_b^{-1}(0) = b$. The derivative formula follows from the chain rule, or by taking the reciprocal of the gain in M06. □

For a prescribed critical *point* c, a different formula is appropriate:

$$B_c^{\mathrm{cp}}(w) = m_{c^2}\big(m_c(w)^2\big).$$

It fixes zero, has critical point c, and has critical value $-c^2$. The source extremal argument needs the omitted point to be the critical value. These two specifications must be distinguished in an implementation of the argument.

4.2.4 Green rays and equipotentials

Proposition 4.8 (M08: conformal Green coordinates)**.** *Let $\phi : \mathbb{D} \to \Omega$ be a conformal equivalence with $\phi(0) = a$. On $\Omega \setminus \{a\}$ define*

$$G_a(z) = -\log|\phi^{-1}(z)|.$$

Its level curves and the images of radial lines are orthogonal analytic foliations. Near a,

$$G_a(z) = -\log|z-a| + \log \operatorname{crad}(\Omega, a) + o(1).$$

Proof. On a local logarithm chart for ϕ^{-1}, the function $-\log\phi^{-1}$ is holomorphic. Its real and imaginary parts have orthogonal gradients by the Cauchy–Riemann equations. The derivative does not vanish away from a. The Taylor expansion $\phi^{-1}(z) = (z-a)/\phi'(0) + O((z-a)^2)$ gives the displayed singularity. Rotating the disk coordinate changes the angular origin and leaves the two unlabelled foliations unchanged. □

The formula describes interior coordinates. Boundary continuity on an arbitrary planar boundary requires additional regularity or a prime-end formulation. This volume makes no unproved boundary-extension assertion.

4.2.5 Quantitative comparison and limitations

The existence proof is an extremal compactness argument. It does not specify a terminating algorithm for an arbitrary domain representation. A certified numerical map needs a domain description, normalization, boundary control, an injectivity certificate and a computable approximation error. For an explicitly given map $\phi(z) = z + az^2$, $|a| < 1/2$, these obligations can be verified directly. For an arbitrary polygon or a rough boundary they constitute separate numerical problems.

The finite-jet viewpoint describes local data in $\mathbb{C}[z]/(z^{N+1})$. The global domain is encoded by its embedding and boundary conditions. Two functions with the same finite jet can have different images. Therefore a theorem about a finite coefficient table cannot establish global equivalence without additional geometric hypotheses.

Exercise 4.3 (Extremal improvement at a specified omitted value)**.** Take $|b| = 1/4$. Compute the derivative-improvement factor in M06 and the derivative modulus of B_b at zero.

Solution. The improvement is $(1 + 1/4)/(2 \cdot 1/2) = 5/4$. The covering derivative modulus is its reciprocal, $4/5$. Thus an extremizer omitting this value would be improved by exactly 25% in derivative modulus, contradicting maximality of that functional. □

Exercise 4.4 (Conformal radius of a disk)**.** For $\Omega = D(c, R)$ and $a \in \Omega$, determine $\operatorname{crad}(\Omega, a)$ and compare it with $\operatorname{dist}(a, \partial\Omega)$.

Solution. Set $u = (a-c)/R$ and $\phi(z) = c + R(z+u)/(1+\overline{u}z)$. Then $\operatorname{crad}(\Omega, a) = R(1-|u|^2)$ and $d = R(1-|u|)$. Their ratio is $1 + |u| \in [1, 2)$. The class-wide factor four in the quarter theorem is not attained by round disks. □

4.3 Reference §4.3: spherical normal families

4.3.1 Source Exercise 4.3: a complete metric

Theorem 4.9 (M09: complete compact-open metric)**.** *Let U be a second-countable Riemann surface, and choose a compact exhaustion $K_m \subset \operatorname{int} K_{m+1}$, $m \geq 1$. On the space $\mathcal{M}(U)$ of*

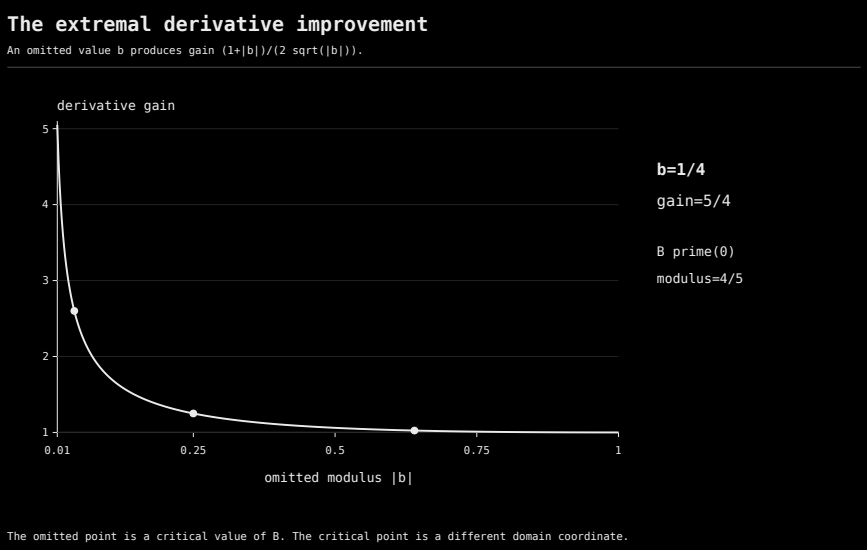


Figure 4.3: The extremal improvement factor $(1+t)/(2\sqrt{t})$ for an omitted value of modulus t. It is strictly greater than one for $0 < t < 1$ and approaches one as $t \uparrow 1$.

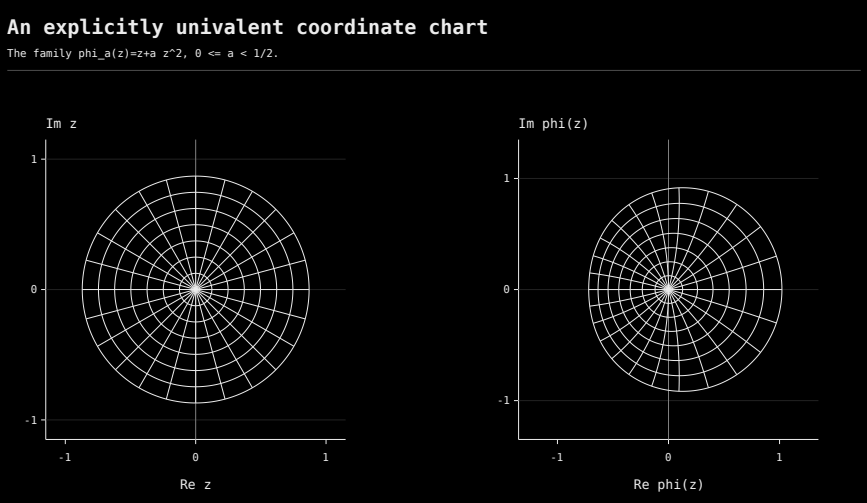


Figure 4.4: Conformal coordinate curves under $\phi(z) = z + z^2/5$. Injectivity follows from the factored difference and $|z + w| < 2$. The image is a worked explicit domain for the computational and electrical branches.

holomorphic maps $U \to \mathbb{P}^1(\mathbb{C})$, including the constant map ∞, set

$$d(f,g) = \sum_{m=1}^{\infty} 2^{-m} \sup_{z \in K_m} \chi(f(z), g(z)).$$

Then d is a complete metric. Its convergent sequences are exactly the locally uniformly spherically convergent sequences.

Proof. Each summand lies between zero and 2^{-m}, so the series converges. The triangle inequality follows from the triangle inequality for the chordal metric. If $d(f,g) = 0$, each compact supremum is zero; exhaustion then gives $f = g$.

If $d(f_j, f) \to 0$, the mth summand gives $\sup_{K_m} \chi(f_j, f) \le 2^m d(f_j, f) \to 0$. Conversely, choose M so that $\sum_{m>M} 2^{-m} < \varepsilon/2$, then use uniform convergence on the finite collection $K_1, \dots, K_M$ to control the remaining terms.

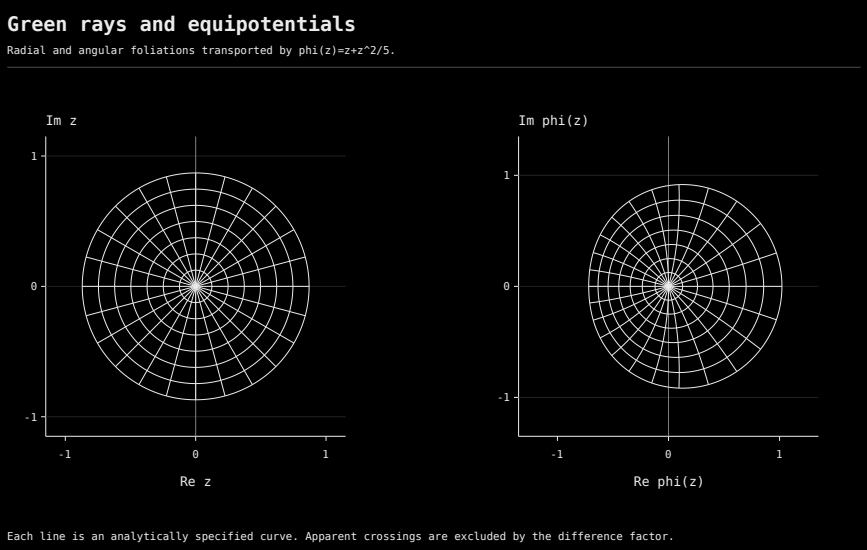


Figure 4.5: Green rays and equipotentials on the same explicit domain. Equal potential increments correspond to geometric radial ratios in the disk coordinate.

For completeness, a d-Cauchy sequence is uniformly Cauchy on every K_m. The target sphere is complete, so the limits exist and agree on overlaps. They define a continuous map $f : U \to \mathbb{P}^1(\mathbb{C})$. Fix $z_0 \in U$ and a target chart containing $f(z_0)$. Continuity provides a closed coordinate disk about z_0 whose image lies in a compact subset of this chart. Uniform spherical convergence places all sufficiently late f_j in the same chart. On that compact subset the Euclidean and chordal metrics are uniformly equivalent. The local coordinate functions therefore converge uniformly, and their limit is holomorphic by Cauchy's formula. These disks cover U, so $f \in \mathcal{M}(U)$. □

The two affine charts w and $1/w$ on $\mathbb{P}^1(\mathbb{C})$ are the algebraic-geometric structure underlying this proof. The transition function is rational. Compact-open completeness is a statement about maps into the compact projective target; it does not confer a linear-space structure on $\mathcal{M}(U)$.

4.3.2 Source Exercise 4.4: a limit taking the value infinity

Proposition 4.10 (M10: holomorphic escape alternative)**.** *Suppose U is connected, $f_n : U \to \mathbb{C}$ are holomorphic, and $f_n \to f$ locally uniformly in the spherical metric. If $f(z_0) = \infty$ at any point, then $f \equiv \infty$.*

Proof. On a sufficiently small disk about z_0, the limit image lies in a chart about infinity. For all large n, f_n has no zero there and $g_n = 1/f_n$ is holomorphic and zero-free. These functions converge uniformly in that chart to $g = 1/f$, with $g(z_0) = 0$. If g were nonconstant, choose a small circle around an isolated zero and use Rouché's theorem to force a zero of g_n, a contradiction. Hence g vanishes on a neighbourhood. The identity theorem for the meromorphic map gives $f \equiv \infty$ on the connected domain. □

For the constant sequence of functions $f_n(z) = n$, the chordal distance to infinity is $(1+n^2)^{-1/2}$, uniformly in z. Its Euclidean values diverge. A finite sample of those values can be large while the spherical approximation error is small and explicitly quantified.

4.3.3 Source Exercise 4.5: local normality

Proposition 4.11 (M11: local-to-global normality)**.** *A family $\mathcal{F} \subset \mathcal{M}(U)$ is normal if every point has a neighbourhood on which its restrictions form a normal family.*

Proof. Second countability supplies a countable subcover of such neighbourhoods. Given any sequence, successively extract a subsequence convergent on the first neighbourhood, then the second, and so on. The diagonal subsequence converges locally uniformly on every member of the subcover. Limits agree on overlaps because the target is Hausdorff. A compact subset of U has a finite subcover, so the convergence is uniform there. Completeness in M09 identifies the resulting limit as meromorphic. □

The argument is sequential, and all extractions are made from the preceding subsequence. A separate limit chosen independently in every chart need not define a global function.

4.3.4 Source Exercise 4.6: Marty's criterion

For a meromorphic function on a planar domain define its spherical derivative by

$$f^{\#}(z) = \frac{|f'(z)|}{1 + |f(z)|^2}$$

in a finite chart. At a pole the equivalent expression for $1/f$ extends the definition continuously.

Theorem 4.12 (M12: spherical derivative criterion)**.** *A family of meromorphic functions on $U \subset \mathbb{C}$ is normal if and only if its spherical derivatives are uniformly bounded on each compact subset of U.*

Proof. The displayed expression is the norm of the differential from the Euclidean source metric to ds_{sph}. Inverting the target coordinate leaves it unchanged:

$$\frac{|(1/f)'|}{1 + |1/f|^2} = \frac{|f'|/|f|^2}{(1 + |f|^2)/|f|^2} = f^{\#}.$$

A uniform bound on a slightly larger compact neighbourhood bounds spherical image lengths of short source line segments. It gives local equicontinuity. The target sphere is compact, so Arzelà–Ascoli and a diagonal exhaustion yield spherical compact-open convergence. M09 supplies holomorphicity of the limit.

Conversely, suppose the derivative bounds fail. There exist $f_n \in \mathcal{F}$ and z_n in one compact set with $f_n^{\#}(z_n) \to \infty$. Pass to subsequences with $z_n \to z$ and $f_n \to f$ spherically near z. Choose a spherical isometry placing $f(z)$ at zero. On a small disk about z, all sufficiently late transformed functions remain in one bounded affine chart and converge Euclidean-uniformly. Cauchy's estimate bounds their derivatives on a smaller disk. Spherical isometries preserve $f^{\#}$, contradicting divergence at z_n. □

Two examples distinguish the hypotheses. The constant family $f_n = n$ has spherical derivative zero. The family $g_n(z) = nz$ has $g_n^{\#}(0) = n$ and is nonnormal at zero. The derivative concentration has width of order $1/n$; uniform coarse sampling can miss its peak if the sampling nodes avoid the origin.

4.3.5 Source Exercise 4.7: the subsequence criterion

Proposition 4.13 (M13: bounded or escaping subfamilies)**.** *A sequence of holomorphic functions on a connected domain is normal precisely when every subsequence admits a further subsequence that is locally bounded in the Euclidean metric or converges locally uniformly to infinity.*

Proof. Normality supplies a spherically convergent further subsequence. M10 makes its limit either finite and holomorphic everywhere or identically infinite. In the finite case, convergence takes place in bounded charts on compact sets and gives local Euclidean boundedness. In the infinite case the second alternative holds. Conversely, apply M02 to each locally bounded further subsequence. The escaping further subsequences already converge in the compact target metric. This is sequential relative compactness in M09's metric. □

4.3.6 Source Theorem 4.8: three fixed omitted values

Theorem 4.14 (M14: Big Montel Theorem)**.** *Let $\mathcal{F} \subset \mathcal{M}(U)$ and let $a, b, c \in \mathbb{P}^1(\mathbb{C})$ be three fixed, distinct values omitted by every member of $\mathcal{F}$. Then $\mathcal{F}$ is normal.*

Proof. M11 reduces the problem to a source coordinate disk. Endow that disk with its hyperbolic metric and the thrice-punctured target sphere with its complete hyperbolic metric. The source uses the uniformization and Schwarz lemma developed in §2 and §7. In the target, the hyperbolic metric dominates a positive constant times the spherical metric. On a compact set away from the punctures this follows from continuity and positivity of the density ratio. Near a cusp, in a coordinate vanishing at the puncture, the hyperbolic density has cusp behaviour comparable to $1/(|w| \log(1/|w|))$, while the spherical density is bounded. Thus the comparison extends globally with a positive lower constant.

Schwarz–Pick makes every map nonexpanding between the hyperbolic metrics. Composing with the target metric comparison gives a common bound for the spherical derivatives on each smaller source disk. M12 yields normality. □

The hyperbolic uniformization and cusp comparison are explicit imported dependencies in this proof. The present section does not claim a new foundational construction of that metric.

The number of omitted values has a genuine threshold. The functions e^{nz} omit 0 and ∞, but their spherical derivatives at zero equal $n/2$. They are nonnormal there. A rule allowing each member its own three unrelated omitted values would also fail; the source theorem requires a common triple.

4.3.7 Source Theorem 4.9: separated moving values

Theorem 4.15 (M15: moving-target normalization)**.** *Let $\psi_1, \psi_2, \psi_3 : U \to \mathbb{P}^1(\mathbb{C})$ be meromorphic and pairwise distinct at every point. Suppose $f_n(z) \neq \psi_i(z)$ for every z, n, i. Then $\{f_n\}$ is normal.*

Proof. Locally choose nonvanishing holomorphic lifts $v_i(z) \in \mathbb{C}^2$ of the three target points. Their pairwise determinants are nonzero. The projective transformation

$$H_z([u]) = [\det(u, v_1) \det(v_2, v_3) : \det(u, v_3) \det(v_2, v_1)]$$

sends ψ_1, ψ_2, ψ_3 to $0, 1, \infty$. Changing a lift multiplies the two coordinates by the same nonzero factor. Thus the formula defines a holomorphic family of projective transformations without dependence on the chosen lifts.

The functions $g_n(z) = H_z(f_n(z))$ omit $0, 1, \infty$, so M14 gives normality. On a compact source set, $(z, w) \mapsto H_z^{-1}(w)$ is uniformly continuous on the product with the target sphere. Therefore locally uniform spherical convergence of a subsequence of g_n transfers to the corresponding subsequence of f_n. M11 completes the global statement. □

In an affine chart with finite target values, the same transformation is

$$H_z(w) = \frac{(w - \psi_1(z))(\psi_2(z) - \psi_3(z))}{(w - \psi_3(z))(\psi_2(z) - \psi_1(z))}.$$

Its numerical conditioning depends on lower bounds for the pairwise chordal separations. Compactness guarantees a positive minimum for a fixed separated triple on a compact set. An implementation still needs a computable certificate for that minimum.

4.3.8 Source Exercise 4.10: isolated collisions

Theorem 4.16 (M16: distinct meromorphic targets with collisions)**.** *Assume the three meromorphic functions ψ_i are pairwise nonidentical on a connected domain. Their values may coincide at isolated points. If $f_n(z) \neq \psi_i(z)$ for every z, n, i, the family $\{f_n\}$ is normal on the entire domain.*

Proof. The coincidence sets of nonidentical meromorphic maps are discrete, as seen from the zeros of their local projective determinants. Away from these sets, M15 applies. Fix a collision point and choose a source coordinate disk with that point at zero and no other collision.

After a fixed projective change of target chart, one colliding target ψ_1 is holomorphic and finite near zero. The moving transformation $w \mapsto 1/(w - \psi_1(z))$ sends it to infinity. The transformed g_n are holomorphic on the full disk: avoidance of ψ_1 excludes poles, and poles of the original f_n become zeros. The transformed second colliding target, denoted ψ, has a pole of order $m \geq 1$ at zero and no other pole in a sufficiently small disk.

By normality on the punctured disk and M13, every sequence has a further subsequence that is locally bounded there or diverges locally uniformly to infinity. In the bounded case, a uniform bound on one circle transfers to its interior by the maximum principle, and M02 supplies normality across zero.

Suppose instead that $g_n \to \infty$ uniformly on a small circle. For large n, $|g_n| > |\psi|$ on that circle. The curves g_n and $g_n - \psi$ then have the same winding about zero. The first winding is nonnegative because g_n is holomorphic inside. The second winding equals $-m$: the meromorphic function $g_n - \psi$ has no zero, by the omission condition, and has one pole of order m. This contradiction excludes the escaping case. Returning through the moving projective transformation preserves spherical normality. □

This argument expands the source hint on page 633. It gives a direct pole-count contradiction and avoids an unsupported inference from the value of a holomorphic function at a single point to a minimum-modulus estimate on a disk.

4.3.9 Normality sets and algebraic-topological bookkeeping

The union of all open sets on which a family is normal is itself a normality set, by M11. It is the unique maximal such open set. Its complement is closed. For iterates of rational maps this local construction becomes part of the Fatou–Julia decomposition later in the book; no dynamical classification is inferred in the present unit.

Projective determinants locate target collisions, winding counts repair local degeneracies, and a countable exhaustion assembles local convergence. These are separate operations in the proof graph. In particular, the existence of a finite-dimensional projective representation does not remove the analytic compactness argument.

Exercise 4.5 (A nonnormal family with two omitted values)**.** Compute $(e^{nz})^{\#}(0)$. Explain why omitting two sphere values does not meet M14.

Solution. At zero the numerator is n and the denominator is 2, so the spherical derivative is $n/2$. M12 excludes normality near zero. The omitted values are exactly 0 and ∞; a third common value is absent. □

Exercise 4.6 (A complete-metric truncation)**.** If two meromorphic maps differ by at most ε in chordal distance on each of $K_1, \dots, K_M$, bound their metric distance in M09.

Solution. The first M summands total at most $\varepsilon(1 - 2^{-M})$, and the remaining summands total at most 2^{-M}. Thus $d(f, g) \le \varepsilon + 2^{-M}$. This finite certificate assumes full compact-set bounds, with the unobserved tail controlled by the metric definition. □

Exercise 4.7 (Collision normalization)**.** At zero let $\psi_1(z) = 0$, $\psi_2(z) = z^3$ and $\psi_3(z) = 1$. Describe the local transformed functions and the pole order used in M16.

Solution. The transformation is $w \mapsto 1/w$. A family omitting $0, z^3, 1$ becomes a holomorphic family $g_n = 1/f_n$ omitting z^{-3} and 1, with poles of f_n filled as zeros. The moving target z^{-3} has pole order three. In a hypothetical escaping subsequence, the argument principle would give winding -3 for $g_n - z^{-3}$ and nonnegative winding for g_n, contradicting the boundary homotopy. □

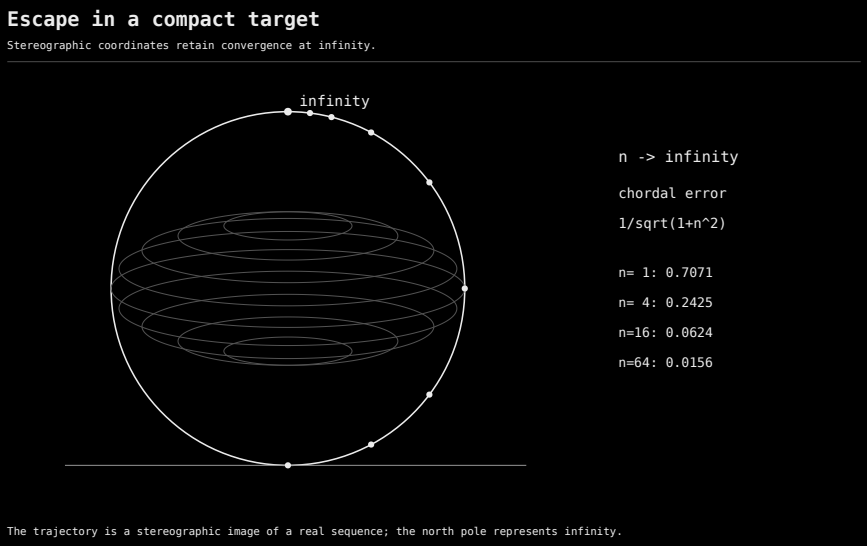


Figure 4.6: Stereographic compactification of the constant sequence $f_n = n$ and the exact chordal error to infinity. The geometric panel and error curve use the same sequence.

4.4 Reference §4.4: normalized univalent functions and distortion

4.4.1 Source Lemma 4.13 and Theorem 4.12

Let $\mathcal{U}$ consist of all univalent $f : \mathbb{D} \to \mathbb{C}$ with $f(0) = 0$ and $|f'(0)| = 1$. Compactness is taken in the topology of local Euclidean-uniform convergence.

Lemma 4.17 (M17: an omitted point on the unit circle)**.** *For every $f \in \mathcal{U}$, the image $f(\mathbb{D})$ omits a point of the unit circle.*

Proof. Suppose the entire unit circle lies in $f(\mathbb{D})$. This image is simply connected and contains zero. The Jordan curve therefore bounds its full closed disk inside the image. Compactness of the closed disk and openness of the image give $D_R \subset f(\mathbb{D})$ for some $R > 1$. The inverse $f^{-1} : D_R \to \mathbb{D}$ fixes zero. Schwarz's lemma after rescaling gives $|(f^{-1})'(0)| \le 1/R < 1$, contradicting $|f'(0)| = 1$. □

Theorem 4.18 (M18: compactness of the normalized class)**.** *The family $\mathcal{U}$ is compact.*

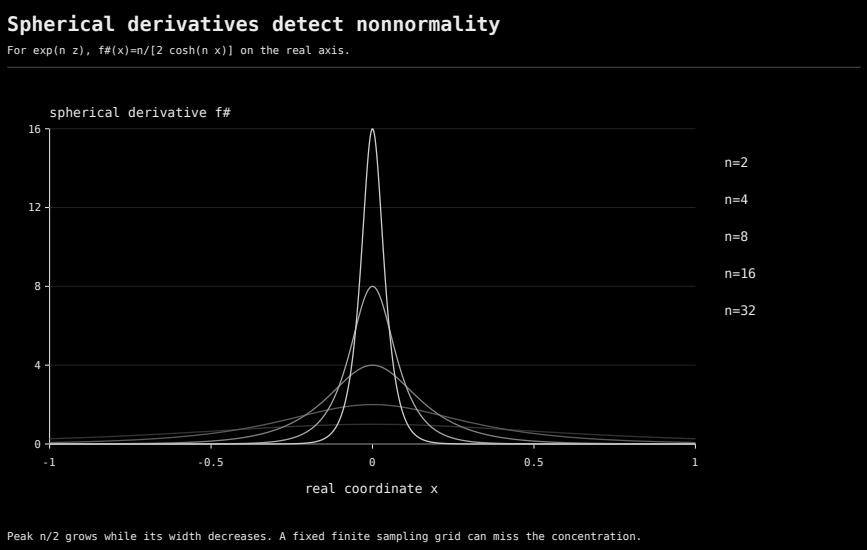


Figure 4.7: Spherical derivatives of $f_n(z) = e^{nz}$ on the real axis. The height is $n/2$ and the width contracts as n^{-1}. The image gives an explicit failure of uniform equicontinuity at zero.

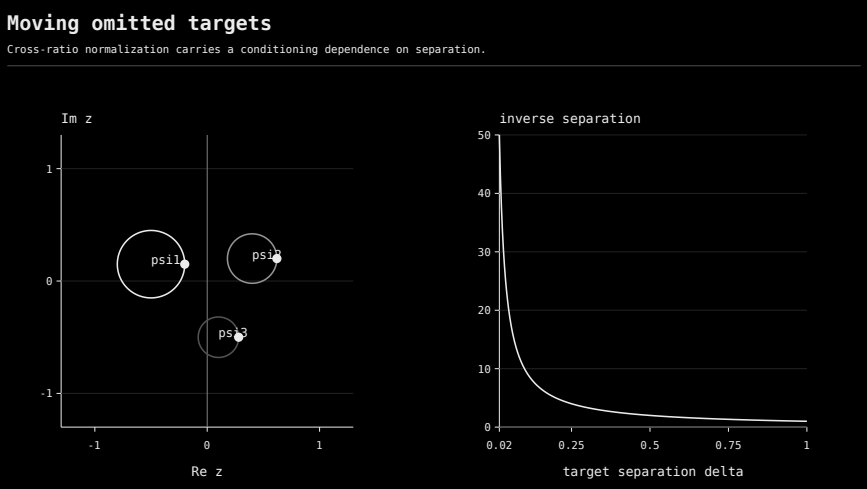


Figure 4.8: A moving triple and its pairwise separation. The normalization is projectively well defined where the determinants are nonzero. Collision locations are recorded explicitly.

Proof. Given a sequence in $\mathcal{U}$, rotate each target so that the corresponding omitted unit-circle point is 1. A subsequence of these rotations converges, so it suffices to treat the rotated family $\mathcal{U}_0$ omitting 1.

On $\mathbb{D} \setminus \{0\}$ every member omits $0, 1, \infty$. Big Montel gives a spherically convergent subsequence there. M10 makes the limit finite holomorphic or identically infinite on this connected punctured domain. The second possibility is excluded as follows. On any circle $|z| = r$, the image loop winds once around zero and zero times around 1. Hence it meets the segment $[0, 1]$. A sequence of such loops cannot tend uniformly to infinity.

The finite limit gives a uniform bound on a fixed circle. The maximum principle extends the bound to its interior, and M02 gives normality across zero. Cauchy's formula preserves the derivative normalization. M05 excludes a constant limit and preserves injectivity. Hurwitz applied to $f_n - 1$ also preserves omission of 1. Thus every sequence has a convergent subsequence with limit in $\mathcal{U}_0$. Undoing the convergent rotations proves compactness of $\mathcal{U}$. □

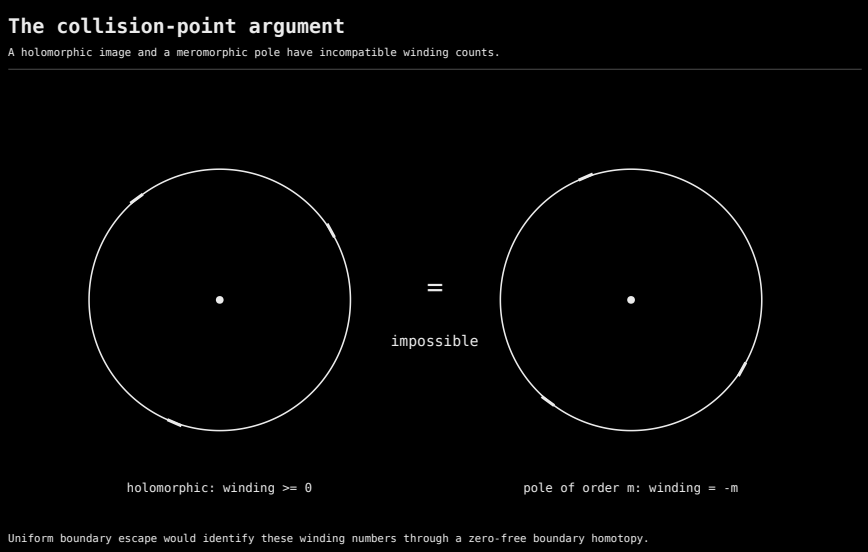


Figure 4.9: The local pole-count contradiction in M16. The annular boundary carries the winding calculation; the pole at the collision point supplies its signed contribution.

The proof is a subsequence argument. Normality of a family on a punctured domain does not, without that extraction, divide the entire family into a single bounded or escaping regime.

4.4.2 Source Theorem 4.11: qualitative distortion

Theorem 4.19 (M19: compactness produces universal interior constants)**.** *For $0 < r < 1$ there exist positive constants A_r, B_r, C_r such that every $f \in \mathcal{U}$ satisfies*

$$A_r \le |f'(z)| \le B_r \quad (|z| \le r), \qquad \frac{|f'(z)|}{|f'(w)|} \le C_r \quad (|z|, |w| \le r).$$

The inner and outer radii of $f(D_r)$ about zero are bounded below and above by positive constants depending only on r.

Proof. Differentiation is continuous for compact-open holomorphic convergence by Cauchy's formula. The functions $(f, z) \mapsto |f'(z)|$ and $(f, z) \mapsto |f(z)|$ are continuous on the product of the compact family with a compact source set. Injectivity implies $f' \neq 0$. Hence the derivative has a positive minimum and a finite maximum on $\mathcal{U} \times \overline{D_r}$.

On $|z| = r$, injectivity and $f(0) = 0$ imply $f(z) \neq 0$. Compactness again gives a uniform positive minimum of $|f(z)|$. The Jordan domain $f(D_r)$ contains the centred disk with that radius. Its outer radius is bounded by the maximum of $|f(z)|$ on the same circle. Dividing derivative bounds proves the quotient estimate. □

For arbitrary univalent ϕ, apply M19 to

$$f(z) = \frac{\phi(z) - \phi(0)}{\phi'(0)}.$$

This gives exactly the scale dependence in source estimates (4.1)–(4.3). The sharp constants require the quantitative arguments below.

4.4.3 Source Exercise 4.14: precompactness and omitted values

Proposition 4.20 (M20: precompactness inside the univalent class)**.** *A family of univalent maps on $\mathbb{D}$ has compact closure consisting entirely of univalent maps if and only if*

$$|f(0)| \le C, \qquad C^{-1} \le |f'(0)| \le C$$

for a common $C > 0$.

Proof. Normalize each member by subtracting its value at zero and dividing by its derivative. M18 supplies compactness of the normalized functions. The affine parameters range over compact sets, with derivative bounded away from zero, so all limits remain univalent. Conversely, evaluation and derivative evaluation are continuous on the compact closure. The derivative is nonzero at every member of that closure and therefore has a positive minimum as well as a finite maximum. □

Proposition 4.21 (M21: omitted-value criterion on a proper domain)**.** *Let $\Omega \subsetneq \mathbb{C}$ be a domain, $a \in \Omega$, and let $\mathcal{F}$ be a family of univalent maps satisfying $|f(a)| \le C$. The family is normal if and only if there is $\rho > 0$ such that every $f \in \mathcal{F}$ omits a finite value of modulus less than ρ.*

Proof. Assume first the omitted-value bound. Choose $D(a, r) \Subset \Omega$. The quarter theorem M26 applied to the restriction gives

$$D\left(f(a), \frac{r|f'(a)|}{4}\right) \subset f(\Omega).$$

An omitted value of modulus less than ρ consequently implies $|f'(a)| \le 4(C + \rho)/r$. The growth and derivative bounds on disks propagate the bounds for f and f' through a finite chain of overlapping disks from a to any compact set in Ω. A finite family of such chains covers the compact set. Thus the family is locally bounded, and M02 applies.

Conversely, normality and the bounded value at a exclude an infinite spherical limit by M10. Every sequence therefore has locally Euclidean-bounded further subsequences. Cauchy's estimate implies a common bound $|f'(a)| \le M$; an unbounded derivative sequence would contradict that subsequence conclusion.

Choose $b \notin \Omega$ and put $d = |a - b|$. If $D(f(a), R) \subset f(\Omega)$, apply M26 to the inverse on this disk. Its image contains the disk about a of radius $R/(4|f'(a)|)$, which must omit b. Hence $R \le 4Md$. The same argument excludes $f(\Omega) = \mathbb{C}$. A complementary value can therefore be found at distance at most $4Md + 1$ from $f(a)$. The common bound $\rho = C + 4Md + 2$ suffices. □

The hypothesis $\Omega \ne \mathbb{C}$ is necessary in the literal formulation. The singleton family $\{z \mapsto z\}$ on $\mathbb{C}$ is normal and omits no finite value. This is an explicit qualification of source Exercise 4.14(ii).

4.4.4 An area proof for the quantitative constants

Theorem 4.22 (M22: the exterior area theorem)**.** *Suppose F is univalent on $\{|\zeta| > 1\}$ and has Laurent expansion*

$$F(\zeta) = \zeta + b_0 + \sum_{n \ge 1} b_n \zeta^{-n}.$$

Then

$$\sum_{n \ge 1} n|b_n|^2 \le 1.$$

Proof. For $R > 1$, the image of $|\zeta| = R$ is a positively oriented analytic Jordan curve. Its enclosed area is

$$\frac{1}{2i}\int_{F(|\zeta|=R)} \overline{w}\, dw = \pi\left(R^2 - \sum_{n\geq 1} n|b_n|^2 R^{-2n}\right).$$

To obtain the identity, substitute $\zeta = Re^{it}$, differentiate the Laurent series and integrate. The Fourier cross terms vanish. Uniform convergence of the series and its derivative on this circle justifies the operation. The enclosed area is nonnegative. Thus every finite partial sum of the nonnegative series is at most R^2. Let $R \downarrow 1$, then let the partial-sum length tend to infinity. □

This is a functional-analytic coefficient budget with a geometric origin. Its coefficients belong to a weighted ℓ^2 ball. The geometry of univalence is used in the positivity of the enclosed Jordan area.

Corollary 4.23 (M23: the second coefficient bound)**.** *If $f(z) = z + a_2 z^2 + \cdots$ is univalent on $\mathbb{D}$, then $|a_2| \leq 2$.*

Proof. The function $f(z^2)/z^2$ extends as a zero-free holomorphic function taking value 1 at zero. Its square root with value 1 at zero is even. Therefore

$$g(z) = \sqrt{f(z^2)} = z + \frac{a_2}{2} z^3 + \cdots$$

is an odd holomorphic function. It is univalent: equality of two values implies equality of the squares of the source points by injectivity of f; the opposite-point possibility is excluded by oddness unless both points are zero.

The exterior function $1/g(1/\zeta) = \zeta - (a_2/2)\zeta^{-1} + \cdots$ is univalent. M22 gives $|a_2/2|^2 \leq 1$. □

4.4.5 Sharp derivative and growth bounds

Theorem 4.24 (M24: derivative distortion)**.** *For a normalized univalent function $f(0) = 0$, $f'(0) = 1$ and $|z| = r < 1$,*

$$\frac{1-r}{(1+r)^3} \leq |f'(z)| \leq \frac{1+r}{(1-r)^3}.$$

Consequently, on $\overline{D_r}$,

$$\frac{|f'(z)|}{|f'(w)|} \leq \left(\frac{1+r}{1-r}\right)^4.$$

Proof. Fix $a \in \mathbb{D}$. The normalized Koebe transform is

$$F_a(\zeta) = \frac{f((a+\zeta)/(1+\overline{a}\zeta)) - f(a)}{(1-|a|^2)f'(a)}.$$

Its second coefficient is

$$\frac{1-|a|^2}{2}\frac{f''(a)}{f'(a)} - \overline{a}.$$

M23 bounds its modulus by two. For $a = te^{i\theta}$, take the real part after multiplication by $e^{i\theta}$ and obtain

$$\frac{2t-4}{1-t^2} \leq \frac{d}{dt}\log|f'(te^{i\theta})| \leq \frac{2t+4}{1-t^2}.$$

Integrate from zero to r, using $|f'(0)| = 1$. The two antiderivatives are the logarithms of the displayed bounds. The quotient estimate follows by division. The Koebe function at r and $-r$ attains the quotient constant in the limiting closed-disk statement. □

Theorem 4.25 (M25: growth, inner radius and outer radius)**.** *Under the same normalization,*

$$\frac{r}{(1+r)^2} \le |f(z)| \le \frac{r}{(1-r)^2}, \qquad |z| = r < 1.$$

In particular, the inner and outer radii of $f(D_r)$ satisfy

$$\frac{r}{(1+r)^2} \le r_{f(D_r)}(0) \le R_{f(D_r)}(0) \le \frac{r}{(1-r)^2}.$$

The corresponding shape ratio is at most $((1+r)/(1-r))^2$.

Proof. Integrate the upper derivative bound along the radial segment from zero to z; its integral is $r/(1-r)^2$.

For the lower bound, choose a point of minimum modulus on $f(|z| = r)$. The radial segment from zero to that image point is contained in the closure of $f(D_r)$. Lift its interior by f^{-1} to a curve γ from zero to $|z| = r$. Its image length equals the selected modulus. M24 gives

$$\text{length}(f \circ \gamma) \ge \int_\gamma \frac{1-|z|}{(1+|z|)^3}|dz| \ge \int_0^r \frac{1-t}{(1+t)^3}\,dt = \frac{r}{(1+r)^2}.$$

For the second inequality, compose a primitive of the positive radial weight with $|\gamma|$ and use total variation. No monotonicity of the lifted radial coordinate is assumed. Since the selected image point has minimum modulus, the estimate holds on the entire image circle. The Jordan-domain inclusions and shape estimate follow. □

The domain reduction $r < 1$ is quantitative. As $r \uparrow 1$, the class-wide upper growth and derivative bounds diverge. These bounds are compatible with compactness on every fixed smaller disk.

4.4.6 The quarter theorem and the extremal slit

Theorem 4.26 (M26: Koebe quarter theorem)**.** *If f is univalent on $\mathbb{D}$, $f(0) = 0$ and $f'(0) = 1$, then $D_{1/4} \subset f(\mathbb{D})$.*

Proof. Let $b \notin f(\mathbb{D})$. It is nonzero. The function

$$g(z) = \frac{f(z)}{1 - f(z)/b} = z + (a_2 + 1/b)z^2 + \cdots$$

is normalized and univalent. Applying M23 to f and g gives $|a_2| \le 2$ and $|a_2 + 1/b| \le 2$. Hence $|1/b| \le 4$, so every omitted value has modulus at least $1/4$. □

Proposition 4.27 (M27: source Exercise 4.15)**.** *The Koebe function*

$$k(z) = \frac{z}{(1-z)^2}$$

is a conformal equivalence

$$k : \mathbb{D} \longrightarrow \mathbb{C} \setminus (-\infty, -1/4].$$

Its inverse on that slit plane is

$$k^{-1}(w) = \frac{\sqrt{1+4w} - 1}{\sqrt{1+4w} + 1} = \frac{4w}{(1+\sqrt{1+4w})^2},$$

where the square root has positive real part.

Proof. The Cayley map $s = (1+z)/(1-z)$ sends $\mathbb{D}$ bijectively to the right half-plane. Squaring sends that half-plane bijectively to the plane cut along the nonpositive real axis. The identity

$$4k(z) + 1 = \left(\frac{1+z}{1-z}\right)^2$$

proves the image statement. Solving for z gives the inverse. The second expression rationalizes the numerator and extends continuously at $w = 0$.

Alternatively, the algebraic difference

$$k(z) - k(w) = \frac{(z-w)(1-zw)}{(1-z)^2(1-w)^2}$$

proves injectivity directly on the disk. The slit endpoint is approached at $z = -1$, and $k(-r) \to -1/4$ as $r \uparrow 1$. Therefore no class-wide replacement of the quarter radius by a larger constant is possible. □

The rational map $k : \mathbb{P}^1(\mathbb{C}) \to \mathbb{P}^1(\mathbb{C})$ has degree two. Its branch values are $-1/4$ and ∞, corresponding to the critical point -1 and the double pole at 1. The algebraic involution $z \mapsto 1/z$ exchanges the two generic sheets. Restriction to $\mathbb{D}$ selects the sheet that maps zero to zero. This relates the present extremal function to the degree and monodromy analysis in the preceding corpus.

4.4.7 Source Corollary 4.18

Corollary 4.28 (M28: conformal radius and Euclidean clearance)**.** *Let $\phi : \mathbb{D} \to D$ be a conformal equivalence with $\phi(0) = a$, and let $d = \operatorname{dist}(a, \partial D)$. Then*

$$d \le |\phi'(0)| \le 4d.$$

Proof. For every $r < d$, the inverse map restricts to $D(a, r) \to \mathbb{D}$. Schwarz's lemma gives $1/|\phi'(0)| \le 1/r$. Let $r \uparrow d$. For the other inequality, normalize ϕ and apply M26 to obtain $D(a, |\phi'(0)|/4) \subset D$. □

The target must be the image domain. The inclusion map $\mathbb{D} \hookrightarrow D_{100}$ has derivative one and target clearance one hundred; it violates the left inequality when surjectivity is dropped.

Exercise 4.8 (Interior distortion at one half)**.** For a normalized univalent function, give derivative bounds on $\overline{D_{1/2}}$, a derivative-ratio bound, and a shape bound.

Solution. The derivative bounds are $4/27$ and 12. Their ratio is 81. The growth radii are $2/9$ and 2, so the shape is bounded by 9. Derivative distortion and shape are different quantities, with different powers of the same radius ratio. □

Exercise 4.9 (Degeneration to a constant)**.** Explain why $f_n(z) = z/n$ is normal but fails the compactness conclusion inside the space of univalent maps.

Solution. It converges uniformly on compact sets to zero. The limit is constant, hence nonunivalent. The lower derivative bound in M20 fails because $f_n'(0) = 1/n$. Its normality in a larger function space is fully consistent with that failure. □

Exercise 4.10 (An exact Koebe truncation error)**.** Let $k_N(z) = \sum_{j=1}^{N} jz^j$. Prove

$$\sup_{|z| \le r} |k(z) - k_N(z)| = \frac{r^{N+1}((N+1) - Nr)}{(1-r)^2}, \quad 0 < r < 1.$$

Solution. The tail is $\sum_{j>N} jz^j$. The triangle inequality bounds its modulus by the positive real tail. Differentiating the geometric-series tail gives the displayed expression. Equality holds at $z = r$, where all terms are nonnegative real numbers. Thus the formula is an exact supremum, including its extremizing point. □

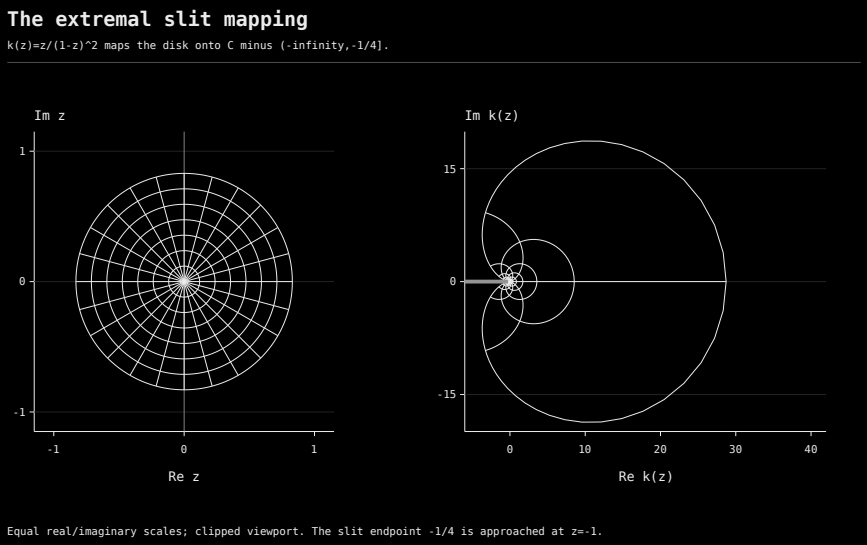


Figure 4.10: Conformal images of radial and circular coordinate curves under the Koebe function. The omitted slit is marked, and the singularity at the boundary point 1 is excluded from the plotted parameter range.

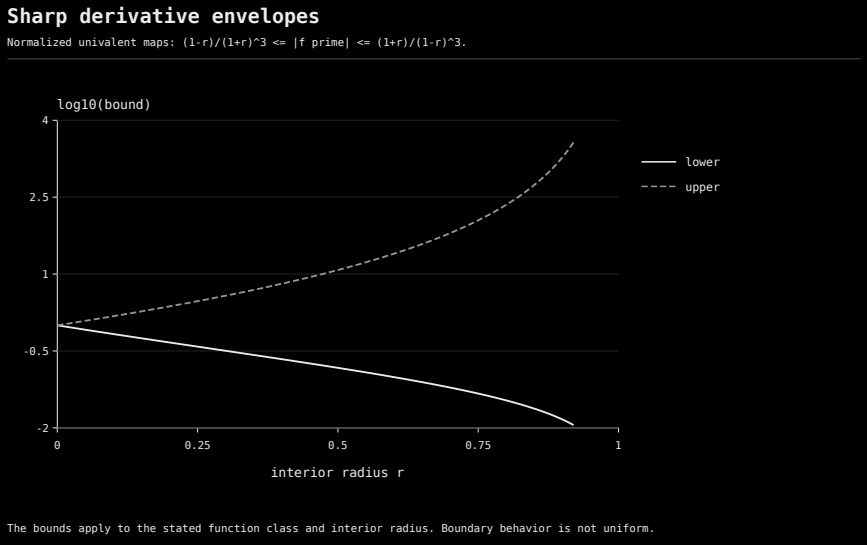


Figure 4.11: Sharp derivative envelopes and their values for the Koebe extremizer on the positive and negative real axes.

Mathematical reference. [3], Theorems 4.11–4.12, Lemma 4.13, Exercise 4.14, the quarter theorem, Exercise 4.15 and Corollary 4.18, pp.113–115. The area-theorem derivation and explicit sharp inequalities are supplied developments.

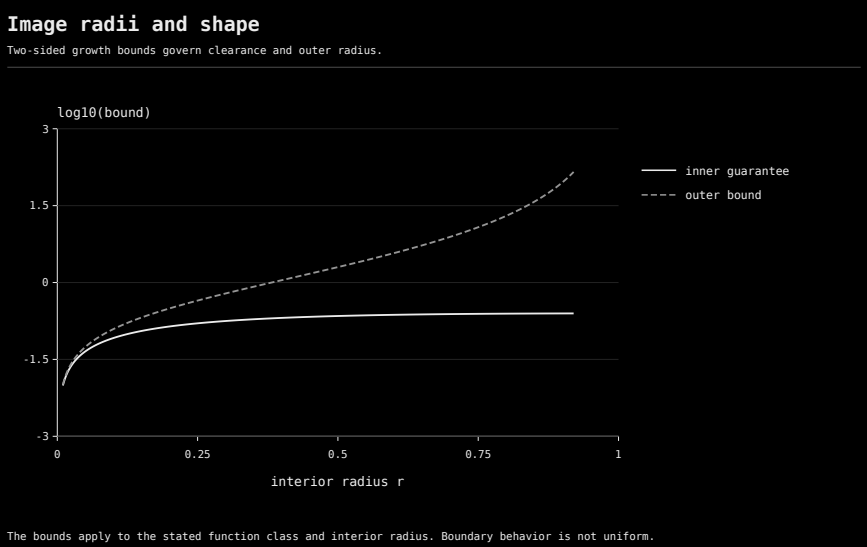


Figure 4.12: Growth and radius envelopes for normalized univalent functions. The lower and upper radial limits produce different geometric constraints.

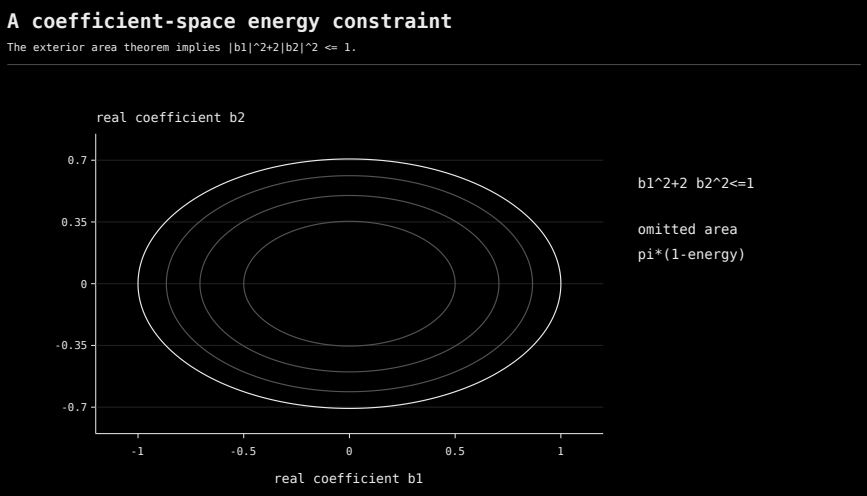


Figure 4.13: The exterior area theorem as a weighted coefficient budget. The displayed finite slice is a necessary projection of the full infinite constraint.

4.5 Annular collars and source Lemma 4.17

4.5.1 The domain of the modulus statement

For the ring-domain interpretation, let $K \Subset \mathbb{D}$ be a nondegenerate full continuum containing zero. Here "full" means that its planar complement is connected. Then $\mathbb{D} \setminus K$ is doubly connected. A disconnected arbitrary compact set can produce several complementary components; its complement need not possess the single annular modulus used in the source proof. The point continuum is a degenerate limiting case.

We use the variational condenser capacity

$$\operatorname{Cap}(K, \mathbb{D}) = \inf_{u} \int_{\mathbb{D}} |\nabla u|^2 \, dA,$$

where the infimum is over smooth or Lipschitz test functions with compact support in $\mathbb{D}$ and

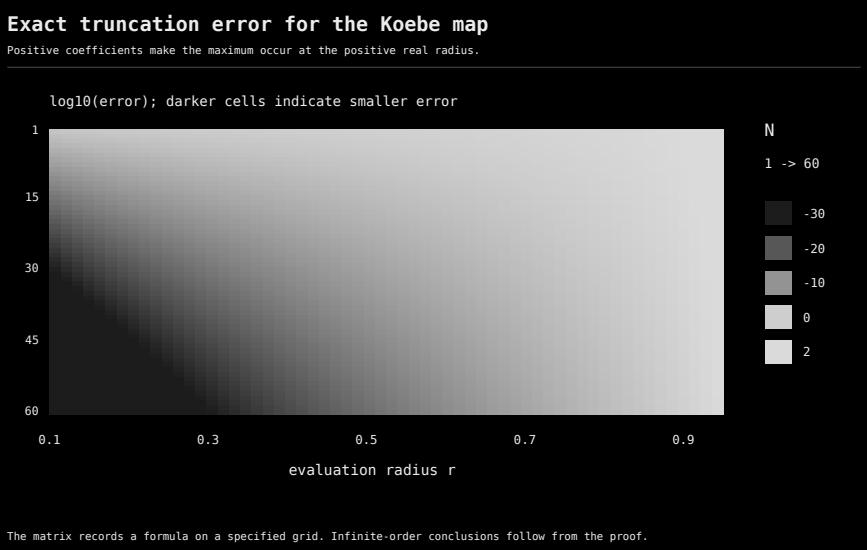


Figure 4.14: Exact uniform errors of Koebe Taylor truncations on four closed subdisks. These values include an analytic extremum certificate.

value one on a neighbourhood of K. Equivalent Sobolev formulations follow by approximation. Conformal invariance of Dirichlet energy and the round-annulus calculation give

$$\operatorname{Cap}(K, \mathbb{D}) = \frac{1}{\operatorname{mod}(\mathbb{D} \setminus K)}.$$

Indeed, on $A(\rho, 1)$ the radial potential is $u(r) = \log r / \log \rho$ and its energy is $2\pi / \log(1/\rho)$. Cauchy–Schwarz on every radial segment shows that no admissible function has smaller energy. Smooth truncation near the two boundary circles approaches the same value. Under a conformal uniformization, the truncated functions are constant near the ends and extend to admissible functions in the physical domain. The Jacobian factor cancels the squared gradient factor.

Mathematical reference. The modulus convention is from [3], pp.73–74. Ring uniformization and the extremal-length/Dirichlet principle are explicit dependencies from source §2 and §6.

4.5.2 A quantitative substitute for the nondegeneration step

Theorem 4.29 (M29: an explicit collar-to-radius bound)**.** *Let K satisfy the preceding hypotheses, let $r = \max_{z \in K} |z|$, and suppose* $\operatorname{mod}(\mathbb{D} \setminus K) \geq \mu > 0$. *Then*

$$r \leq k(\mu) := 1 - \exp\left(-\frac{\pi}{4\mu}\right) < 1.$$

This bound is universal; no sharpness is asserted for its constant.

Proof. Choose $z_* \in K$ with $|z_*| = r$ and put $\xi = z_*/r \in \partial\mathbb{D}$. Connectedness of K, together with $0, z_* \in K$, implies that the distance function $z \mapsto |z - \xi|$ assumes every value between $1 - r$ and 1 on K.

For $1 - r < t < 1$, the part of the circle $|z - \xi| = t$ lying in $\mathbb{D}$ is an arc of length

$$\ell(t) = 2t \arccos(t/2) \leq \pi t.$$

It meets K. Any admissible u vanishes at the two arc endpoints and takes value one at a point of K. If that point divides the arc into lengths a, b, one-dimensional Cauchy–Schwarz gives

$$\int_{\text{arc}} |\partial_s u|^2 \, ds \geq \frac{1}{a} + \frac{1}{b} \geq \frac{4}{a+b} \geq \frac{4}{\pi t}.$$

Integrating over t in polar coordinates centred at ξ yields

$$\int_{\mathbb{D}} |\nabla u|^2 dA \geq \frac{4}{\pi} \log \frac{1}{1-r}.$$

Taking the infimum and using capacity reciprocity gives

$$\frac{4}{\pi} \log \frac{1}{1-r} \leq \frac{1}{\operatorname{mod}(\mathbb{D} \setminus K)} \leq \frac{1}{\mu}.$$

Solving for r proves the assertion. □

The source normal-family proof passes from annulus uniformizations to a claimed Jordan limit. Normality alone permits a constant limit. M29 supplies a variational proof that includes the missing nondegeneration control.

The source additionally identifies radial slits as extremal continua. The circular rearrangement proof in M38 establishes that assertion and identifies the optimal radius through the radial-slit modulus. M29 supplies a separate explicit elementary upper bound.

4.5.3 Source Theorem 4.16: invariant distortion

Theorem 4.30 (M30: modulus-controlled distortion)**.** *Let D be a proper simply connected planar domain and let $\Delta \Subset D$ be a domain with full compact closure such that $D \setminus \overline{\Delta}$ has modulus at least $\mu > 0$. Every univalent $\phi : D \to \mathbb{C}$ satisfies*

$$\frac{|\phi'(z)|}{|\phi'(w)|} \leq \left(\frac{1+k(\mu)}{1-k(\mu)} \right)^8, \qquad z, w \in \Delta,$$

where $k(\mu)$ is from M29.

Proof. Choose a uniformization $h : \mathbb{D} \to D$ with $h(0) \in \Delta$. The compact full continuum $K = h^{-1}(\overline{\Delta})$ contains zero and has the same complementary modulus. M29 places it in $\overline{D_k}$. Both h and $F = \phi \circ h$ are univalent on the unit disk. Apply M24 to each, after affine normalization. If $z = h(u)$ and $w = h(v)$, then

$$\frac{|\phi'(z)|}{|\phi'(w)|} = \frac{|F'(u)|}{|F'(v)|} \frac{|h'(v)|}{|h'(u)|}.$$

Each factor is at most $((1+k)/(1-k))^4$. □

The exponent eight comes from two independently controlled coordinate maps. A known round collar allows a substantially tighter radius substitution: for $\Delta = D_r \subset D = \mathbb{D}$, the exact modulus is $(2\pi)^{-1} \log(1/r)$ and only the one-map exponent four is needed. The comparison records the extra geometric information used in that improvement.

4.5.4 A coefficient-specific geometric model

Proposition 4.31 (M31: global bounds for a quadratic embedding)**.** *Let $\phi_a(z) = z + az^2$ with $|a| < 1/2$. For $z, w \in \overline{D_r}$, $r \leq 1$,*

$$(1 - 2|a|r)|z - w| \leq |\phi_a(z) - \phi_a(w)| \leq (1 + 2|a|r)|z - w|.$$

In particular ϕ_a is univalent, and its derivative ratio on this disk is at most

$$Q(a, r) = \frac{1 + 2|a|r}{1 - 2|a|r}.$$

Proof. Factor the difference as $(z-w)(1+a(z+w))$. Apply the triangle and reverse triangle inequalities, using $|z+w| \le 2r$. The derivative is $1+2az$, so the same bounds apply to its modulus. Positivity of the lower bound proves injectivity. □

For $a = 1/5$, $r = 3/5$, the specialized ratio is $31/19$, while the class-wide Koebe ratio is 256. The latter enclosure is $4864/31 \approx 156.90$ times wider as a multiplicative upper bound. Both are valid for the same known map; the smaller bound uses its exact polynomial coefficients. This is an enclosure comparison, with no associated runtime claim.

4.5.5 Approximation and injectivity certificates

Proposition 4.32 (M32: derivative error from a function error)**.** *Let f, g be holomorphic on D_R and satisfy $|f-g| \le \varepsilon$ there. For $0 < r < R$,*

$$\sup_{D_r} |f' - g'| \le \frac{\varepsilon}{R-r}.$$

If $|f'| \ge m > \varepsilon/(R-r)$ on D_r, then g' has no zero there.

Proof. Apply M01 to $f-g$ on disks of radius less than $R-r$ centred in D_r. The second assertion follows from the reverse triangle inequality. □

A nonzero derivative establishes local conformality. Global injectivity needs a two-point estimate.

Proposition 4.33 (M33: global perturbation of a certified embedding)**.** *Let f be holomorphic on a convex domain V and satisfy $|f(z)-f(w)| \ge m|z-w|$ for $z, w \in V$, with $m > 0$. If g is holomorphic and $\sup_V |g'-f'| \le \delta < m$, then*

$$|g(z)-g(w)| \ge (m-\delta)|z-w|.$$

Thus g is univalent on V.

Proof. Integrate $g'-f'$ along the line segment joining z and w. The resulting difference is at most $\delta|z-w|$. Subtract this from the lower bound for f. □

For the model in M31, take $m = 1-2|a|r$. Combining M32 and M33 yields a complete global certificate on D_r whenever $\varepsilon/(R-r) < 1-2|a|r$. Every quantity in this implication has an explicit domain and a strict margin.

Exercise 4.11 (A numerical-domain certificate)**.** Let $a = 1/5$, $r = 3/5$, $R = 4/5$. An analytic approximation g to ϕ_a satisfies $|g-\phi_a| \le 1/100$ on D_R. Prove a global lower Lipschitz bound for g on D_r.

Solution. M31 gives $m = 1-6/25 = 19/25$. M32 gives $\delta \le (1/100)/(1/5) = 1/20$. Therefore $m-\delta = 71/100 > 0$. By M33, $|g(z)-g(w)| \ge (71/100)|z-w|$. This proves global injectivity on the smaller disk. □

Exercise 4.12 (Two independent collar bounds)**.** For a concentric disk $K = \overline{D_{1/2}}$, compute its exact complementary modulus and compare the exact radius with the bound in M29.

Solution. The modulus is $\mu = (\log 2)/(2\pi)$. M29 gives $r \le 1-\exp(-\pi^2/(2\log 2))$, approximately 0.99919. The exact radius is $1/2$. The large discrepancy quantifies the price of using only the general continuum estimate. It does not invalidate either bound. □

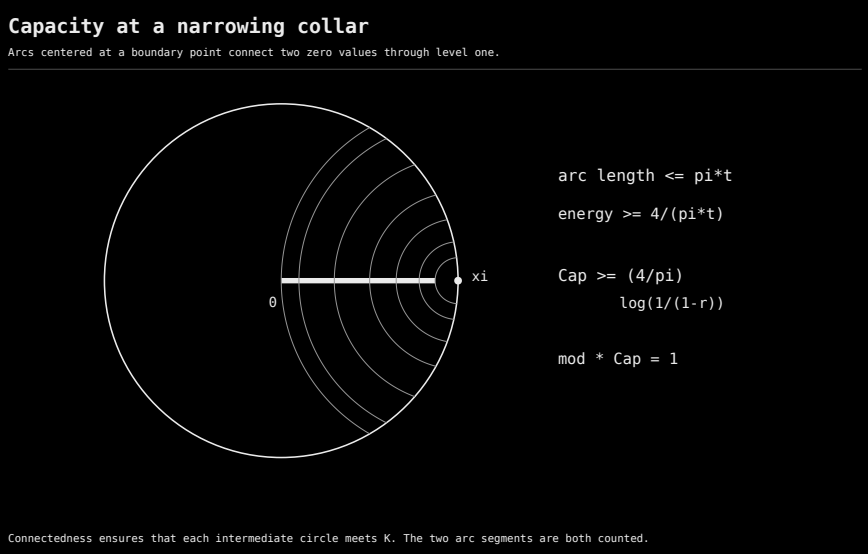


Figure 4.15: The arc family used in M29. Each arc meets the connected inner continuum and has two boundary endpoints with prescribed potential zero.

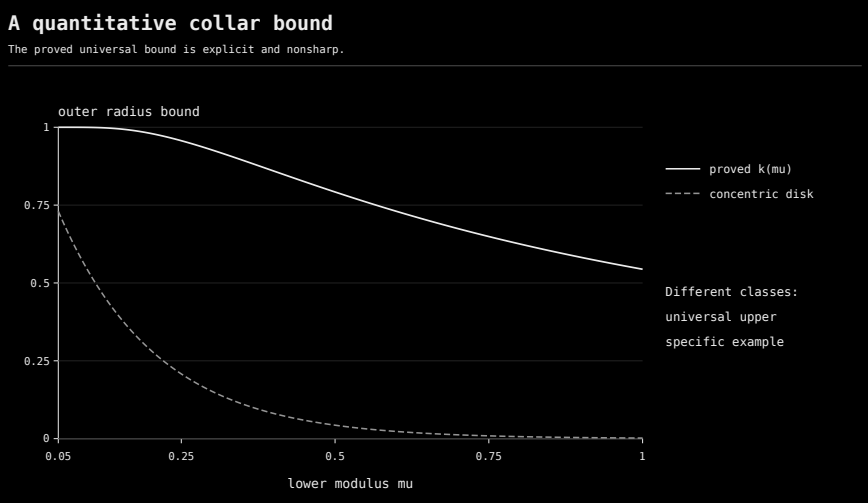


Figure 4.16: Exact concentric-disk radii and the universal continuum enclosure as functions of the modulus. The displayed normalization is $(2\pi)^{-1}\log(R/r)$.

4.6 Additional functional and projective derivations

4.6.1 A Hilbert model for interior restriction

Define $H^2(\mathbb{D})$ by the coefficient norm

$$f(z)=\sum_{j\ge 0} a_j z^j, \qquad \|f\|_{H^2}^2=\sum_{j\ge 0}|a_j|^2<\infty.$$

Cauchy–Schwarz shows that these series converge uniformly on each smaller disk:

$$|f(z)|\le \frac{\|f\|_{H^2}}{\sqrt{1-|z|^2}}.$$

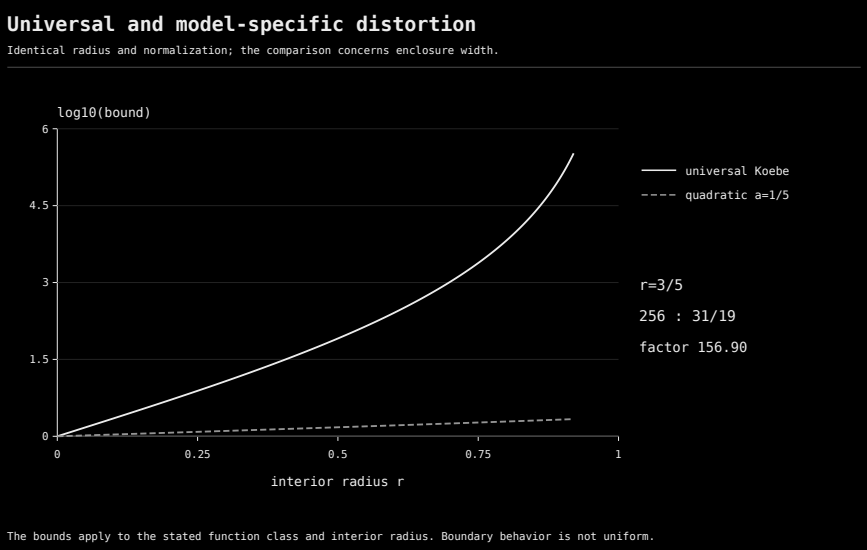


Figure 4.17: Universal Koebe distortion and the coefficient-specific quadratic bound. The two curves concern the same explicit family; their input information is recorded in M31.

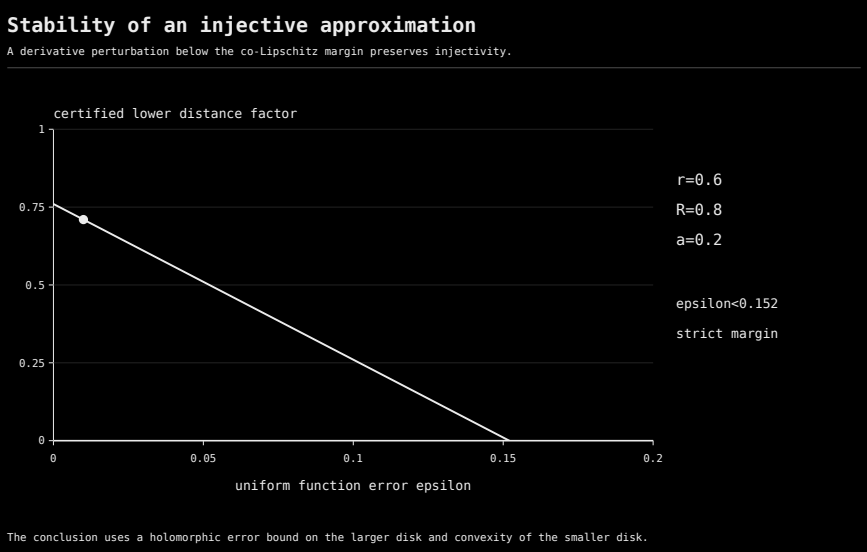


Figure 4.18: The admissible approximation-error region from M32–M33. Points below the threshold retain a positive two-point separation margin.

The coefficient map is a Hilbert-space isometry with $\ell^2(\mathbb{N}_0)$. For $0 < r < 1$, the rescaled restriction operator is

$$T_r f(z) = f(rz), \qquad T_r(a_0, a_1, \ldots) = (a_0, ra_1, r^2 a_2, \ldots).$$

Proposition 4.34 (M34: compactness and operator ideals)**.** *The operator T_r is compact, positive and self-adjoint in coefficient coordinates. Its singular values are $1, r, r^2, \ldots$. Its Hilbert–Schmidt norm squared is $(1-r^2)^{-1}$ and its trace is $(1-r)^{-1}$.*

Proof. The monomials form an orthonormal basis and are eigenvectors with positive eigenvalues r^j. Truncation to the first N coordinates has operator error $r^N \to 0$, proving compactness. The Hilbert–Schmidt and trace sums are the corresponding geometric series. Self-adjointness follows directly from the real diagonal coefficients. □

This model supplies a quantitative instance of the compact restriction mechanism in Little

Montel. Its hypotheses are those of the Hardy coefficient ball. They do not describe every locally bounded holomorphic family with the same numerical constants.

4.6.2 An exact optimality theorem

Theorem 4.35 (M35: optimal rank for Hardy restriction)**.** *For $N \geq 0$,*

$$\inf_{\substack{A:H^2\to H^2 \text{ bounded linear} \\ \operatorname{rank} A\leq N}} \|T_r - A\|_{H^2\to H^2} = r^N.$$

The infimum is attained by retaining the first N Taylor coefficients and applying the corresponding diagonal multipliers.

Proof. Let P_N retain coordinates $0, \ldots, N-1$. Then

$$\|T_r - T_r P_N\|^2 = \sup_{\|a\|_2=1} \sum_{j\geq N} r^{2j}|a_j|^2 = r^{2N},$$

with equality for the unit vector at index N.

For the lower bound, let A have rank at most N and restrict it to $W = \operatorname{span}\{1, z, \ldots, z^N\}$. This domain has dimension $N+1$, so the restricted map has a nonzero kernel vector. Normalize it to $\|u\|_{H^2} = 1$. Then $Au = 0$ and

$$\|(T_r - A)u\|^2 = \|T_r u\|^2 = \sum_{j=0}^{N} r^{2j}|u_j|^2 \geq r^{2N}.$$

The lower and upper bounds agree. □

For a prescribed operator tolerance $0 < \varepsilon < 1$, the least admissible rank is

$$N_{\min} = \left\lceil \frac{\log \varepsilon}{\log r} \right\rceil.$$

The cost model counts complex linear coordinates. Quantization, coefficient acquisition, nonlinear encoders and hardware latency require additional models. Within the stated class, M35 is a genuine global optimum with an explicit lower-bound witness.

4.6.3 Bergman space and conformal radius

Let $A^2(D)$ be the space of holomorphic functions square-integrable with respect to Euclidean area. For a conformal equivalence $\phi : \mathbb{D} \to D$, define

$$(U_\phi F)(z) = F(\phi(z))\phi'(z).$$

Theorem 4.36 (M36: unitary pullback and evaluation bounds)**.** *The map $U_\phi : A^2(D) \to A^2(\mathbb{D})$ is unitary. If $a = \phi(0)$, the Bergman kernel satisfies*

$$K_D(a, a) = \frac{1}{\pi \operatorname{crad}(D, a)^2}.$$

With $d = \operatorname{dist}(a, \partial D)$,

$$\frac{1}{16\pi d^2} \leq K_D(a, a) \leq \frac{1}{\pi d^2}.$$

Proof. Conformal change of variables gives

$$\int_{\mathbb{D}} |F(\phi(z))\phi'(z)|^2 dA(z) = \int_D |F(w)|^2 dA(w).$$

The inverse uses ϕ^{-1} and its derivative, so the isometry is onto. On the unit disk, polar integration gives the orthonormal basis $\sqrt{(j+1)/\pi}\, z^j$. Summing its reproducing series yields

$$K_{\mathbb{D}}(z,w) = \frac{1}{\pi(1-z\overline{w})^2}.$$

The unitary transformation gives

$$K_D(\phi(z),\phi(w))\phi'(z)\overline{\phi'(w)} = K_{\mathbb{D}}(z,w).$$

Set $z = w = 0$ and use M28 for the distance comparison. □

The factor ϕ' is the transformation law of a holomorphic differential. An unweighted composition is generally not an isometry for area measure. In geometric terms, the pullback acts on sections of the canonical bundle; in Hilbert-space terms, it preserves a norm and hence reproducing evaluation functionals.

4.6.4 The Schwarzian as a projective differential invariant

For a locally univalent function define

$$S_f = \frac{f'''}{f'} - \frac{3}{2}\left(\frac{f''}{f'}\right)^2.$$

Direct differentiation gives the chain rule

$$S_{f\circ g} = (S_f \circ g)(g')^2 + S_g.$$

A Möbius map has zero Schwarzian, so postcomposition by such a map leaves S_f unchanged.

Proposition 4.37 (M37: a projectively invariant univalence bound)**.** *Every univalent map on* $\mathbb{D}$ *satisfies*

$$(1-|z|^2)^2|S_f(z)| \le 6.$$

The constant is attained by the Koebe function on the real diameter.

Proof. Normalize $f(z) = z + a_2z^2 + a_3z^3 + \cdots$. The exterior function

$$\frac{1}{f(1/\zeta)} = \zeta - a_2 + (a_2^2 - a_3)\zeta^{-1} + \cdots$$

is univalent. M22 gives $|a_3 - a_2^2| \le 1$. Since $S_f(0) = 6(a_3 - a_2^2)$, the estimate holds at zero. Apply the normalized Koebe transform at an arbitrary point and use the chain rule to obtain the factor $(1-|z|^2)^2$. For $k(z) = z/(1-z)^2$, differentiation gives $S_k(z) = -6/(1-z^2)^2$, proving equality on the real diameter. □

On a simply connected domain where $f' \neq 0$, choose a square root of f'. The functions

$$y_1 = \frac{f}{\sqrt{f'}}, \qquad y_2 = \frac{1}{\sqrt{f'}}$$

solve $y'' + \frac{1}{2}S_f y = 0$, and their quotient is f. The identity follows by differentiation. It connects projective mapping data to a second-order linear equation. A converse univalence criterion would require a separate theorem; the bound above is used only in its proved direction.

Exercise 4.13 (A rank lower-bound witness)**.** For $r = 1/2$ and a rank-at-most-three operator A, identify the finite-dimensional space used to prove $\|T_r - A\| \geq 1/8$.

Solution. Use $W = \operatorname{span}\{1, z, z^2, z^3\}$. Its dimension is four. The restriction of A has a nonzero kernel vector $u \in W$. After normalization, all four diagonal factors of $T_{1/2}$ on W are at least $1/8$, so $\|(T_{1/2} - A)u\| \geq 1/8$. □

Exercise 4.14 (A point-evaluation certificate)**.** Suppose $F \in A^2(D)$, $\|F\|_{A^2} \leq 1$, and $\operatorname{crad}(D, a) = 2$. Bound $|F(a)|$ and identify an extremizer up to phase.

Solution. The reproducing property gives $|F(a)| \leq \sqrt{K_D(a, a)} = 1/(2\sqrt{\pi})$. Equality is attained by the normalized reproducing kernel $K_D(\cdot, a)/\sqrt{K_D(a, a)}$, multiplied by a unimodular constant. □

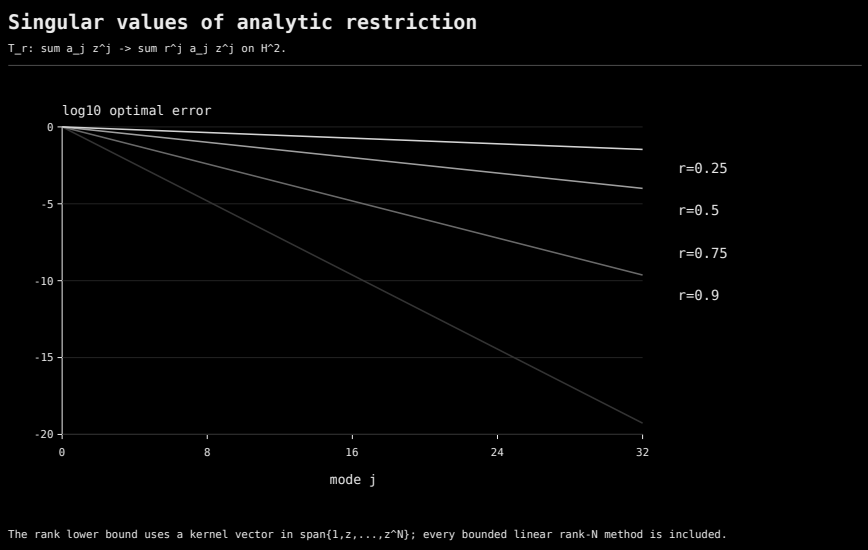


Figure 4.19: Singular values of interior restriction on Hardy space. The first omitted singular value is both an achievable error and the lower bound for every linear representation of the same rank.

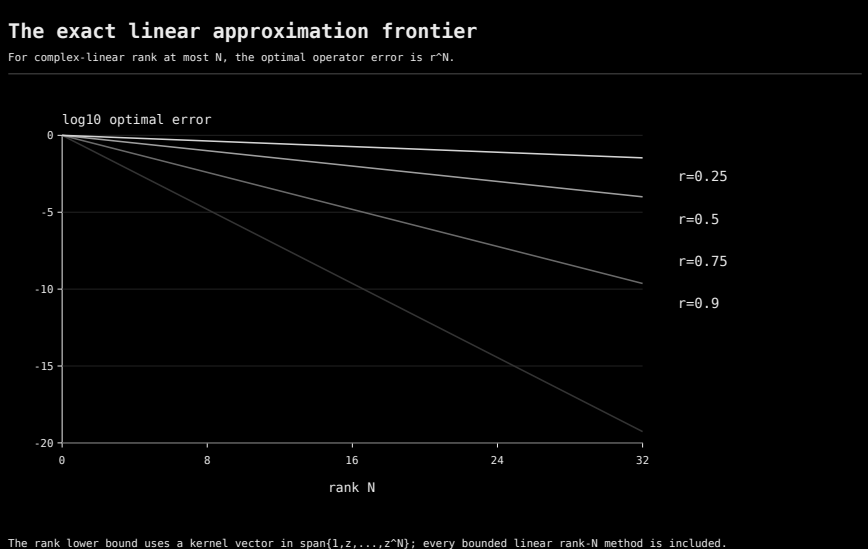


Figure 4.20: The exact rank–error frontier from M35. The horizontal coordinate counts complex coefficients; the error is an operator norm over the full Hardy unit ball.

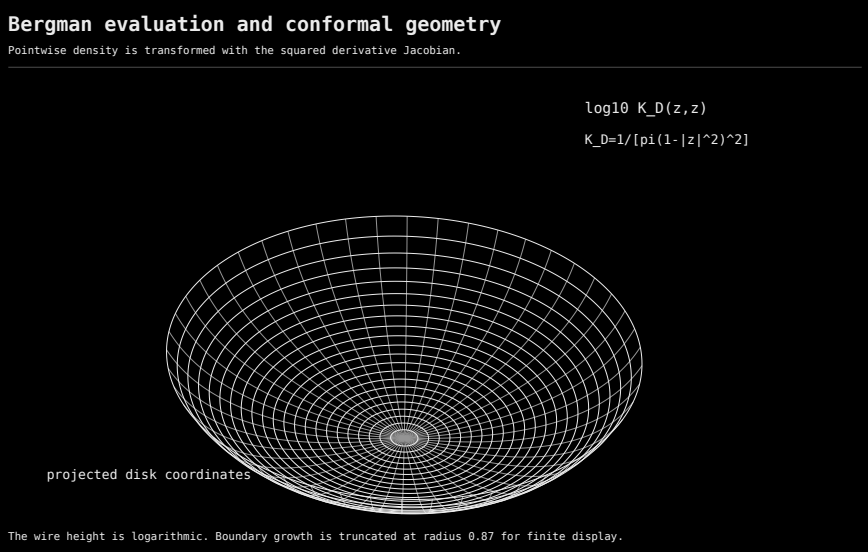


Figure 4.21: The evaluation norm transported by the quadratic conformal map. The Jacobian factor is explicit in the unitary pullback.

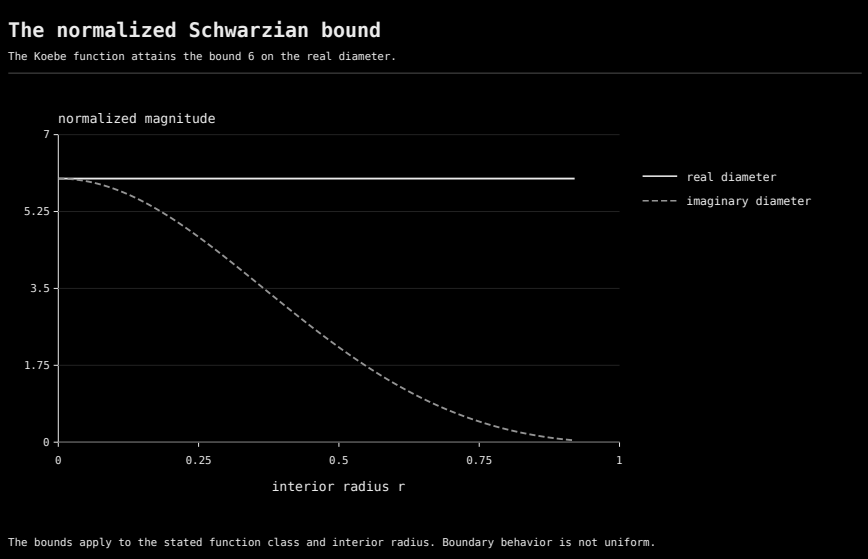


Figure 4.22: The Koebe Schwarzian and its invariantly weighted modulus. The unit-disk boundary is excluded from the plotted domain.

4.7 Completion of the radial-slit extremal argument

Mathematical reference. The extremal remark following source Lemma 4.17, physical p.115. The following derivation uses the variational definition of planar condenser capacity.

The explicit collar bound M29 is sufficient for the invariant distortion theorem. The source also identifies a radial slit as the extremal continuum. A circular rearrangement argument proves that assertion without assuming a nondegenerate limit of annulus uniformizations.

Theorem 4.38 (M38: radial-slit extremality)**.** *Let $K \Subset \mathbb{D}$ be a nondegenerate full continuum containing zero, and put $r = \max_{z\in K} |z|$. Then*

$$\operatorname{Cap}([0,r],\mathbb{D}) \le \operatorname{Cap}(K,\mathbb{D}), \qquad \operatorname{mod}(\mathbb{D}\setminus K) \le \operatorname{mod}(\mathbb{D}\setminus[0,r]).$$

Equality is attained by $K = [0,r]$ and by its rotations. The statement asserts attainment; it does not classify every possible equality case.

Proof. Use the capacity convention

$$\operatorname{Cap}(K,\mathbb{D}) = \inf\left\{ \int_{\mathbb{D}} |\nabla u|^2 dA : u \in C_c^\infty(\mathbb{D}),\ u \geq 1 \text{ on a neighbourhood of } K \right\}.$$

Truncation to $0 \leq u \leq 1$, followed by the usual Sobolev approximation, leaves the infimum unchanged. We first rearrange a nonnegative smooth test function with value one near K.

For each radius ρ, replace the angular function $u(\rho,\theta)$ by its symmetric decreasing rearrangement $u^*(\rho,\theta)$ on the circle, centred at $\theta = 0$. Thus every nontrivial superlevel set of $u^*(\rho,\cdot)$ is a centred arc with the same angular length as the corresponding superlevel set of $u(\rho,\cdot)$. The rearrangement preserves the L^2 norm of each angular slice.

Angular energy. Let v be a smooth nonnegative angular slice and let $\ell(s) = |\{\theta : v(\theta) > s\}|$. At a regular level with a nonempty proper superlevel set, the coarea formula gives

$$-\ell'(s) = \sum_{v(\theta)=s} \frac{1}{|v'(\theta)|}, \qquad \int_{\mathbb{T}} |v'|^2 d\theta = \int \sum_{v(\theta)=s} |v'(\theta)|\, ds.$$

There are at least two boundary points at such a level. Cauchy–Schwarz therefore gives

$$\sum_{v=s} |v'| \geq \frac{4}{-\ell'(s)}.$$

The symmetric decreasing rearrangement has two boundary points with equal derivative magnitudes, and attains the expression on the right. Integrating proves

$$\int |\partial_\theta u^*|^2 d\theta \leq \int |\partial_\theta u|^2 d\theta.$$

Plateaux and critical levels are handled by approximation and lower semicontinuity. This extension is needed because rearrangement of a smooth function need not be smooth.

Radial energy. For two nonnegative angular slices v, w, layer-cake integration gives

$$\int vw = \int_0^\infty \int_0^\infty |\{v > s\} \cap \{w > t\}|\, ds\, dt \leq \int v^* w^*.$$

Indeed, the intersection of two sets has length at most the smaller length, and the two centred arcs attain that intersection length. Since the two individual squared norms are preserved,

$$\|v^* - w^*\|_{L^2(\mathbb{T})} \leq \|v - w\|_{L^2(\mathbb{T})}.$$

Apply this contraction to slices at ρ and $\rho + h$, divide by h^2, and integrate on a compact radial interval away from zero. The difference-quotient characterization of Sobolev derivatives and lower semicontinuity imply

$$\int_\varepsilon^1 \int_{\mathbb{T}} |\partial_\rho u^*|^2 \rho\, d\theta\, d\rho \leq \int_\varepsilon^1 \int_{\mathbb{T}} |\partial_\rho u|^2 \rho\, d\theta\, d\rho.$$

The harmless endpoint intervals are removed before passing to the difference-quotient limit. Letting $\varepsilon \downarrow 0$ gives the radial inequality on the disk. Combining both components in polar coordinates yields

$$\int_{\mathbb{D}} |\nabla u^*|^2 dA = \int_0^1 \int_{\mathbb{T}} \left(|\partial_\rho u^*|^2 + \rho^{-2} |\partial_\theta u^*|^2 \right) \rho\, d\theta\, d\rho \leq \int_{\mathbb{D}} |\nabla u|^2 dA.$$

Admissibility. The continuous image $|K|$ is a connected compact subset of $\mathbb{R}$ containing zero and r; it therefore contains $[0, r]$. Every circle of radius $\rho \in (0, r]$ meets K. Since $u = 1$ on a fixed neighbourhood of the compact set K, the angular slice equals one on an arc of uniformly

positive angular width away from the centre. Rearrangement places that arc around angle zero. The neighbourhood of zero and the endpoint at r supply the remaining local pieces. Consequently u^* equals one on a neighbourhood of the radial segment $[0, r]$ and vanishes near the unit circle. Sobolev approximation, or mollification inside the two separated neighbourhoods, supplies admissible smooth functions with energies approaching that of u^*. Hence

$$\operatorname{Cap}([0, r], \mathbb{D}) \leq \int |\nabla u^*|^2 \leq \int |\nabla u|^2.$$

Take the infimum over u. The annular identity $\operatorname{mod} = 1/\operatorname{Cap}$ reverses the inequality and proves the modulus assertion. A radial slit is itself admissible and attains equality. □

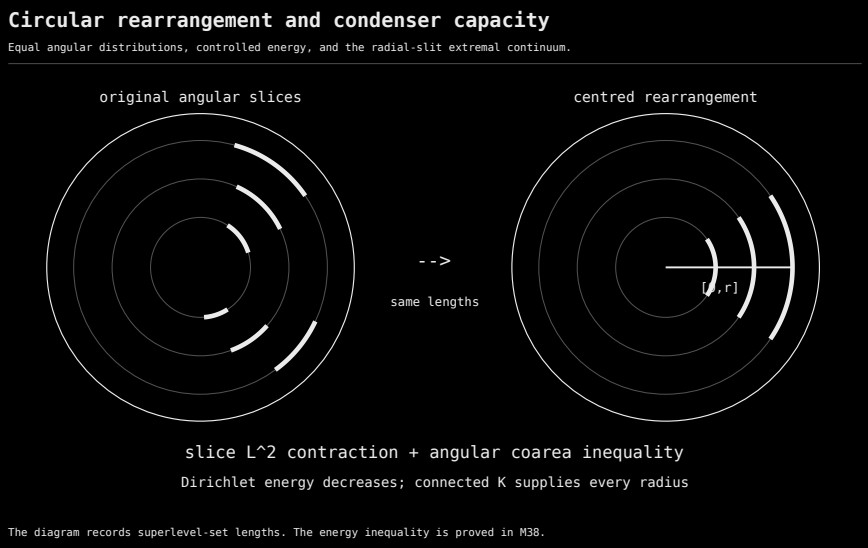


Figure 4.23: Circular rearrangement preserves the distribution of each angular slice. The contraction and coarea estimates in M38 control the radial and angular components of Dirichlet energy.

4.7.1 Optimal-radius interpretation

Let $G(r) = \operatorname{mod}(\mathbb{D} \setminus [0, r])$, and define

$$k_*(\mu) = \sup\{r \in (0, 1) : G(r) \geq \mu\}.$$

The theorem shows that every admissible K with modulus at least μ has maximal radius at most $k_*(\mu)$. Conversely, radial slits realize every admissible radius in this defining set, so no smaller universal upper bound is possible. The explicit estimate M29 gives $k_*(\mu) \leq 1 - e^{-\pi/(4\mu)} < 1$ and hence proves finiteness of the optimal-radius bound without requiring a special-function formula for G.

The supremum formulation suffices for optimality even before continuity is considered. Continuity of condenser capacity for nested radial segments gives the usual inverse-function description and endpoint attainment. That continuity follows from the variational definition: decreasing compact segments are controlled by a fixed neighbourhood test function; increasing segments are treated by weak compactness of the minimizing potentials, with the missing endpoint having zero planar capacity. The zero capacity of a singleton follows from logarithmic radial cutoffs with energies tending to zero. These are properties of the same Dirichlet-capacity model used in the proof.

4.7.2 Scope of the analytical foundations

The proof uses the one-dimensional coarea formula, the Sobolev difference-quotient characterization, and lower semicontinuity of the Dirichlet norm. They are explicit foundational dependencies. Their use does not establish an accepted Lean implementation. The independent exact checker verifies the finite two-arc inequality; the infinite rearrangement and capacity arguments are written analytical proofs.

Exercise 4.15 (Why connectedness matters)**.** Explain why the radial projection of K covers every radius from zero to r, and identify the step that fails for a disconnected set consisting of zero and one other point.

Solution. The modulus map is continuous, so the image of a connected set is an interval. A two-point set projects to only two radii. Intermediate angular slices then need not have value one anywhere, and the rearranged test function need not be admissible for the full radial segment. A two-point set also does not give the ring-domain geometry assumed in the theorem. □

Chapter 5

Constructive examples in algebraic geometry and topology

5.1 Local multiplicity without the vocabulary barrier

5.1.1 MU1: turn a nonlinear expression into a power

Consider $f(z) = z^2(1+z)$ near zero. The equation $f(z) = 0$ has a double zero at zero and another zero at -1. Near zero the factor $1 + z$ never vanishes. Choose the analytic square root agreeing with one at zero and put

$$u = z\sqrt{1+z} = z + \tfrac{1}{2}z^2 - \tfrac{1}{8}z^3 + \cdots .$$

Then $f(z) = u^2$ exactly, not just to low Taylor order. Moreover $u'(0) = 1$, so sufficiently near zero this is an invertible change of coordinate. The expansion helps us recognize the coordinate, but the identity and inverse-function theorem establish the local result.

Why the change matters. For a small nonzero target w, there are two nearby preimages, even though at $w = 0$ their positions coincide. The two positions have not vanished from the count; their combined multiplicity is two. A small positively oriented circle around zero maps to a loop making two turns around zero. Thus polynomial multiplicity and boundary winding are two descriptions of the same local phenomenon. The algebraic-geometric version retains the two-dimensional quotient $\mathbb{C}[u]/(u^2)$, in which u is nonzero but $u^2 = 0$. Counting distinct coordinate values would miss this information.

Exercise 5.1 (MU1a)**.** For $g(z) = (z-2)^3(5+z-2)$, find a local coordinate at $z = 2$ that makes g a cube. State why it is a coordinate and identify the local degree.

Solution. Put $v = (z-2)(5+z-2)^{1/3}$, choosing the local cube root positive at $z = 2$. Its derivative there is $5^{1/3} \neq 0$, so it is locally invertible. The identity $g = v^3$ gives local degree three. No global cube-root branch on the entire punctured plane is claimed.

5.1.2 MU2: test a disk-preserving fraction exactly

Take $b(z) = (z - 1/3)/(1 - z/3)$. The useful identity is

$$|1 - z/3|^2 - |z - 1/3|^2 = \tfrac{8}{9}(1 - |z|^2).$$

For $|z| < 1$, the right side is positive, so the numerator has smaller modulus than the denominator. On the unit circle the moduli are equal. The pole is at three, outside the disk. These three observations establish disk preservation, boundary preservation and absence of an interior pole.

The inverse is $(w+1/3)/(1+w/3)$. The product $B(z)=z^2b(z)$ has degree three, not two: the zero at zero has multiplicity two and the zero at $1/3$ contributes one. Its boundary phase makes three full turns. The derivative at a point is a local feature; the number three is a global fibre count. These should not be confused.

Exercise 5.2 (MU2a)**.** Replace $1/3$ by one in the fraction. What happens to the formula, and why is this not a proper degree-one disk map?

Solution. For $z\neq 1$, $(z-1)/(1-z)=-1$. Cancellation gives a constant function. Its image is on the boundary instead of in the open disk. The admissible hypothesis $|a|<1$ has failed, so continuity of the unreduced formula in the parameter does not preserve degree.

5.2 Euler characteristic as a branch-count ledger

5.2.1 MU3: a genus-two curve from a quadratic equation

Let

$$X:\quad y^2=(x-1)(x-2)(x-3)(x-4)(x-5)(x-6).$$

Projection to x has two preimages for a generic value. At each of $1,\dots,6$ the two sheets meet in a simple branch point. Locally, because the remaining factors are nonzero, the equation is equivalent to $y^2=x-a$ after a nonvanishing coordinate adjustment. The contribution $e-1$ is one at each such point.

Infinity must be checked instead of guessed. Put $u=1/x$ and $v=y/x^3$. Then

$$v^2=\prod_{j=1}^{6}(1-ju).$$

At $u=0$ there are two points, $v=+1$ and $v=-1$. The derivative of $v^2-\prod(1-ju)$ with respect to v is nonzero at both. Each has local coordinate u, and projection has local degree one there. Infinity contributes no branching. The compact smooth surface therefore has

$$\chi(X)=2\chi(S^2)-6=4-6=-2,\qquad 2-2g=-2,\quad g=2.$$

The core proves the Riemann–Hurwitz relation; the undergraduate calculation uses it with an explicitly checked branch list.

Exercise 5.3 (MU3a)**.** For $y^2=\prod_{j=1}^{5}(x-j)$, determine the contribution at infinity and the genus of the smooth compactification.

Solution. Set $u=1/x$ and $v=y/x^3$. The equation becomes $v^2=u\prod_{j=1}^{5}(1-ju)$. Near $(u,v)=(0,0)$, the factor multiplying u is nonzero, so projection to u is locally the square of a coordinate. Infinity is one simple branch point. There are five finite branch points plus one at infinity; the same total six gives genus two. An odd polynomial degree does not mean there is an odd total number of simple branch values on the sphere.

Exercise 5.4 (MU3b)**.** Why can the same six-point count not be applied unchanged when two factors collide, for example $(x-1)^2\prod_{j=3}^{6}(x-j)$?

Solution. The affine curve is singular at $(1,0)$: both partial derivatives of its defining equation vanish. On its normalization, divide by the repeated factor through $\widetilde{y}=y/(x-1)$ to obtain $\widetilde{y}^2=\prod_{j=3}^{6}(x-j)$. The normalized cover has four simple branch points and genus one. Applying the smooth-surface theorem to the singular model without normalization counts the wrong object.

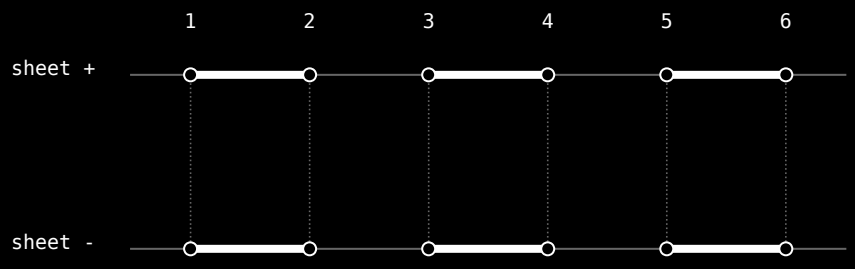


Figure 5.1: A schematic cut construction for the smooth degree-six hyperelliptic example. The six branch values and the two points at infinity are counted in the proof; the displayed cross-sheet lines are a gluing guide, not an embedding of a genus-two surface.

5.3 Boundary information can count an interior

5.3.1 MU4: compute a winding by continuation of angle

For $\gamma(t) = 2e^{3it}$, $0 \le t \le 2\pi$, a continuous argument is $3t$. The total change is 6π, so the winding is three. Reversing the path changes the sign. For a general closed loop avoiding zero, locally choose a continuous angle and continue it; the final angle differs from the initial one by an integer multiple of 2π.

The loop can deform without changing this integer as long as it never passes through zero. For instance $2e^{it} + c$, $c \in [0, 1]$, always avoids zero because its modulus is at least $2 - c \ge 1$. It has winding one for every such c. At $c = 2$ the exclusion fails and the loop touches zero. No amount of numerical angle unwrapping can repair the missing hypothesis at that parameter.

5.3.2 MU5: count quartic roots in two regions

Let $p(z) = z^4 + 5z + 1$. On $|z| = 1$, the term $5z$ has modulus five, whereas $|z^4 + 1| \le 2$. The homotopy $5z + t(z^4 + 1)$ is nonzero on the circle for all $t \in [0, 1]$. Consequently p has one zero inside, counted with multiplicity. On $|z| = 2$, the leading term has modulus sixteen and $|5z + 1| \le 11$. The same argument gives four interior zeros there. Thus exactly three roots lie in $1 < |z| < 2$, and no roots lie on either circle.

This uses the argument principle through a strict triangle inequality. Finding decimal approximations to four roots would answer a different question and would need error bounds to certify these strict regions. For one quartic, a numerical solver may still be the best constructor of approximate values; the two inequalities are a small independent checker.

Exercise 5.5 (MU5a)**.** Show that $p_t(z) = z^4 + 5z + t$ has exactly one root in the unit disk for every complex $|t| < 4$. Explain the strict boundary of this sufficient condition.

Solution. On $|z| = 1$, $|z^4 + t| \le 1 + |t| < 5 = |5z|$. The nonvanishing homotopy and the degree of $5z$ give one. At $|t| = 4$ the bound is no longer strict, so this proof alone stops. It does not assert that every such endpoint parameter has a boundary root; a sufficient inequality failing is weaker than the conclusion failing.

Exercise 5.6 (MU5b)**.** A continuous real map can have a target in its image but zero boundary degree. Give the reason using $F(x, y) = (x^2, y)$ near a small positive target $(\varepsilon, 0)$.

Solution. The target has two preimages $(\pm\sqrt{\varepsilon}, 0)$. The Jacobian determinant is $2x$, positive at one and negative at the other. The local indices sum to zero. Holomorphic nonconstant maps preserve orientation locally away from branching, so their local multiplicities do not cancel this way. The converse of "nonzero degree implies existence" is false for arbitrary continuous maps.

5.4 Quotient algebras: an algebraic-geometric object you can compute

5.4.1 MU6: multiplication by a coordinate

For $p(z) = z^3 - 2$, consider the three-dimensional vector space $A = \mathbb{C}[z]/(p)$ with basis $1, z, z^2$. Reducing z^3 to two gives the multiplication matrix

$$M_z = \begin{pmatrix} 0 & 0 & 2 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \qquad M_z^3 = 2I.$$

For any polynomial observable r, multiplication by $r(z)$ is represented by $r(M_z)$. The eigenvalues of M_z are the roots of p, but traces can be calculated without assigning or following individual roots.

Exercise 5.7 (MU6a)**.** Compute $\operatorname{tr} M_z$, $\operatorname{tr} M_z^2$, $\operatorname{tr} M_z^3$ and interpret them as power sums of the roots. Then explain what changes for the quotient by z^3.

Solution. The first two traces are zero and the third is six because $M_z^3 = 2I$. Thus the sums of first, second and third powers of the roots are $0, 0, 6$. For the quotient by z^3, multiplication is a nilpotent three-by-three shift. The only root position is zero, but the quotient still has dimension three. This dimension records multiplicity that a set of distinct roots cannot record. It is the finite local-ring idea underlying ramification, not just an arbitrary matrix analogy.

5.5 Following a root is a path problem

5.5.1 MU7: a three-sheet lift

Solve $z^3 = w$ along $w(t) = e^{it}$ starting with $z(0) = 1$. Continuity selects $z(t) = e^{it/3}$. At $t = 2\pi$ the endpoint is $e^{2\pi i/3}$, not one. A second loop reaches the other nontrivial cube root, and a third closes the lift. A routine that returns the principal root independently at each sample may jump at its branch cut; that is not the continued solution.

Exercise 5.8 (MU7a)**.** For a loop with winding m in $\mathbb{C}^*$, when does its lift under $z \mapsto z^d$ close? Prove the criterion by a continuous argument, not by a picture.

Solution. Write the loop as $\rho(t)e^{i\theta(t)}$ with a continuous argument along its interval parameter. Its lifted root has argument $(\theta(t) + 2\pi k)/d$, where k is fixed by the initial root. The endpoint argument changes by $2\pi m/d$. The endpoint equals the starting root exactly when m/d is an integer, or $d \mid m$. This is the subgroup $d\mathbb{Z}$ in the fundamental group of the punctured plane.

5.5.2 MU8: permutations do not generally commute

For three sheets label the two elementary swaps $a = (12)$ and $b = (23)$. Applying b then a sends $1 \mapsto 2 \mapsto 3 \mapsto 1$ under their product convention, producing a three-cycle; reversing the order gives its inverse. The generated group has six elements: identity, three transpositions and two three-cycles. A regular cover records all six permutation states, while tracking one root records only its orbit of size three. The forgotten information is its stabilizer.

Exercise 5.9 (MU8a)**.** Use the convention $(ab)(j) = a(b(j))$ to write ab and ba explicitly and show that $aba = (13)$.

Solution. Direct evaluation gives $ab = (123)$, $ba = (132)$ and $aba = (13)$. These calculations specify the convention needed by a permutation program. Multiplying symbols from left to right without deciding what they mean can invert a monodromy record.

5.6 Composition: a formula is not its domain

5.6.1 MU9: a two-stage pullback

Let $U_0 = \{|z| < 1\}$, $U_1 = \{|w| < 1\}$, $U_2 = \{|v| < 1/2\}$, and let $\phi_0(z) = 2z$, $\phi_1(w) = 3w$. The terminal inverse branch gives $w = v/3$, then $z = v/6$. Thus the correct domain mapped onto U_2 is $|z| < 1/12$. The formula $6z$ exists algebraically everywhere; its existence is not what the telescope theorem asserts.

The intermediate admissibility domain for merely reaching U_1 is $|z| < 1/2$. The terminal domain is contained in it with a strict margin. These two nested sets encode separate obligations. An input at $z = 1/5$ passes the first stage but does not land in U_2 after both stages.

Exercise 5.10 (MU9a)**.** Replace the second target by $|v - 1/4| < 1/8$, leaving both stage maps unchanged. Determine the initial domain and check that it lies in U_0 and that the intermediate values lie in U_1.

Solution. Pulling back through $v = 6z$ gives $|z - 1/24| < 1/48$. Its maximal modulus is $1/16 < 1$. The intermediate variable $w = 2z$ satisfies $|w - 1/12| < 1/24$, hence $|w| < 1/8 < 1$. Each claim is a numerical domain check, not an assumption based on the composite expression.

5.6.2 MU10: the error budget of an inverse iteration

Suppose a stage contracts errors by $q = 1/2$ and has a fresh evaluation discrepancy at most $\eta = 1/1000$. Starting from exact input, after n steps

$$e_n \le \frac{1}{500}(1 - 2^{-n}).$$

After ten steps the bound is $1023/512000$, below $1/500$. More iterations do not force this worst-case bound to zero, because new error is inserted at every stage. The qualitative word "contractive" is not a substitute for this calculation.

Exercise 5.11 (MU10a)**.** A claimed implementation replaces the bound by ηq^n. Give a two-step counterexample consistent with the recurrence $e_{n+1} \le q e_n + \eta$.

Solution. Choose actual errors $e_0 = 0$, $e_1 = \eta$, $e_2 = q\eta + \eta$. These satisfy the recurrence with equality. The claimed bound at two steps is ηq^2, smaller than η when $0 < q < 1$, so it fails. The correct geometric sum has contributions of different ages.

5.7 MG1: finite fibres as algebras, including nilpotents

Let $p(z) = \prod_{j=1}^{r}(z-a_j)^{e_j}$ with distinct a_j and degree $d = \sum e_j$. The Chinese remainder theorem gives

$$\mathbb{C}[z]/(p) \cong \bigoplus_{j=1}^{r} \mathbb{C}[z]/((z-a_j)^{e_j}) \cong \bigoplus_{j=1}^{r} \mathbb{C}[\varepsilon]/(\varepsilon^{e_j}).$$

To prove the first isomorphism, distinct linear factors generate comaximal ideals; Bezout identities construct mutually orthogonal idempotents projecting onto each summand. The second map uses $\varepsilon = z - a_j$. Each summand has basis $1, \varepsilon, \ldots, \varepsilon^{e_j-1}$, so its vector dimension is e_j. This is the algebraic reason fibre multiplicities add to degree.

Multiplication by a polynomial h on the jth local summand is upper triangular with diagonal value $h(a_j)$ repeated e_j times. Consequently

$$\operatorname{tr} M_h = \sum_{j=1}^{r} e_j h(a_j), \qquad \det M_h = \prod_{j=1}^{r} h(a_j)^{e_j}.$$

The nilpotent upper diagonals depend on derivatives of h, but trace and determinant retain multiplicity without choosing a basis of root eigenvectors. At colliding roots, a diagonal eigenbasis may fail while this algebra still has dimension d.

Exercise 5.12 (MG1a)**.** For $p(z) = (z-1)^2(z+2)$, find $\operatorname{tr} M_z$, $\operatorname{tr} M_{z^2}$ and $\det M_{z-3}$. Explain why a sum over distinct roots is wrong.

Solution. The local dimensions are two and one. Thus the traces are $2(1) + (-2) = 0$ and $2(1)^2 + (-2)^2 = 6$, while the determinant is $(1-3)^2(-2-3) = -20$. Distinct-root summation would give -1 and five for the two traces. It discards the doubled local algebra at one.

5.8 MG2: Hardy interpolation detects coalescing data

Use $H^2(\mathbb{D}) = \{f(z) = \sum_{n\geq 0} c_n z^n : \sum |c_n|^2 < \infty\}$, with the inner product linear in the first variable. Evaluation is bounded:

$$|f(a)| \leq \|f\|_{H^2} \left(\sum_{n\geq 0} |a|^{2n}\right)^{1/2} = \frac{\|f\|_{H^2}}{\sqrt{1-|a|^2}}.$$

The representing kernel is $k_a(z) = 1/(1-\bar{a}z)$. The coefficient proof gives $\langle f, k_a\rangle = f(a)$ directly. Derivative evaluations are also bounded at interior points because the appropriate polynomially weighted geometric sums converge.

Let a finite Blaschke product B have zeros a_j of multiplicities e_j. Then BH^2 consists of those f with $f^{(m)}(a_j) = 0$ for $0 \leq m < e_j$. One direction is divisibility. In the reverse direction, divide successively by the elementary factors; removal of the zeros gives a holomorphic quotient, and the boundary modulus-one relation preserves the H^2 norm. Thus the finite-dimensional orthogonal complement $H^2 \ominus BH^2$ is spanned by the representing kernels of these jet evaluations. Hermite interpolation by polynomials shows the $d = \sum e_j$ functionals are independent, so this complement has dimension d.

5.8.1 A quantified collision of two interpolation sites

For real sites $a, -a$, $0 < a < 1$, the kernel Gram matrix is

$$G = \begin{pmatrix} (1-a^2)^{-1} & (1+a^2)^{-1} \\ (1+a^2)^{-1} & (1-a^2)^{-1} \end{pmatrix}.$$

The symmetric and antisymmetric eigenvectors give eigenvalues

$$\lambda_+ = \frac{2}{1-a^4}, \qquad \lambda_- = \frac{2a^2}{1-a^4}.$$

Hence $\kappa_2(G) = a^{-2}$. As the sites coalesce, two value samples become nearly redundant. The rescaled antisymmetric functional $(f(a) - f(-a))/(2a)$ tends to $f'(0)$, revealing the jet that should replace the duplicate value. This is a functional-analytic explanation of the local nilpotent fibre in MG1, not an assertion that every ill-conditioned matrix is a branch point.

Exercise 5.13 (MG2a)**.** Find the minimal H^2 norm of a function with $f(a) = 1$ and $f(-a) = -1$, and identify the divergence as $a \downarrow 0$.

Solution. The minimum-norm interpolant lies in $\operatorname{span}\{k_a, k_{-a}\}$. Solving $Gc = (1, -1)^T$ gives coefficient vector $c = \lambda_-^{-1}(1, -1)^T$. Its squared norm is $(1, -1)G^{-1}(1, -1)^T = 2/\lambda_- = (1 - a^4)/a^2$. Thus the norm diverges like $1/a$. The interpolation data demand a slope of order $1/a$, so this is a genuine data-conditioning effect, not merely a poor numerical solver.

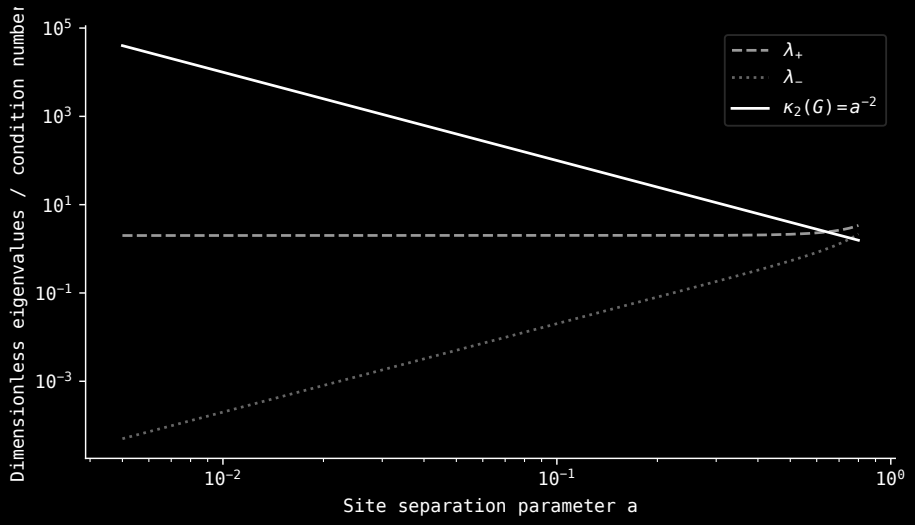


Figure 5.2: The exact two-site Hardy Gram eigenvalues and condition number under coalescence. The graph uses positive a only. A jet basis resolves the limiting information type, but it does not make incompatible fixed value data uniformly bounded.

5.9 MG3: an explicit differential basis on a double cover

Let p be squarefree of degree $2g+2$, $g \geq 1$, and let X be the smooth compactification of $y^2 = p(x)$. By the fully counted Riemann–Hurwitz formula, X has genus g. The differentials

$$\omega_k = \frac{x^k\,dx}{y}, \qquad 0 \leq k \leq g-1,$$

are holomorphic. At a simple finite branch point a, a local coordinate t satisfies $x - a = t^2$ and $y = t\,u(t)$ with $u(0) \neq 0$; therefore $dx/y = 2dt/u(t)$ is regular. Away from branch points there is no problem. At infinity use $u = 1/x$, $v = y/x^{g+1}$, obtaining

$$\frac{x^k dx}{y} = -\frac{u^{g-1-k}}{v}\,du,$$

regular exactly for the displayed range. The g forms are linearly independent because a polynomial relation in x would vanish on an open set. The theorem $\dim H^0(X, K_X) = g$ then shows they form a basis; that dimension theorem is an explicitly imported Riemann-surface result, not established by this local computation alone.

For odd degree $2g+1$, infinity is one branch point. Set $x = t^{-2}$ locally; then $y = t^{-(2g+1)}u(t)$ with nonvanishing u, so $x^k dx/y$ is a constant times $t^{2g-2k-2}dt/u(t)$. The same range is regular. The coordinate at infinity changes, but the basis dimension agrees.

Exercise 5.14 (MG3a)**.** For a smooth curve $y^2 = x^8 - 1$, list a holomorphic differential basis and compute the total zero order of dx/y.

Solution. The genus is three. A basis is $dx/y, x\,dx/y, x^2dx/y$. The first differential has order two at each of the two infinity points and no zero or pole at a finite branch point. Its total zero order is four, matching $\deg K_X = 2g - 2 = 4$. This is an explicit consistency check between local coordinates and a global divisor degree.

5.10 MG4: a finite cover defines an averaging adjoint

Let $f : X \to Y$ be a degree-d cover away from its critical values. Choose a finite measure ν on Y giving those critical values measure zero. Define μ on X by uniformly averaging over the d distinct preimages:

$$\int_X h\,d\mu = \int_Y \frac{1}{d} \sum_{f(x)=y} h(x)\,d\nu(y).$$

For measurable functions let $Ug = g \circ f$ and $Ph(y) = d^{-1}\sum_{f(x)=y} h(x)$ almost everywhere. Then

$$\|Ug\|_{L^2(\mu)} = \|g\|_{L^2(\nu)}, \quad \langle Ug, h\rangle = \langle g, Ph\rangle, \quad PU = I.$$

The identities follow by substituting the defining measure formula. Jensen's finite-sum inequality gives $\|P\| \le 1$. Consequently UP is the orthogonal projection onto the fibre-constant functions. This is a function-space meaning of passing from sheet-level data to the base.

The measure was chosen to make this statement true. Ordinary area measure on the source has Jacobian weights, and Dirichlet energy instead scales by degree as in M10. Reusing the same symbol L^2 while changing the measure would invalidate the adjoint formula.

Exercise 5.15 (MG4a)**.** For $z \mapsto z^3$ on the unit circle with normalized angular measure, compute $P(z^n)$, $n \in \mathbb{Z}$.

Solution. The three preimages of w are $\zeta^j z$ with $z^3 = w$ and $\zeta^3 = 1$. Their average is $z^n(1 + \zeta^n + \zeta^{2n})/3$. It equals $w^{n/3}$ when three divides n and zero otherwise. Thus averaging removes Fourier modes that do not descend to the base. This is a roots-of-unity identity, not a claim that arbitrary signal downsampling includes an ideal antialiasing filter.

5.11 MG5: the complete data of a branched cover

For a punctured target with chosen base point, loops around branch values induce permutations $\sigma_1, \ldots, \sigma_r \in S_d$. A connected cover requires a transitive generated action. Relabelling the fibre conjugates every permutation by the same element. The cycle lengths of σ_i are the local degrees over the ith value. The total ramification over it is

$$d - \#\,\mathrm{cycles}(\sigma_i).$$

The product constraint depends on the target: on the punctured sphere, the ordered product including the loop around infinity is the identity. On a punctured disk, its inverse describes the boundary monodromy under a fixed orientation convention. Matching only the list of cycle lengths is not sufficient for equivalence; the relative arrangement of the permutations also matters.

Exercise 5.16 (MG5a)**.** A transitive degree-four cover of the sphere has local permutations with cycle types (4), (2, 2) and (2, 1, 1). Compute its genus assuming a consistent product relation exists. What does the calculation not prove?

Solution. The defects are three, two and one, totaling six. Hence $2-2g = 8-6 = 2$ and $g = 0$. The arithmetic does not construct permutations with product one or prove transitivity. In this example a dihedral realization exists, but a genus ledger alone is not the existence theorem. A passport must include and verify the actual permutations.

5.12 MG6: a regular cover much smaller than the factorial bound

For $d \geq 3$ define polynomials $P_0(x) = 2$, $P_1(x) = x$ and $P_{n+1} = xP_n - P_{n-1}$. Induction using $x = t + t^{-1}$ gives

$$P_d(t + t^{-1}) = t^d + t^{-d}.$$

The rational map $g(t) = t^d + t^{-d}$ has degree $2d$. Its deck transformations are $t \mapsto \zeta^k t$ and $t \mapsto \zeta^k/t$, where $\zeta = e^{2\pi i/d}$. They form a dihedral group of order $2d$: rotations compose by addition, reflection has square one and conjugates a rotation to its inverse. There are $2d$ transformations, equal to the map degree, so this is a regular cover. It factors through $x = t+t^{-1}$, a degree-two quotient, followed by P_d, degree d.

The stabilizer of x is the reflection subgroup. For $d \geq 3$ its intersection with all conjugates is trivial: a nontrivial common element would have to be the same reflection in every conjugate, impossible because the conjugates include a distinct reflection. Thus the induced action on the d cosets is faithful, and the dihedral cover is the minimal normal closure of the polynomial cover. For $d = 2$ that argument fails: the degree-two polynomial is already Galois, so the four-sheet construction is regular but not minimal. The exception is part of the theorem.

Exercise 5.17 (MG6a)**.** Compute the normal-closure degree for P_5 and compare it with the general factorial bound. Check the genus of its regular closure from inertia orders $2, 2, 5$.

Solution. The degree is ten, compared with the universal upper bound $5! = 120$. On the ten-sheet regular cover the ramification contribution of inertia order e is $10(1 - 1/e)$. The three contributions are five, five and eight, totaling eighteen. Riemann–Hurwitz gives $\chi = 20 - 18 = 2$, hence genus zero. The bound $d!$ is a worst-case possibility, not an instruction to allocate $d!$ states for every polynomial.

5.13 MG7–MG8: analytic operators and local inverse certificates

If $\phi : U \to V$ is holomorphic, pullback $C_\phi h = h \circ \phi$ maps $H^\infty(V)$ to $H^\infty(U)$ with norm at most one. For composable maps the operator order reverses: $C_{\psi\circ\phi} = C_\phi C_\psi$. The theorem is almost immediate from a supremum, but the domains are its substance. The source telescope constructs exactly those intermediate domains on which all factors are defined.

For a quantitative local inverse of $z^2 = w$, choose a center $c \neq 0$ and radius $r > 0$. If w ranges over a compact parameter cell Q and

$$\sup_{w\in Q} |c^2 - w| + r^2 < 2|c|r,$$

then $z^2 - w = 2c(z-c) + (c^2 - w) + (z-c)^2$ has exactly one zero in $|z-c| < r$ for every $w \in Q$. This is the core argument-principle theorem applied to a parameter-uniform linear domination. The selected root varies analytically on any open parameter domain on which the same strict inequality holds. The nonzero derivative and the uniqueness of the root patch local implicit solutions together.

Exercise 5.18 (MG8a)**.** Certify a square-root branch near $w = 1$ using $c = 1$, $r = 1/4$ and $|w - 1| \le 1/4$. Give a derivative bound for the certified root.

Solution. The remainder bound is $1/4 + 1/16 = 5/16 < 1/2 = 2|c|r$. Hence there is exactly one root in the disk. Every such root satisfies $|z| \ge 3/4$, so $|dz/dw| = 1/(2|z|) \le 2/3$. The strict inequality persists on a slightly larger open parameter disk. The branch is selected by the neighborhood of one, not by a global sign rule on $\mathbb{C}^*$.

5.14 MP1: a discriminant margin gives a separation bound

Let $p(z) = \prod_{j=1}^{d}(z - r_j)$ be monic of degree $d \ge 2$. Assume all roots satisfy $|r_j| \le R$, $R > 0$, and

$$|\operatorname{Disc}(p)| = \prod_{i<j} |r_i - r_j|^2 \ge \delta > 0.$$

Write $N = d(d-1)/2$. Then every root pair satisfies

$$|r_i - r_j| \ge s := \frac{\sqrt{\delta}}{(2R)^{N-1}}.$$

Indeed, isolate the chosen pair in the discriminant product and bound every other separation by $2R$. The result is often pessimistic, but it is uniform and computable from declared data. It proves simplicity and gives $|p'(r_i)| \ge s^{d-1}$.

5.14.1 A coefficient-perturbation certificate

Suppose $\widetilde{p} = p + \sum_{k=0}^{d-1} \eta_k z^k$ with $|\eta_k| \le \eta$. Choose $0 < r < s/2$. On $|z - r_i| = r$,

$$|p(z)| \ge r(s-r)^{d-1}, \qquad |\widetilde{p}(z) - p(z)| \le \eta \sum_{k=0}^{d-1} (R+r)^k.$$

Thus the strict inequality

$$\eta \sum_{k=0}^{d-1} (R+r)^k < r(s-r)^{d-1}$$

certifies exactly one perturbed root in each of the disjoint root disks. The count is multiplicity-sensitive and global: there are d disks and a degree-d polynomial, so no further roots remain. The assumptions include a certified root bound and discriminant lower bound; estimating either from a root plot is not the same theorem.

Exercise 5.19 (MP1a)**.** For $p_w(z) = z^3 - 3z - w$, $|w| \le 1$, derive a uniform root bound and a separation lower bound. Compare that bound with the actual minimum separation at $w = 0$.

Solution. Cauchy's monic root bound gives $R = 4$. The discriminant is $108 - 27w^2$, whose modulus is at least $108 - 27 = 81$. With $N = 3$, $s \ge 9/64$. At $w = 0$ the roots are $-\sqrt{3}, 0, \sqrt{3}$, with minimum separation $\sqrt{3}$. The global estimate is about twelve times smaller. This difference quantifies the cost of a uniform invariant-based guarantee; it is not a failed proof or evidence that pointwise estimates are uniformly safe.

Research extension. A potential contribution is a sharper certificate combining local resultants, adaptive root clusters and a fixed coefficient uncertainty model. The required comparison is the number of certified continuation cells and bit complexity at a common error tolerance. No such new superiority result is asserted by the elementary discriminant estimate above.

5.15 MP2: a branch collision has a logarithmic period scale

Consider the elliptic integral

$$K(k) = \int_0^{\pi/2} \frac{d\theta}{\sqrt{1 - k^2 \sin^2 \theta}}, \qquad \varepsilon = \sqrt{1 - k^2} \in (0, 1/2].$$

It is a period of the smooth double-cover curve associated with the four branch points of $(1 - x^2)(1 - k^2 x^2)$. As $k \uparrow 1$, branch points collide. We now derive a uniform logarithmic bound without using the full asymptotic expansion.

Set $t = \cos\theta$:

$$K(k) = \int_0^1 \frac{dt}{\sqrt{(1 - t^2)(\varepsilon^2 + (1 - \varepsilon^2)t^2)}}.$$

For $0 \le t \le 1/2$, put $C(t) = (\varepsilon^2 + t^2)^{-1/2}$. The integrand is at least $C(t)$. For $0 \le u \le 1/4$, $(1 - u)^{-1/2} \le 1 + u$; squaring verifies this since $(1 + u)^2(1 - u) - 1 = u - u^2 - u^3 \ge 0$. Applying the inequality to $u = t^2$ and to $u = \varepsilon^2 t^2/(\varepsilon^2 + t^2) \le t^2$ bounds the integrand by

$$(1 + t^2)^2 C(t) \le (1 + \tfrac{9}{4} t^2) C(t).$$

The resulting excess integral on $[0, 1/2]$ is at most $\frac{9}{4} \int_0^{1/2} t\, dt = 9/32$. On $[1/2, 1]$, the integrand is at most $2/(\sqrt{3}\, t\sqrt{1 - t^2})$, whose integral is $(2/\sqrt{3}) \log(2 + \sqrt{3})$. Therefore

$$\operatorname{arsinh} \frac{1}{2\varepsilon} \le K(k) \le \operatorname{arsinh} \frac{1}{2\varepsilon} + \frac{9}{32} + \frac{2}{\sqrt{3}} \log(2 + \sqrt{3}).$$

Since $\operatorname{arsinh}(1/(2\varepsilon)) - \log(1/\varepsilon)$ is uniformly bounded and nonnegative for this range, we obtain $K(k) = \log(1/\varepsilon) + O(1)$ with an explicit uniform constant. In particular the leading logarithmic coefficient is one.

This calculation identifies the mechanism and a certified scale, but does not derive the sharper constant $\log 4$ in the classical expansion. It also applies to an ideal conservative period model; damping and forcing remove the conserved algebraic energy curve unless additional structure is supplied.

Exercise 5.20 (MP2a)**.** Why is a uniform absolute error bound $|K(k) - A(k)| \le C$ compatible with a relative error tending to zero as $k \uparrow 1$? Why is it not a uniform derivative-error bound?

Solution. The denominator $K(k)$ diverges like $\log(1/\varepsilon)$, so $C/K(k) \to 0$. A bounded function difference can still have a rapidly oscillating or singular derivative; differentiating a big-$O(1)$ statement is invalid without a differentiated integral estimate. The norm and the requested observable must be specified separately.

5.16 MP3: Fredholm index survives norm perturbations, not arbitrary limits

For $0 < a < 1$ consider the inner multiplier M_{b_a} on H^2, with $b_a(z) = (z - a)/(1 - az)$. Its range has codimension one and its kernel is zero, so its Fredholm index is -1. For each fixed interior z, $b_a(z) \to -1$ as $a \uparrow 1$.

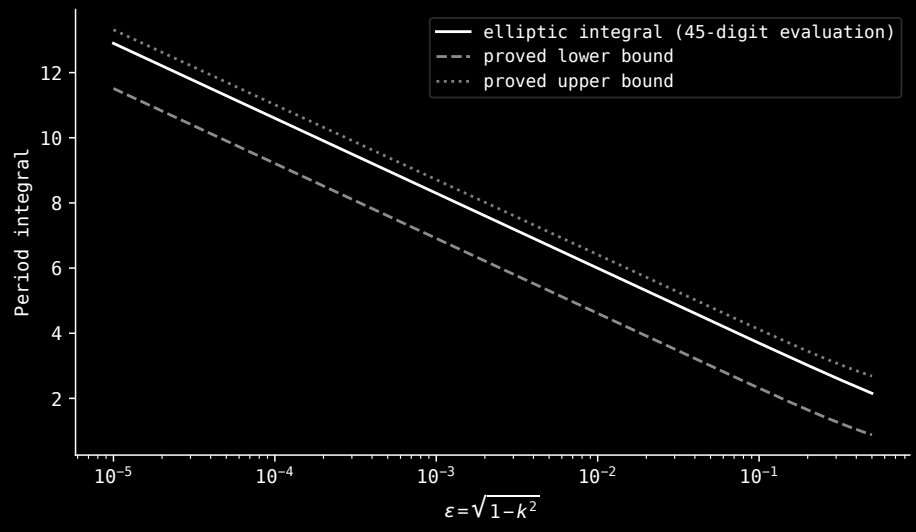


Figure 5.3: The period integral and the uniform analytic bounds from MP2. The central curve is a high-precision numerical display; the inequality is proved by integral comparison and is not inferred from the plotted samples.

The convergence is strong as an operator, but not in operator norm. Since $\langle b_a, 1\rangle = b_a(0) = -a$ and $\|b_a\| = 1$,

$$\|b_a + 1\|_{H^2}^2 = 2(1-a) \to 0.$$

For every polynomial h, multiplication by the bounded polynomial gives $\|(b_a+1)h\|_{H^2} \leq \|h\|_\infty \|b_a+1\|_{H^2} \to 0$. Uniform operator boundedness by two and density of polynomials extend the convergence to every $h \in H^2$. On the other hand

$$\|M_{b_a} + I\| = \|b_a + 1\|_{H^\infty} = 2,$$

because values approaching $z = 1$ approach two and the triangle inequality gives the reverse bound. The limiting operator $-I$ has index zero. This does not contradict local constancy of Fredholm index in operator norm: that topology's hypothesis fails.

5.16.1 The disappearing one-dimensional model space

The normalized kernel $e_a(z) = \sqrt{1-a^2}/(1-az)$ has norm one. Every fixed Taylor coefficient tends to zero; boundedness and density give $e_a \rightharpoonup 0$. Hence the rank-one projection $P_a h = \langle h, e_a\rangle e_a$ tends strongly to zero, while $\|P_a\| = 1$. A fixed dimension at every finite parameter does not imply that a strong limit has the same rank. This is a concrete functional-analytic model of a zero escaping through the boundary of a proper-map parameter space.

Exercise 5.21 (MP3a)**.** What is the boundary-degree behavior of b_a for $a < 1$, and why may it not be computed by first setting $a = 1$ in the cancelled rational expression?

Solution. The interior family has degree one and boundary winding one. At $a = 1$ the numerator and denominator cancel away from $z = 1$, giving -1, while the original boundary expression has a simultaneous zero and pole at that point. The continuous nonvanishing boundary homotopy required to preserve winding fails at the limiting parameter. Operator topology and boundary topology identify the same excluded degeneration through different, explicitly checked hypotheses.

5.17 MP4: a parameter-uniform continuation patch

Let $p(z, w)$ be a polynomial in z with analytic coefficients in w. For a parameter cell Q and a fixed contour Γ, suppose

$$\inf_{w\in Q, z\in\Gamma} |p(z, w)| \geq m > 0.$$

Then no root crosses Γ anywhere in Q. If Q is connected, the argument-principle root count is constant. A finite set of sampled w does not prove this lower bound; one can instead certify a convex coefficient family or a tensor-product Bernstein enclosure as developed in the CSE research branch.

A root count of one inside a simply connected isolating disk gives a single local branch when the root is simple. A root count of two does not define two global labels. For $z^2 = w$ over an annulus around zero, every fibre has two simple roots and their separation is bounded away from zero on the annulus, yet a loop exchanges them. Uniform local separation is not global triviality of the covering. The additional obstruction is the monodromy representation of the parameter domain.

Exercise 5.22 (MP4a)**.** Give a domain on which $z^2 = w$ has uniformly separated roots but no global continuous single-root selection, and prove both assertions.

Solution. Take $1 < |w| < 2$. The roots differ by $2\sqrt{w}$ in modulus, so separation exceeds two. If a continuous root selection existed, restrict it to $w = re^{it}$ with $1 < r < 2$. The winding of its square would be twice an integer, whereas the base loop has winding one. Contradiction. The obstruction is topological, not small root separation.

5.18 MP5: the normal core and changes of branch generators

A degree-d connected unbranched cover corresponds to a subgroup H of index d in a fundamental group Π. Its normal core

$$N = \bigcap_{g\in\Pi} gHg^{-1}$$

is the kernel of the coset action $\Pi \to S_d$. This reduces the apparently infinite intersection to a finite permutation representation: Π/N is the generated monodromy group. The degree of the minimal regular closure is $|\Pi/N|$, bounded by $d!$ but often smaller.

A Hurwitz move on adjacent branch generators has the form

$$(\sigma, \tau) \longmapsto (\sigma\tau\sigma^{-1}, \sigma).$$

It preserves the ordered product, the generated subgroup and the unordered list of conjugacy classes. The first statement is multiplication; the second follows because each old generator can be recovered from the new pair; the third is conjugation and interchange. These are useful finite checks. They do not settle whether two arbitrary passports lie in the same Hurwitz orbit: that is a distinct orbit problem with its own allowed moves and target topology.

Exercise 5.23 (MP5a)**.** For the degree-five dihedral cover in the graduate branch, why is the normal-closure degree ten instead of 120? Which additional input would be needed before claiming that a computed group proves the intended analytic cover is correct?

Solution. The explicit rational model realizes the faithful dihedral monodromy of order ten. The number 120 is only the size of S_5, not the subgroup actually generated. To identify another numerically presented cover, one must certify the critical values, the chosen loops, their root continuation and the resulting permutations; then one must match the cover's equivalence data. An exact group computation on guessed permutations proves only the group of the guessed data.

5.19 MP6: telescope certificates with nonuniform constants

Suppose inverse stages ψ_j have Lipschitz constants q_j on explicitly certified domains. An implemented stage incurs error at most η_j. Repeated substitution in $e_{j+1} \le q_j e_j + \eta_j$ gives

$$e_n \le \left(\prod_{j=0}^{n-1} q_j \right) e_0 + \sum_{k=0}^{n-1} \eta_k \prod_{j=k+1}^{n-1} q_j.$$

The empty product is one. This formula distinguishes early errors, which undergo more later contraction, from late errors. A uniform bound $q_j \le q < 1$ recovers the familiar geometric sum, but can be wasteful if only a few stages are poorly conditioned.

To make the bound a valid program certificate, pair it with a containment inequality at each stage: the approximate image plus its error radius must lie within the next stage's certified domain. It is invalid to use a derivative bound evaluated on a region from which the error ball has already escaped. The complex-analytic derivative estimates, topological domain construction and arithmetic error recurrence are three separate proof components.

Exercise 5.24 (MP6a)**.** For three stages with $(q_0, q_1, q_2) = (2, 1/4, 1/4)$, equal errors η and $e_0 = 0$, compute the bound. Does the first expanding stage alone preclude an overall useful bound?

Solution. The three new error contributions are $\eta/16$, $\eta/4$ and η, totaling $21\eta/16$. The initial error would be multiplied by $2/16 = 1/8$. An expanding first stage therefore does not preclude overall contraction or a bounded implemented error, provided every domain condition is preserved. A common one-step contraction is sufficient but not necessary for this finite composite.

5.19.1 An optimal allocation of local precision

The weights in M27 also define a resource-allocation problem not solved by choosing the same error at every stage. Write $w_k = \prod_{j=k+1}^{n-1} q_j > 0$. Suppose the continuous precision cost at stage k is $c_k \log(1/\eta_k)$ with $c_k > 0$, and require $\sum_k w_k \eta_k \le E$. In the interior regime without active upper caps on η_k, the unique minimizer is

$$\eta_k = \frac{E c_k}{w_k \sum_j c_j}.$$

To prove this, the cost has positive-definite diagonal Hessian c_k/η_k^2, so it is strictly convex. Its derivative is negative in every coordinate, forcing the error budget to be saturated at an optimum. The Lagrange equations $c_k/\eta_k = \lambda w_k$ and the budget give $\lambda = (\sum c_k)/E$. These conditions are sufficient by convexity. Domain and maximum-error caps change the solution to the corresponding constrained KKT problem; the formula is not applied past those caps.

For the three-stage example with equal costs, the weights on newly inserted errors are $1/16, 1/4, 1$. The optimal tolerances are therefore $16E/3, 4E/3, E/3$ as long as all remain within the allowed domains and error caps. Earlier errors can be larger because later stages damp them. Equal precision is generally not optimal even when every individual implementation uses the same arithmetic. This is a mathematical cost model, not a measured runtime claim; real precision costs and integer bit allocations require a separate implementation study.

Chapter 6

Quantitative conformal examples

6.1 MU01. Geometric scale in a derivative estimate

Mathematical reference. Volume I, normal-family and mapping core M01; source Theorem 4.1

6.1.1 Prerequisites and governing statement

A disk of radius R consists of points at distance less than R from its centre. Cauchy's formula expresses derivatives by an integral on a surrounding circle.

$$f'(a) = \frac{1}{2\pi i} \int_{|z-a|=R} \frac{f(z)}{(z-a)^2}\, dz$$

The circle has length 2π R; the denominator has modulus R^2. The resulting scale is M/R.

$$|f'(a)| \le M/R, \qquad |f^{(j)}(a)| \le j!M/R^j.$$

6.1.2 Derivation and worked instance

Taking the absolute value inside the contour integral gives a product of three factors: the normalization $1/(2\pi)$, the circle length 2π R, and M/R^2. Their cancellation yields M/R. The same calculation with the j-th derivative produces $\mathrm{j!M}/R^j$. A change of length units scales an order-j derivative by the inverse j-th power. This provides a direct dimensional check of a printed estimate.

For f(z)=z on the disk of radius two, M=2 and $f'(0)$=1. The correct bound equals one. The displayed M/R^2 bound in the source equals one half and fails. The statement of normality remains valid after restoring the correct estimate.

6.1.3 Comparison and validity conditions

Baseline. Estimate a derivative from nearby differences, with a step length and data-error budget.

Proposed procedure. Use a bound on an enclosing analytic circle to obtain one derivative certificate throughout a smaller region.

Quantitative conclusion. The certified derivative bound uses the first inverse power of the available radius.

Validity conditions. The surrounding circle must lie within the analytic domain. A finite set of values on the circle does not by itself certify M.

Exercise 6.1. Let |f|≤6 on |z|≤3. Bound $|f'(0)|$ and $|f''(0)|$.

Solution. Cauchy gives 6/3=2 and $2!\,6/3^2$=4/3. Both bounds retain the physical dimension of the derivative. □

Graphical reference. See Figure 4.1 for the shared construction.

6.2 MU02. A finite polynomial with an explicit remainder

Mathematical reference. Volume I, normal-family and mapping core M03; source Little Montel proof

6.2.1 Prerequisites and governing statement

A power series can be truncated after degree n. A geometric series with ratio q in (0,1) has a computable tail.

$$\sum_{j=n+1}^{\infty} q^j = \frac{q^{n+1}}{1-q}$$

At q=1/2, the tail is 2^{-n}.

$$\|f - p_n\|_{\infty,\overline{\mathbb{D}_r}} \le M\frac{(r/R)^{n+1}}{1-r/R}.$$

6.2.2 Derivation and worked instance

Cauchy bounds every coefficient by M/R^j. Summing the absolute values of the omitted terms proves the displayed bound. The resulting stopping rule is valid simultaneously at every point of the smaller disk. The rule depends on the radius ratio, which describes the separation between the evaluation set and the available analytic boundary.

With M=2 and r/R=1/2, the bound is 2^{1-n}. Degree eleven gives 2^{-10}; degree ten gives 2^{-9}. The input consists of both the coefficients and the surrounding analytic bound.

6.2.3 Comparison and validity conditions

Baseline. Increase polynomial degree until several computed values appear stable.

Proposed procedure. Select a degree from a uniform remainder inequality on the entire evaluation disk.

Quantitative conclusion. A whole-disk tolerance replaces an unquantified visual stopping test.

Validity conditions. A smooth real function need not have a holomorphic extension to the proposed disk. A nearby complex pole invalidates the assumed radius.

Exercise 6.2. For M=1 and r/R=1/3, find a sufficient degree for error below 1/1000.

Solution. The bound is $1/(2\,3^n)$. At n=5 it is 1/486, while at n=6 it is 1/1458. Degree six is the smallest degree selected by this bound. □

6.3 MU03. Local uniform convergence

Mathematical reference. Volume I, normal-family and mapping core M02 and M09; source Theorem 4.1 and Exercise 4.3.

6.3.1 Prerequisites and governing statement

Uniform convergence on a set requires one index valid at every point of that set. Compact subdisks stay a positive distance from the unit circle.

$$\sup_{|z|\le r} |z^n| = r^n \quad (0 < r < 1)$$

The same sequence has supremum one on the open unit disk, so whole-domain uniform convergence fails.

$$z^n \longrightarrow 0 \text{ locally uniformly on } \mathbb{D}, \qquad \sup_{\mathbb{D}} |z^n| = 1.$$

6.3.2 Derivation and worked instance

Fix r<1. The elementary geometric limit $r^n \to 0$ proves uniform convergence on the closed disk of radius r. Every compact subset of the unit disk is contained in such a disk. The convergence is therefore locally uniform. Choosing evaluation points tending to the boundary with n changes the quantifiers and can prevent convergence.

At z_n=1-1/n, the values z_n^n approach e^{-1}. This does not contradict the locally uniform result because no fixed compact subdisk contains the entire sequence of evaluation points.

6.3.3 Comparison and validity conditions

Baseline. Check convergence separately at finitely many points.

Proposed procedure. Use a cofinal family of compact subdisks and record a uniform error on each.

Quantitative conclusion. The error on a chosen subdisk is exactly r^n.

Validity conditions. Pointwise convergence alone does not permit unrestricted passage of derivatives through a limit. Local uniform holomorphic convergence supplies that theorem through Cauchy's formula.

Exercise 6.3. At radius r=1/2, how many terms are required to make $|z^n| \le 1/1024$?

Solution. n=10 is both sufficient and necessary for this exact uniform estimate. □

6.4 MU04. Infinity as a target coordinate

Mathematical reference. Volume I, normal-family and mapping core M09–M10; source Exercises 4.3–4.4.

6.4.1 Prerequisites and governing statement

The Riemann sphere adds one point at ∞. Near that point, the coordinate 1/w replaces w.

$$\chi(w, \infty) = \frac{1}{\sqrt{1 + |w|^2}}$$

For w=10, this error is $1/\sqrt{101}$; it is finite and less than one tenth.

$$f_n(z) = n \quad \Longrightarrow \quad \sup_K \chi(f_n, \infty) = \frac{1}{\sqrt{1 + n^2}}.$$

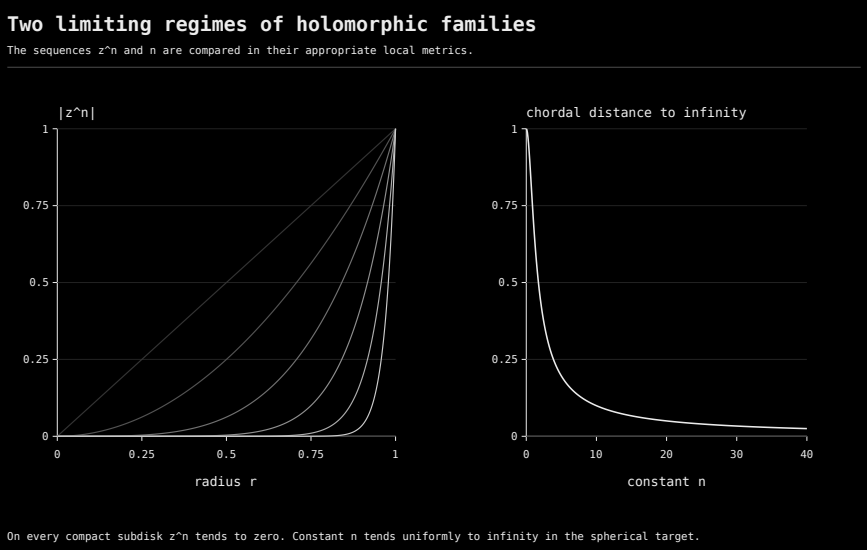


Figure 6.1: The error on a chosen subdisk is exactly r^n. The figure is generated from the stated formula; its pixels have no independent proof status.

6.4.2 Derivation and worked instance

The sequence of constant functions has no finite Euclidean limit. In the sphere it converges uniformly to the constant map ∞. The convergence rate is independent of the source point. This gives an elementary instance of normality with an infinite limiting value. A holomorphic sequence with finite-valued terms cannot converge locally spherically to a nonconstant meromorphic function possessing an isolated pole on a connected domain; core M10 proves the obstruction.

The homogeneous pairs [n:1] can be rescaled to [1:1/n]. The second representation remains bounded as n grows and converges coordinatewise to [1:0].

6.4.3 Comparison and validity conditions

Baseline. Judge convergence from Euclidean magnitudes alone.

Proposed procedure. Use the spherical target metric when a sequence approaches ∞.

Quantitative conclusion. The exact chordal error gives a stopping criterion near ∞.

Validity conditions. The pair [0:0] represents no projective point. Every numerical interface using homogeneous coordinates must reject it.

Exercise 6.4. Give a sufficient integer n for chordal error at most 1/100.

Solution. One needs $1+n^2 \geq 10000$. The smallest integer is n=100, since $99^2+1=9802<10000$. □

6.5 MU05. The algebra behind the extremal mapping argument

Mathematical reference. Volume I, normal-family and mapping core M06–M07; source Riemann mapping proof and Exercise 4.2.

6.5.1 Prerequisites and governing statement

A function is injective when distinct inputs have distinct outputs. Compositions of injective maps remain injective on their stated domains.

$$G(b) = \frac{1+b}{2\sqrt{b}}, \qquad 0 < b < 1$$

Writing q=$\sqrt{b}$ gives 1+b-2$\sqrt{b}$=(1-q)^2>0.

$$G(b) - 1 = \frac{(1-\sqrt{b})^2}{2\sqrt{b}} > 0.$$

6.5.2 Derivation and worked instance

The construction uses a disk automorphism carrying an omitted value to zero, a square-root branch, and another disk automorphism. Its derivative at the marked point is larger by the factor G(b). A maximizing embedding cannot admit this improvement. Therefore it omits no disk point. Existence of the maximizing embedding is a separate compactness argument; the inequality alone does not prove the mapping theorem.

For b=1/4 the gain is 5/4. The associated double covering has derivative modulus 4/5 at zero. These reciprocal quantities can be checked without numerical root finding.

6.5.3 Comparison and validity conditions

Baseline. Attempt to guess a conformal map onto a general simply connected domain.

Proposed procedure. Normalize the derivative, prove an extremum exists, and rule out omitted target points by an explicit improvement.

Quantitative conclusion. The strict gain excludes every omitted point with 0<|b|<1.

Validity conditions. The omitted target point is a critical value of the double covering. Confusing it with the critical point changes the lifting domain.

Exercise 6.5. Show that G(b) approaches one as b approaches one, and interpret the loss of margin.

Solution. The displayed formula tends to zero for G(b)-1. The extremal contradiction stays strict for every interior omitted point, but its quantitative margin degenerates near the boundary. □

Graphical reference. See Figure 4.3 for the shared construction.

6.6 MU06. Solving the Koebe inverse equation

Mathematical reference. Volume I, normal-family and mapping core M27; source Exercise 4.15.

6.6.1 Prerequisites and governing statement

A quadratic equation has two algebraic roots. A single-valued inverse on a domain also requires a consistent branch choice.

$$w = \frac{z}{(1-z)^2} \quad \Longleftrightarrow \quad wz^2 - (2w+1)z + w = 0$$

The discriminant is 1+4w. Its zero occurs at w=-1/4.

$$z(w) = \frac{4w}{(1+\sqrt{1+4w})^2}.$$

6.6.2 Derivation and worked instance

The Cayley identity 1+4k(z)=((1+z)/(1-z))^2 shows that the image avoids the closed ray (-∞,-1/4]. On the complementary slit domain the square root with positive real part is single-valued. Substitution gives the inverse in the display. Its Cayley coordinate lies in the right half-plane, which places the inverse inside the unit disk.

At w=0 the rationalized expression is zero. At w=2 the square root is three and z=8/16=1/2; direct substitution gives k(1/2)=2.

6.6.3 Comparison and validity conditions

Baseline. Apply the quadratic formula independently at each target point.

Proposed procedure. Select the root normalized by z(0)=0 and state the slit-domain square-root branch.

Quantitative conclusion. The inverse is explicit on the entire slit domain and satisfies an exact substitution identity.

Validity conditions. The formula with an unspecified square root is multivalued. The critical value -1/4 belongs to the boundary of the inverse domain.

Exercise 6.6. Find k(-1/2) and verify the inverse value.

Solution. k(-1/2)=(-1/2)/(3/2)^2=-2/9. Then sqrt(1-8/9)=1/3 and the inverse equals (-8/9)/(4/3)^2=-1/2. □

6.7 MU07. Clearance and the derivative at the centre

Mathematical reference. Volume I, normal-family and mapping core M26 and M28; source quarter theorem and Corollary 4.18.

6.7.1 Prerequisites and governing statement

The clearance of a marked point is its distance to the domain boundary. The conformal radius is the derivative modulus of a normalized disk map.

$$c = |\phi'(0)|, \qquad d = \operatorname{dist}(\phi(0), \partial\phi(\mathbb{D}))$$

For a centred round disk of radius R, both quantities equal R.

$$\frac{c}{4} \le d \le c.$$

6.7.2 Derivation and worked instance

The upper bound follows by applying Schwarz's lemma to the inverse on the largest inscribed disk. The lower bound is the Koebe quarter theorem. Both refer to the actual image of the conformal map. Replacing that image by a larger containing set changes the clearance and can invalidate the lower or upper interpretation.

If c lies in [7,9], the resulting clearance enclosure is [7/4,9]. The ratio four is universal. A known geometric subclass can produce a narrower enclosure.

6.7.3 Comparison and validity conditions

Baseline. Estimate a nearest boundary point from a finite drawing.

Proposed procedure. Use a verified derivative enclosure and the universal radius inequalities.

Quantitative conclusion. One derivative enclosure yields a whole-boundary clearance certificate.

Validity conditions. A map into a named domain need not be onto it. An inclusion map into a large disk supplies a counterexample to careless image substitution.

Exercise 6.7. For a disk of radius R and marked point at distance R/2 from its centre, compute c/d.

Solution. Here c=R(1-1/4)=3R/4 and d=R/2, so c/d=3/2. □

Graphical reference. See Figure 4.12 for the shared construction.

6.8 MU08. Normalization prevents degeneration

Mathematical reference. Volume I, normal-family and mapping core M05, M18 and M20; source Theorem 4.12 and Exercise 4.14(i).

6.8.1 Prerequisites and governing statement

A family of injective functions can converge to a constant. A nonzero derivative at a fixed point excludes that particular limit.

$$f_n(z) = z/n, \qquad f_n'(0) = 1/n$$

The functions are injective for every n, but their locally uniform limit is zero.

$$|f(0)| \le C, \qquad C^{-1} \le |f'(0)| \le C.$$

6.8.2 Derivation and worked instance

After subtracting the base value and dividing by the derivative, every function belongs to the normalized compact univalent class. Bounded base values and derivatives allow convergent subsequences of the normalization parameters. The positive lower derivative bound ensures that the resulting limit is nonconstant. Hurwitz's theorem then preserves injectivity.

The family nz violates the upper derivative bound and has no finite-valued locally uniform subsequence. The family z/n violates the lower derivative bound and converges to a constant. Translations z+n violate the base-value bound.

6.8.3 Comparison and validity conditions

Baseline. Assume that a limit retains all properties of every approximating map.

Proposed procedure. Track base values and both upper and lower derivative bounds throughout the compactness argument.

Quantitative conclusion. Three independently necessary controls identify three distinct degenerations.

Validity conditions. Compactness in the sphere and compactness inside the space of finite-valued univalent maps impose different limit requirements.

Exercise 6.8. Give one family satisfying the derivative bounds but failing the base-value bound.

Solution. f_n(z)=z+n has derivative one for every n, while $f_n(0)$=n is unbounded. □

Chapter 7

Analytic spaces and extremal arguments

7.1 MG01. The topology on a space of meromorphic maps

Mathematical reference. Volume I, normal-family and mapping core M09; source Exercise 4.3.

7.1.1 Prerequisites and governing statement

A compact exhaustion K_m is cofinal when every compact set lies in some K_m. A summable weighted series of bounded pseudometrics can separate functions.

$$d(f,g)=\sum_{m\geq 1}2^{-m}\sup_{K_m}\chi(f,g)$$

The tail after m=M is at most 2^{-M}, independently of the functions.

$$d(f,g)\leq \max_{m\leq M}\sup_{K_m}\chi(f,g)+2^{-M}.$$

7.1.2 Derivation and worked instance

A d-Cauchy sequence is uniformly Cauchy in the complete spherical target on each compact K_m. The limits agree on overlaps. Near a point where the limit is finite, a fixed affine target chart contains the nearby image and the usual holomorphic-limit theorem applies. Near an infinite value use the reciprocal chart. This proves closure of the meromorphic-map space and completeness of the metric.

To establish a metric error below 2^{-12}, it is sufficient to control the first thirteen compact sets to error 2^{-13}; the omitted tail contributes at most another 2^{-13}. The finite check must still cover each compact continuum.

7.1.3 Comparison and validity conditions

Baseline. Use an informal statement of convergence on every compact set.

Proposed procedure. Construct a complete metric with explicit tail control and local target-chart closure.

Quantitative conclusion. The metric gives an explicit decomposition into finitely controlled terms and a universal series tail.

Validity conditions. Completeness of the target is only one step. Holomorphy in compatible target charts is needed for completeness of the meromorphic subspace.

Exercise 7.1. Prove that d(f,g)=0 implies f=g.

Solution. Every nonnegative summand vanishes. Each point belongs to some exhaustion compact, so its chordal distance is zero. The chordal metric separates sphere points. □

Graphical reference. See Figure 4.6 for the shared construction.

7.2 MG02. Jets, local algebras, and analytic realization

Mathematical reference. Volume I, normal-family and mapping core M03; source bounded-family compactness; algebraic-geometric extension.

7.2.1 Prerequisites and governing statement

The quotient C[z]/(z^{N+1}) records a germ through order N. Its elements can be represented by coefficient vectors.

$$J_N = \mathbb{C}[z]/(z^{N+1}), \qquad \dim_{\mathbb{C}} J_N = N + 1$$

The class of z is nilpotent of order N+1, although z is not nilpotent as an analytic function.

$$J_{N+1} \longrightarrow J_N, \quad [p] \longmapsto [p] \pmod{z^{N+1}}.$$

7.2.2 Derivation and worked instance

The projection maps form an inverse system of finite-dimensional algebras. A compatible family is a formal power series. Convergence is an extra condition: coefficient growth must permit a positive analytic radius. The sequence j! defines a compatible formal series with radius zero, since successive terms at nonzero z eventually increase in magnitude.

For an analytic bound $|a_j| \leq M/R^j$, every truncation inherits the geometric remainder estimate. Without this bound, increasing N need not improve approximation on any fixed nonzero disk.

7.2.3 Comparison and validity conditions

Baseline. Identify a polynomial table with a complete analytic object.

Proposed procedure. Use the jet algebra for finite operations and a separate analytic-radius bound for realization.

Quantitative conclusion. Finite algebraic compatibility and analytic tail control become separate verifiable obligations.

Validity conditions. The algebraic inverse limit is C[[z]]. It contains many nonconvergent series and is not identified with the analytic germ ring.

Exercise 7.2. Which formal series has coefficients a_j=j!, and what is its radius of convergence?

Solution. The formal series is sum j! z^j. For any z≠0 the ratio of successive term magnitudes is (j+1)|z|, which eventually exceeds one. Its radius is zero. □

7.3 MG03. Attainment in the extremal proof

Mathematical reference. Volume I, normal-family and mapping core M04–M06; source Riemann mapping theorem.

7.3.1 Prerequisites and governing statement

Normality gives convergent subsequences. A derivative normalization can exclude constant limits of univalent maps.

$$s = \sup\{|f'(a)| : f : D \hookrightarrow \mathbb{D},\ f(a) = 0\}$$

A bounded embedding shows s>0; a local Cauchy estimate shows s<∞.

$$|f_n'(a)| \to s > 0 \quad \Longrightarrow \quad |f'(a)| = s.$$

7.3.2 Derivation and worked instance

Local uniform convergence implies derivative convergence by Cauchy's formula. The positive limiting derivative excludes a constant limit. A locally uniform limit of injective holomorphic maps is injective unless constant. The maximum principle then puts the nonconstant limit inside the open unit disk. These facts establish actual attainment, which licenses the omitted-value contradiction.

The chain of implications contains two closure statements: holomorphic closure of the bounded family and injective-or-constant closure of the univalent family. Treating the latter as unconditional injective closure would fail for z/n.

7.3.3 Comparison and validity conditions

Baseline. Write down a supremum and assume that a maximizing map exists.

Proposed procedure. Extract a maximizing sequence and prove that its limit remains in the admissible class.

Quantitative conclusion. The extremum is proved to be attained with a nonzero derivative.

Validity conditions. The proof supplies existence. A numerical construction requires an effective representation of D and independent approximation estimates.

Exercise 7.3. Why does a positive derivative at a imply that the image of the limit lies in the open disk?

Solution. The limit has values in the closed disk. If it attained a point of modulus one at an interior source point, the maximum principle would make it constant, contradicting the nonzero derivative. □

Graphical reference. See Figure 4.3 for the shared construction.

7.4 MG04. Cross-ratios and colliding target sections

Mathematical reference. Volume I, normal-family and mapping core M15–M16; source Theorem 4.9, Exercise 4.10 and hint p.633.

7.4.1 Prerequisites and governing statement

Three distinct projective points determine a unique Möbius transformation to 0,1,∞. Equality of two nonidentical meromorphic functions has isolated zeros.

$$H_z(w) = \frac{(w-\psi_1)(\psi_2-\psi_3)}{(w-\psi_3)(\psi_2-\psi_1)}$$

This affine formula sends psi1 to zero and psi2 to one wherever all denominators make sense.

$$\operatorname{wind}(g_n - \psi) = \operatorname{wind}(g_n) \geq 0, \qquad \operatorname{wind}(g_n - \psi) = -m.$$

7.4.2 Derivation and worked instance

At a collision, transform one omitted section to ∞. The maps become holomorphic and another target becomes a pole of order m. Punctured normality gives either finite local convergence or uniform escape on a small circle. In the escape case the pole target is a smaller boundary perturbation. Winding is preserved, but the meromorphic argument principle gives minus the pole order because the target is omitted. The contradiction extends normality across the collision.

The targets 0,z^3,1 are pairwise nonidentical. The transformation w$\mapsto$1/w converts the collision at zero into the pole z^{-3}, whose winding on a positive circle is -3.

7.4.3 Comparison and validity conditions

Baseline. Normalize three moving values in an affine chart without treating collisions.

Proposed procedure. Use homogeneous determinants off the collision set and a winding argument at each isolated collision.

Quantitative conclusion. An isolated collision is resolved by a finite integer obstruction.

Validity conditions. Pairwise nonidentity as functions is necessary. Repeated identical target sections supply fewer than three omission constraints.

Exercise 7.4. Explain why the transformed g_n are holomorphic even when the original f_n have poles.

Solution. At an original pole the reciprocal coordinate equals zero and extends holomorphically. Omission of the chosen finite target excludes poles of the reciprocal transform. □

7.5 MG05. Area and coefficient inequalities

Mathematical reference. Volume I, normal-family and mapping core M22–M24; quantitative expansion of source Koebe theory.

7.5.1 Prerequisites and governing statement

A univalent exterior function has a Laurent expansion at ∞. Its omitted area is nonnegative.

$$F(\zeta) = \zeta + b_0 + \sum_{n\geq 1} b_n \zeta^{-n}$$

The coefficient b0 is a translation and contributes no area term.

$$\sum_{n\geq 1} n|b_n|^2 \leq 1.$$

7.5.2 Derivation and worked instance

For the image of a circle of radius R>1, Green's formula and Fourier orthogonality give area π times (R^2-sum n$|b_n|^2$ R^{-2n}). Nonnegativity bounds every finite partial sum. Letting R decrease to one and then the partial-sum length increase proves the area inequality. Applied to the exterior transform of the odd square root of f(z^2), it yields the coefficient estimate $|a_2|\leq 2$.

The real slice $b_1^2+2b_2^2\leq 1$ is an ellipse. A coefficient choice $b_1=4/5$ leaves budget 9/25, so $|b_2|\leq 3/(5\text{sqrt}2)$.

7.5.3 Comparison and validity conditions

Baseline. Estimate the first coefficients separately by local Cauchy bounds.

Proposed procedure. Integrate the image boundary and obtain one weighted global coefficient budget.

Quantitative conclusion. One area invariant constrains every retained Laurent coefficient simultaneously.

Validity conditions. An admissible coefficient budget is a necessary condition; arbitrary coefficients within this ellipse do not automatically define a univalent function.

Exercise 7.5. If b1=0 and b2=1/2, what budget remains for 3|b3|^2?

Solution. The first retained contribution is $2(1/2)^2=1/2$. Thus $3|\text{b3}|^2\leq 1/2$ and $|\text{b3}|\leq 1/\text{sqrt}6$. □

Graphical reference. See Figure 4.13 for the shared construction.

7.6 MG06. The Schwarzian as a projective differential invariant

Mathematical reference. Volume I, normal-family and mapping core M37; source Koebe distortion as the analytic starting point.

7.6.1 Prerequisites and governing statement

A locally univalent map has nonzero derivative, allowing the logarithmic derivative f second/f prime.

$$Sf = \frac{f'''}{f'} - \frac{3}{2}\left(\frac{f''}{f'}\right)^2$$

For an affine map, both numerator derivatives vanish and Sf=0.

$$S(f\circ g) = (Sf\circ g)(g')^2 + Sg, \qquad (1-|z|^2)^2|Sf(z)| \leq 6.$$

7.6.2 Derivation and worked instance

The chain rule follows by differentiating the logarithmic derivative and collecting terms. Postcomposition by a Möbius map leaves the Schwarzian unchanged. The coefficient inequality $|a_3\text{-}a_2^2|\leq 1$ bounds S f(0)=6(a_3-a_2^2); disk automorphisms transport the estimate to a general point. A local square root of f prime produces y1=f/sqrt(f prime) and y2=1/sqrt(f prime), both satisfying y second+(Sf/2)y=0.

For k(z)=z/(1-z)^2, Sk=-6/(1-z^2)^2. On the real diameter, the normalized magnitude is exactly six.

7.6.3 Comparison and validity conditions

Baseline. Track three separate derivatives under every coordinate composition.

Proposed procedure. Use a projective invariant with a composition law and a normalized bound.

Quantitative conclusion. The bound is sharp and invariant under postcomposition by Möbius maps.

Validity conditions. The displayed inequality is a necessary condition for normalized univalence. A sufficient criterion requires a separate theorem with its precise hypotheses.

Exercise 7.6. Compute the Schwarzian of f(z)=exp(z).

Solution. All three derivatives equal exp(z), so Sf=1-3/2=-1/2. This local calculation does not establish global injectivity of exp. □

7.7 MG07. The Bergman pullback is unitary

Mathematical reference. Volume I, normal-family and mapping core M36; functional-analytic and differential-form extension.

7.7.1 Prerequisites and governing statement

The Bergman space A^2(D) consists of holomorphic functions square-integrable for planar area measure. A conformal Jacobian equals the derivative modulus squared.

$$\|F\|^2_{A^2(D)} = \int_D |F|^2 \, dA$$

The disk monomials have squared norms $\pi/$(n+1).

$$U_\phi F = (F \circ \phi)\phi', \qquad \|U_\phi F\|_{A^2(\mathbb{D})} = \|F\|_{A^2(D)}.$$

7.7.2 Derivation and worked instance

The change-of-variables theorem gives equality of squared norms directly. Applying the same construction to the inverse supplies surjectivity. The normalized monomials sum to the disk kernel K(z,w)=1/[π(1-z conjugate(w))^2]. Kernel transformation at a=ϕ(0) then gives K_D(a,a)=1/[$\pi|\phi$ prime(0)$|^2$]. Thus an evaluation functional records the conformal radius.

For a domain with conformal radius c=2 at a, the evaluation norm is 1/(2sqrt π). The normalized reproducing kernel attains the bound on the A^2 unit ball.

7.7.3 Comparison and validity conditions

Baseline. Compose functions with a coordinate map and assume that their area norm is unchanged.

Proposed procedure. Include the derivative factor in the pullback of holomorphic one-forms.

Quantitative conclusion. The norm identity is exact; the evaluation bound is attained.

Validity conditions. Changing the area normalization changes the kernel constant. Hardy and Bergman norms have different underlying measures.

Exercise 7.7. Find the normalized constant basis function on the unit disk.

Solution. The disk area is π, so the function 1/sqrt π has squared norm one. □

Graphical reference. See Figure 4.21 for the shared construction.

7.8 MG08. A matching lower bound for finite-rank approximation

Mathematical reference. Volume I, normal-family and mapping core M34–M35; source finite compact restrictions motivate the extension.

7.8.1 Prerequisites and governing statement

A linear map from an (N+1)-dimensional space to an N-dimensional range has a nonzero kernel vector.

$$\dim V = N + 1, \quad \operatorname{rank}(A|_V) \le N \quad \Longrightarrow \quad \ker(A|_V) \neq \{0\}$$

Normalize one kernel vector to norm one.

$$\inf_{\operatorname{rank} A \le N} \|T_r - A\| = r^N.$$

7.8.2 Derivation and worked instance

Let V be spanned by the first N+1 monomials. For a unit vector v in the kernel of A restricted to V, the coefficient weights of T_r on V are all at least r^N. Therefore the residual norm is at least r^N. The first-N-coordinate truncation has residual norm exactly r^N. This proves equality without assuming that the competitor preserves the coefficient basis.

At r=1/2,N=3, every rank-three operator has worst-case error at least 1/8. The truncation retaining 1,z,z^2 attains this value.

7.8.3 Comparison and validity conditions

Baseline. Demonstrate that a particular truncation has small error.

Proposed procedure. Prove that every bounded linear rank-N competitor incurs at least the same worst-case error.

Quantitative conclusion. An attained upper bound and a universal lower bound coincide.

Validity conditions. The proof quantifies bounded complex-linear competitors. A nonlinear compression scheme requires a different width or information model.

Exercise 7.8. Which input attains the error of the coordinate truncation itself?

Solution. The monomial z^N has unit H^2 norm. It is discarded by the truncation and mapped by T_r to $r^N\ z^N$, giving error r^N. □

Chapter 8

Parameter families and optimal approximation

8.1 MP01. Quantitative total boundedness of analytic restrictions

Mathematical reference. Volume I, normal-family and mapping core M01–M03; source Little Montel theorem.

8.1.1 Prerequisites and governing statement

A finite ε-net approximates every member of a metric set. Coefficient truncation leaves a tail controlled by an analytic bound.

$$|a_j| \le MR^{-j}, \quad \|f - p_N\|_r \le M\frac{(r/R)^{N+1}}{1 - r/R}$$

Choose N so the tail occupies at most one half of the target tolerance.

$$\sum_{j=0}^{N} |\Delta a_j| r^j \le \varepsilon/2.$$

8.1.2 Derivation and worked instance

Quantize $b_j = a_j R^j$ in the coefficient disk $|b_j| \le M$ using rectangular cells of side at most $h = \varepsilon(1-q)/4$, where $q = r/R$. Each coordinate error is at most h. The weighted coefficient error is at most $\sqrt{2}h/(1-q) < \varepsilon/2$. At most $(1 + 8M/[\varepsilon(1-q)])^2$ cells per coefficient are required. Independent choices for the first $N+1$ coefficients give an explicit finite covering count.

The construction gives a quantitative upper bound on metric entropy. It also exhibits its deterioration as q approaches one. Compactness at each fixed interior radius is compatible with arbitrarily large covering counts near the boundary.

8.1.3 Comparison and validity conditions

Baseline. Invoke relative compactness without a finite covering estimate.

Proposed procedure. Quantize finitely many normalized coefficients and count a valid covering family.

Quantitative conclusion. A qualitative compactness claim acquires a finite constructive witness for a bounded analytic class.

Validity conditions. The covering count is an explicit upper bound, not a proved entropy asymptotic or sharp minimum.

Exercise 8.1. Why does coefficientwise quantization need a geometric weighting in the error norm?

Solution. A coefficient error contributes |Delta a_j|r^j to the evaluation error. Normalizing by R converts this weight to q^j, whose sum is finite for q<1. □

Graphical reference. See Figure 4.2 for the shared construction.

8.2 MP02. Schatten scales and boundary degeneration

Mathematical reference. Volume I, normal-family and mapping core M34–M35; functional-analytic extension.

8.2.1 Prerequisites and governing statement

A compact positive diagonal operator has singular values equal to its diagonal entries. The Schatten p norm is their l-p norm.

$$\|T\|_{\mathcal{S}_p}^p = \sum_{j\geq 0} s_j(T)^p$$

For geometric singular values r^j, the series is explicit.

$$\|T_r\|_{\mathcal{S}_p}^p = \frac{1}{1-r^p}, \qquad \|T_r\| = 1.$$

8.2.2 Derivation and worked instance

Every T_r with r<1 is in all positive Schatten classes. As r increases to one, the operator norms stay equal to one while each Schatten quantity diverges. Moreover T_r converges strongly to the identity by dominated convergence in coefficient space, but its norm distance from the identity is one for every r<1. The example separates strong convergence, norm convergence, and compactness.

For p=2 and r=1/2, the Hilbert-Schmidt norm squared is 4/3. For p=1 the trace is two. Both can be checked by finite geometric sums with exact tails.

8.2.3 Comparison and validity conditions

Baseline. Use the operator norm alone to measure approximation complexity.

Proposed procedure. Record the entire singular spectrum and its trace and Hilbert-Schmidt scales.

Quantitative conclusion. The spectral calculation gives exact values at every radius and identifies a singular boundary regime.

Validity conditions. The identity operator on infinite-dimensional H^2 is not compact. Strong limits of compact operators need not be compact.

Exercise 8.2. Prove that ||I-T_r||=1 for r<1.

Solution. The diagonal entries are 1-r^j, bounded above by one and tending to one as j tends to ∞. Their supremum is one. □

8.3 MP03. A length-area proof of an explicit collar bound

Mathematical reference. Volume I, normal-family and mapping core M29; source Lemma 4.17.

8.3.1 Prerequisites and governing statement

The condenser capacity is the infimum of Dirichlet energies of functions equal to one near K and zero at the outer boundary.

$$\operatorname{Cap}(K, \mathbb{D}) = 1/\operatorname{mod}(\mathbb{D} \setminus K)$$

The equality uses the source convention mod A=(2pi)$^{-1}$log(R/r).

$$\operatorname{Cap}(K, \mathbb{D}) \geq \frac{4}{\pi} \log \frac{1}{1-r}, \qquad r \leq 1 - e^{-\pi/(4\mu)}.$$

8.3.2 Derivation and worked instance

Choose a point of K with maximal modulus r and its radial endpoint xi on the unit circle. Connectedness and the presence of zero imply that every circle about xi of radius t between 1-r and one meets K. Its portion inside the disk is an arc of length at most π t. A test function rises from zero to one and returns to zero along this arc. Cauchy-Schwarz on both portions gives energy at least $4/(\pi$ t). Integration over t proves the bound.

For a concentric inner disk the exact outer radius is exp(-2pi μ). The explicit bound above is much weaker but applies to arbitrary admissible full continua.

8.3.3 Comparison and validity conditions

Baseline. Use a compactness contradiction with no quantitative radius.

Proposed procedure. Integrate a one-dimensional energy lower bound over circles centered on the outer boundary.

Quantitative conclusion. The proof supplies a formula for k(μ) and removes an uncontrolled limiting-map step.

Validity conditions. The set K must satisfy the ring-domain hypotheses. Core M38 supplies a separate circular-rearrangement proof of radial-slit extremality.

Exercise 8.3. Where does the factor four enter the energy estimate?

Solution. If the two arc lengths are a and b, the energy is at least 1/a+1/b. Since (a+b)$^2 \geq$4ab, this is at least 4/(a+b). □

Graphical reference. See Figure 4.15 for the shared construction.

8.4 MP04. Moving-target separation and local conditioning

Mathematical reference. Volume I, normal-family and mapping core M15–M16; source Theorem 4.9 and Exercise 4.10.

8.4.1 Prerequisites and governing statement

The cross-ratio is rational in its target coordinates. Its coefficients involve pairwise differences.

$$H(w) = \frac{(w-a)(b-c)}{(w-c)(b-a)}$$

The denominator coefficient b-a is nonzero for a distinct triple.

$$|b-a| \geq \delta > 0 \quad \Longrightarrow \quad |(b-a)^{-1}| \leq \delta^{-1}.$$

8.4.2 Derivation and worked instance

On a compact parameter set with pairwise distinct target sections, their chordal separations have a positive minimum. A finite chart cover then bounds both the Möbius transformations and their inverses in suitable spherical Lipschitz constants. The constants deteriorate when the separation tends to zero. At isolated collisions, the meromorphic winding argument proves qualitative normality, but it supplies no uniform inverse-separation constant across the collision.

The triple (0,δ,1) gives H(w)=w(δ-1)/[(w-1)δ]. Its coefficient of the small separation is explicit. Homogeneous rescaling can prevent overflow while preserving the intrinsic sensitivity.

8.4.3 Comparison and validity conditions

Baseline. Treat Möbius normalization as numerically uniform over all distinct triples.

Proposed procedure. Separate exact projective existence from quantitative coefficient conditioning.

Quantitative conclusion. The dependency on target separation is exhibited algebraically and in the graphical diagnostic.

Validity conditions. Rescaling coordinates does not remove an ill-conditioned geometric configuration.

Exercise 8.4. For δ=1/100, write H(2).

Solution. H(2)=2(-99/100)/[(1)(1/100)]=-198. The large affine value is compatible with a finite spherical point. □

8.5 MP05. Global injection under analytic perturbation

Mathematical reference. Volume I, normal-family and mapping core M31–M33; source Koebe distortion and additional quantitative model.

8.5.1 Prerequisites and governing statement

A co-Lipschitz map satisfies a lower distance inequality. A derivative bound on a convex set integrates along straight segments.

$$|f(z) - f(w)| \geq m|z-w|, \qquad \sup|e'| \leq \delta < m$$

The perturbed map f+e has lower distance factor at least m-δ.

$$|(f+e)(z) - (f+e)(w)| \geq (m-\delta)|z-w|.$$

8.5.2 Derivation and worked instance

The proof uses the triangle inequality and the integral of e prime along the segment. If e is bounded by ε on a larger disk, Cauchy bounds δ by $\varepsilon/$(R-r). This introduces an explicit analytic collar into a topological injection certificate. For nonconvex domains, one must bound the ratio of connecting-path length to Euclidean distance; the same numerical margin cannot be transferred without that geometric factor.

The quadratic map with a=1/5,r=3/5 has m=19/25. With R=4/5 and ε=1/100, δ=1/20 and m-δ=71/100.

8.5.3 Comparison and validity conditions

Baseline. Test the Jacobian determinant at finitely many mesh points.

Proposed procedure. Combine a whole-domain lower distance witness with an analytic perturbation enclosure.

Quantitative conclusion. The admissible perturbation threshold is an explicit positive scalar.

Validity conditions. A nonzero derivative is a local condition. The proof needs the global distance inequality for the reference map.

Exercise 8.5. What happens when δ=m?

Solution. The lower bound becomes zero and the proof no longer certifies injection. Equality in a sufficient-condition threshold is not a proved counterexample for every model. □

Graphical reference. See Figure 4.18 for the shared construction.

8.6 MP06. A bounded observation model and an attained optimum

Mathematical reference. Volume I, normal-family and mapping core M35; framework optimality theorem.

8.6.1 Prerequisites and governing statement

Worst-case error is a supremum over an explicitly specified input class. An admissible algorithm class must also be fixed.

$$\mathcal{A}_N = \{A : H^2 \to H^2 : \ A \text{ bounded complex-linear, } \operatorname{rank} A \le N\}$$

A lower bound over A_N does not automatically quantify nonlinear methods.

$$\min_{A \in \mathcal{A}_N} \sup_{\|f\| \le 1} \|T_r f - Af\| = r^N.$$

8.6.2 Derivation and worked instance

The canonical truncation attains the upper bound. Rank-nullity supplies an adversarial unit input for every competitor. The equality is therefore a complete optimality statement in the model. The proof identifies both quantifier order and the space carrying the error norm. Enlarging the input class or changing the output norm requires a new theorem; restricting the class can improve the value.

For a prior supported on span1,z, rank two gives zero error. The universal H^2-unit-ball lower bound uses directions absent from that prior.

8.6.3 Comparison and validity conditions

Baseline. Compare one proposed approximation with one selected implementation.

Proposed procedure. Prove a minimax identity over every operator in the declared competitor class.

Quantitative conclusion. The minimum is attained and the lower bound covers the full stated competitor class.

Validity conditions. Algorithmic optimality is always relative to its admissible class and objective. A document-coverage ratio is a different mathematical quantity.

Exercise 8.6. Why does the lower-bound vector depend on A?

Solution. It is selected from the kernel of A on the first N+1 modes. Different operators can have different kernel directions; the worst-case supremum permits this dependence. □

8.7 MP07. Conformal kernels and geometric observables

Mathematical reference. Volume I, normal-family and mapping core M26, M28 and M36.

8.7.1 Prerequisites and governing statement

The reproducing kernel represents evaluation in a Hilbert space of holomorphic functions. Its diagonal is the squared evaluation norm.

$$|F(a)|^2 \le K_D(a,a)\|F\|^2_{A^2(D)}$$

The normalized kernel vector attains equality.

$$\frac{1}{16\pi d^2} \le K_D(a,a) = \frac{1}{\pi c^2} \le \frac{1}{\pi d^2}.$$

8.7.2 Derivation and worked instance

The core establishes c/4≤d≤c and the exact identity K_D(a,a)=1/(π c^2). Taking reciprocal squares gives the displayed bounds, with inequality directions tracked carefully. A verified conformal-radius enclosure yields a sharper kernel enclosure. The estimate concerns the infinite-dimensional A^2 space and the whole domain, even when its certificate uses one normalized derivative.

If c lies in [2,3], the kernel diagonal lies in [1/(9pi),1/(4pi)]. A finite-basis numerical kernel underestimates the full evaluation norm when the finite space is a subspace with the inherited inner product.

8.7.3 Comparison and validity conditions

Baseline. Approximate a point-evaluation norm numerically from a finite function basis.

Proposed procedure. Use the conformal transformation law and clearance inequalities to enclose the evaluation functional.

Quantitative conclusion. A geometric radius controls an attained infinite-dimensional evaluation bound.

Validity conditions. Numerical quadrature that changes the inner product can destroy the subspace monotonicity argument.

Exercise 8.7. Which assumption allows monotonicity under increasing approximation spaces?

Solution. The spaces must be nested and use the same Hilbert norm. The evaluation norm is the supremum over their unit balls, which increase by inclusion. □

Graphical reference. See Figure 4.21 for the shared construction.

8.8 Undergraduate comparative workshop: calculate, explain, challenge

8.8.1 MU11: a root count and a root list are different deliverables

Problem. For $p(z) = z^4 + 5z + 1$, prove how many roots lie in $|z| < 1$ and in $1 < |z| < 2$. The conventional route computes approximate roots, then classifies their moduli. To turn that route into a proof, each root must receive an enclosure with multiplicity and a separation from both circles. The proposed route uses two whole-boundary inequalities, as MU5 already derived.

Compare the exact work. On the unit circle, $|z^4 + 1| \le 2 < 5 = |5z|$. On the radius-two circle, $|5z + 1| \le 11 < 16 = |z^4|$. The dominant terms have known counts one and four. The annular difference is three. Two inequalities replace solving a quartic *for the counting task*; they do not return the four coordinates. If the next question asks for those coordinates to six decimal places, the numerical enclosure route still has work the counting certificate cannot replace.

Why a dense picture is not enough. For $f(z) = z^m$ on the unit circle, N equal samples are $e^{2\pi imj/N}$. When $m = N$, every sampled value is one, although the image winds N times. Thus an arbitrarily large but fixed sampling count has no universal ability to identify degree without a variation bound. This is an exact counterexample, independent of floating-point error.

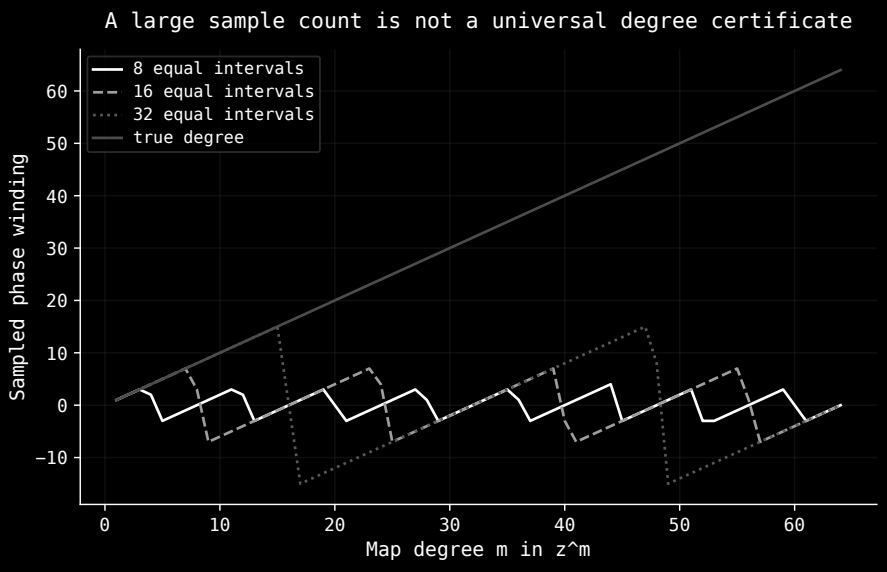


Figure 8.1: Sampled phase can alias the exact degree. Increasing a fixed sample count does not prove correctness for all maps; whole-arc information or a derivative bound changes the guarantee.

Exercise 8.8 (MU11a)**.** Show that the maps $z \mapsto z$ and $z \mapsto z^{N+1}$ agree on the Nth roots of unity but have different winding numbers. What missing datum would prevent this ambiguity?

Solution. At $z^N = 1$, $z^{N+1} = z$. Their continuous boundary windings are one and $N + 1$. A known bound on angular variation between adjacent samples, strict enough to rule out an extra turn, or a whole-segment exclusion/homotopy certificate supplies the missing information. Mere equality at the samples does not.

8.8.2 MU12: a branch is a path with an initial value

Consider $w(t) = e^{it}$, $0 \le t \le 2\pi$, with $z(0) = 1$ and $z(t)^2 = w(t)$. The continuous solution is $z(t) = e^{it/2}$ and ends at -1. A principal square-root calculation at the final value $w = 1$ returns $+1$. Both solve the final algebraic equation; only one solves the whole initial-value path problem.

Finite calculation before topology. With $N = 8$ equal intervals, the continuous candidates are $e^{ij\pi/8}$. At every step the distance to the correct next candidate is $2\sin(\pi/16)$, while the distance to its negative is $2\cos(\pi/16)$. The inequality between them proves that nearest matching follows the correct sheet in exact arithmetic. This introduces covering-space lifting as a precise continuation rule, not as unexplained vocabulary.

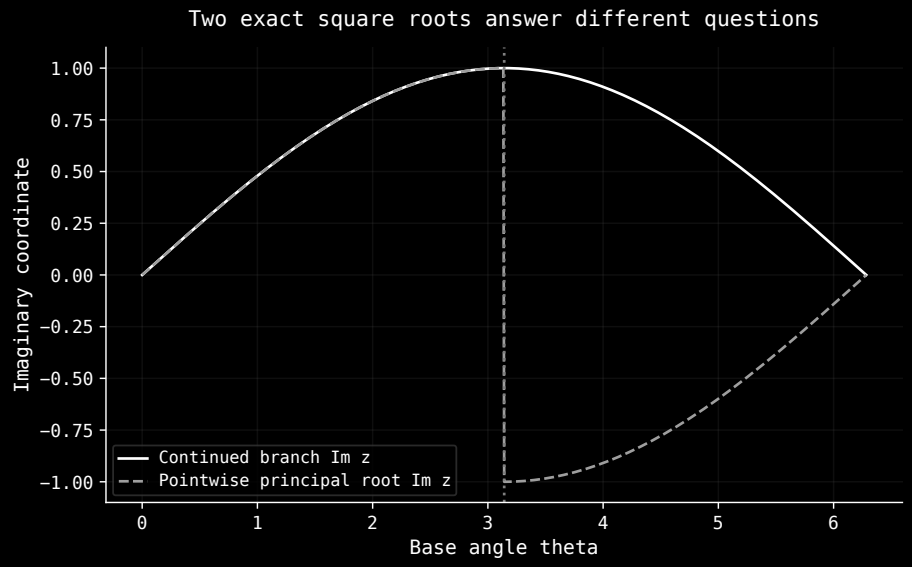


Figure 8.2: The imaginary coordinate of the continued root and the pointwise principal root. The jump is not a failure of the principal-root routine's own specification; it is a mismatch with the continued-path task.

The metric of improvement. The natural metric is satisfaction of the continuation contract and the number of required sheet corrections, not a residual: both candidates have zero algebraic residual. For the exact unit-circle example, a single extra bit recording the current sheet, together with the previous value and admissible step bound, carries the missing information. This does not imply that every algebraic family can be represented by one bit.

Exercise 8.9 (MU12a)**.** For $z^3 = e^{it}$ with $z(0) = 1$, compute the endpoint after one, two and three full base loops. State the subgroup of integer windings that close the lift.

Solution. The endpoints are $e^{2\pi i/3}$, $e^{4\pi i/3}$ and 1. A loop of integer winding m changes the lifted angle by $2\pi m/3$, so closure means $3 \mid m$. The subgroup is $3\mathbb{Z}$. This is the elementary calculation underlying the general covering subgroup criterion.

8.8.3 MU13: distinguish one repeated point from one dimension

At $\lambda = 0$, the equation $x^2 = \lambda$ has only one coordinate solution. The algebra $\mathbb{C}[x]/(x^2)$ still has basis $1, x$. A root-deduplication procedure discards the second direction; a quotient representation retains it through $x^2 = 0$ but $x \neq 0$.

A directly testable consequence. Let $J = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$. Then $e^J = I + J$. Replacing the algebra by the one-point set $\{0\}$ would suggest $e^0 = 1$ and lose the nonzero upper-right entry. The nilpotent direction stores first-derivative information. This is a concrete use of an algebraic-geometric object in a matrix calculation; it is not the assertion that all matrices are algebraic varieties.

Exercise 8.10 (MU13a)**.** For $N^3 = 0$ but $N^2 \neq 0$, express e^N and $(I - N)^{-1}$. What data beyond the single eigenvalue zero are needed?

Solution. The exponential is $I + N + N^2/2$ and the inverse is $I + N + N^2$, verified by multiplication. The matrix or its nilpotent algebraic structure is needed; an unordered set containing only the eigenvalue zero does not determine either answer.

8.8.4 MU14: equal arithmetic precision need not be efficient

Suppose four local errors are bounded by $2^{-b_1}, \ldots, 2^{-b_4}$, with $b_j \geq 2$, and their contributions to final error are weighted by $(1, 4, 16, 64)/64$. The target is $1/64$. Uniform precision requires $85 \cdot 2^{-b} \leq 1$, so the least uniform integer is $b = 7$, costing 28 bits in the declared sum-of-bits model.

The allocation $(2, 4, 6, 8)$ gives four equal contributions of $1/256$ and total $1/64$, with cost 20. This demonstrates a feasible reduction of eight bits, or $2/7$ of the uniform cost, without assuming that twenty is optimal. The research branch supplies the separate lower-bound argument. The distinction between a better example and a proof of optimality is part of the exercise.

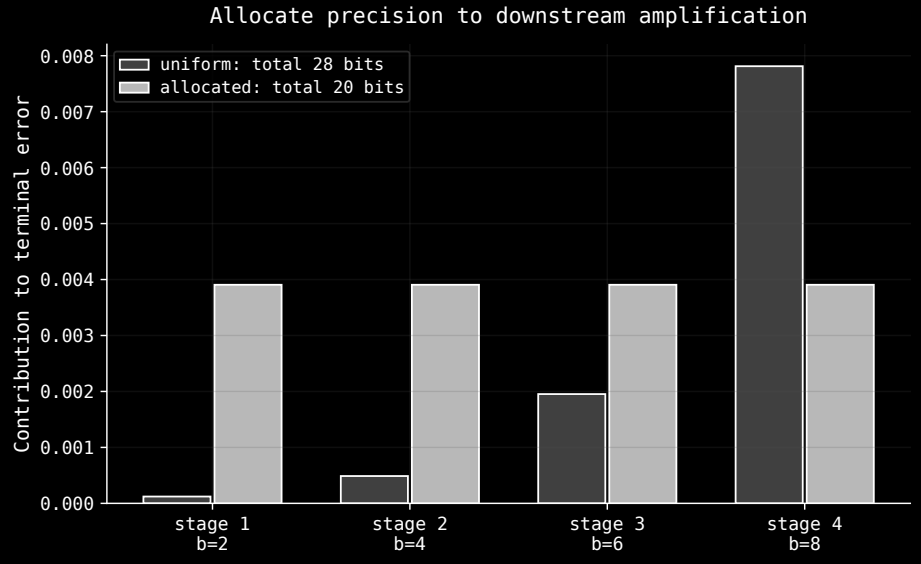


Figure 8.3: The uniform and allocated bounds satisfy the same target but spend precision differently. The displayed quantity is worst-case propagated error, not observed error of a circuit or solver.

Domain gate. A loose early tolerance is allowed only if it fits inside the next valid map domain. If a local tube admits error at most $1/32$, then $b_1 = 2$ is illegal there even if the terminal weighted

sum passes. The allocation must incorporate that cap; a lower terminal error does not legalize an invalid intermediate step.

Exercise 8.11 (MU14a)**.** Add the constraint $b_1 \geq 5$. Is $(2, 4, 6, 8)$ still admissible? Give one feasible replacement without claiming it is optimal.

Solution. The old allocation is inadmissible. Replacing it by $(5, 4, 6, 8)$ reduces the first contribution and preserves the other three, so the total remains at most $1/64$. Its cost is 23. Finding the least cost under the new cap requires another optimization calculation.

8.8.5 A compact comparison checklist

A completed solution identifies the problem's output, the baseline's valid guarantee, the proposed representation, a proved or measured difference, and one case in which the change is not justified. For source 3.2 this means checking infinity before computing genus; for source 3.5 it means distinguishing one sheet from a whole ordering; for source 3.6 it means preserving domains before optimizing precision. These are additional mathematical capabilities, not an assertion that the longer exposition has already been shown to improve student performance.

8.9 Graduate comparison: quotient, Hilbert geometry and covering data

The graduate criterion is not the number of theorem names used. It is the ability to derive a representation, state the norm or topology in which it is valid, and quantify what is gained or lost. The following units connect local algebra, global surface invariants and functional analysis without repeating the full proofs in the core.

8.9.1 MG9: a spectral decomposition is a coordinate choice

For $A(\lambda) = \left(\begin{smallmatrix} 0 & 1 \\ \lambda & 0 \end{smallmatrix}\right)$ and positive $\lambda \leq 1$, eigenvectors can be arranged as columns of

$$V_\lambda = \begin{pmatrix} 1 & 1 \\ \sqrt{\lambda} & -\sqrt{\lambda} \end{pmatrix}.$$

The rows are orthogonal and have lengths $\sqrt{2}$ and $\sqrt{2\lambda}$, so the spectral condition number is $\lambda^{-1/2}$. This is a property of the chosen eigenbasis. The original matrix remains bounded, and its exponential is entire in λ.

Proposed comparison. Retain the finite algebra $A^2 = \lambda I$ and use its even/odd functional calculus. The metric is the absence of the singular eigenbasis factor, together with a series-tail bound. This does not prove that every observable of A is well-conditioned: the eigenvalues themselves have a square-root singularity, and spectral projectors onto individual branches can blow up. The target observable determines the representation.

Exercise 8.12 (MG9a)**.** For $\lambda > 0$, write the two spectral projectors using A and $\sqrt{\lambda}$. Explain why their sum is regular but the individual projectors are not uniformly bounded as $\lambda \downarrow 0$.

Solution. The projectors are $P_\pm = (I \pm A/\sqrt{\lambda})/2$. Their sum is I. The upper-right entry of $P_\pm$ is $\pm 1/(2\sqrt{\lambda})$, so each norm diverges. A regular sum can conceal singular components. Replacing individual projectors by the full identity would answer a different spectral-selection question.

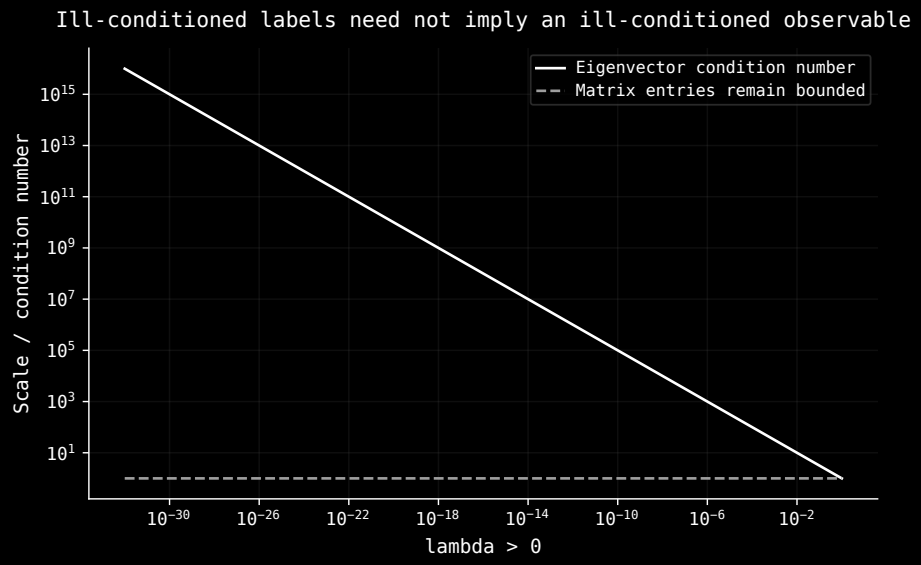


Figure 8.4: The exact eigenvector condition number diverges while the matrix entries stay bounded. The plot isolates representation conditioning from the behavior of the underlying family.

8.9.2 MG10: differentials turn the genus count into a constructive basis

For a smooth odd-degree hyperelliptic curve $y^2 = P_{2g+1}(x)$, the branch ledger gives genus g. To construct holomorphic differentials, test $\omega_k = x^k dx/y$. At a finite simple branch $x = a$, a local coordinate t has $x - a = t^2$ and $y = t\,u(t)$ with $u(0) \neq 0$. Then $dx/y = 2\,dt/u(t)$ is regular. At infinity, use $x = t^{-2}$ and $y = t^{-(2g+1)}v(t)$ with $v(0) \neq 0$:

$$\omega_k = -2t^{2g-2-2k}v(t)^{-1}\,dt.$$

Regularity selects $0 \leq k \leq g-1$. Their linear independence follows because a polynomial combination vanishing as a differential on a nonempty chart must have all coefficients zero. The dimension theorem for holomorphic differentials then proves that these g elements are a basis.

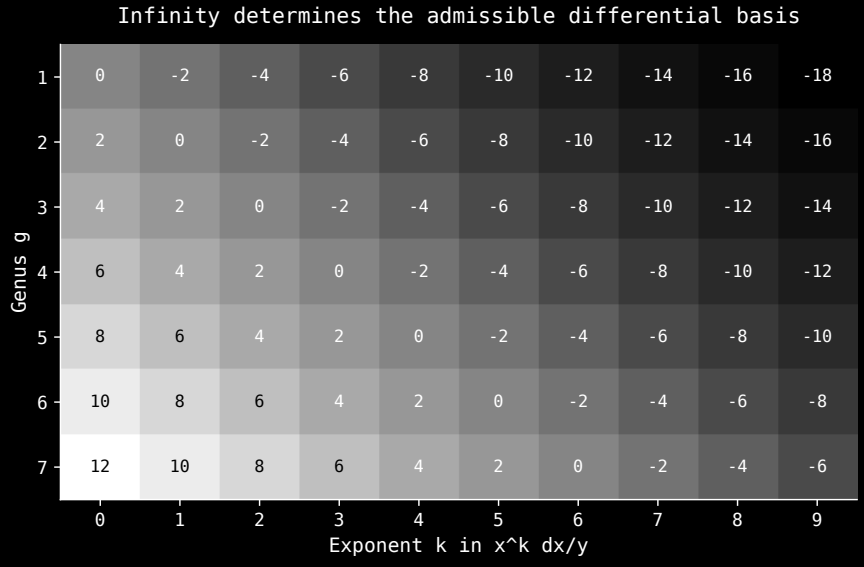


Figure 8.5: Orders in the infinity chart. Entries are exact integers; negative entries identify poles and therefore inadmissible holomorphic-basis candidates.

Traditional versus proposed task. A genus calculation alone returns one integer. The differential calculation returns a basis for period and energy problems. It is a richer output, not a faster method for the same integer. When the requested output is only genus, stop at the branch ledger. When a period matrix is needed, a genus value is insufficient; the chart computation is the necessary next step.

Exercise 8.13 (MG10a)**.** For degree nine, identify the basis and the infinity order of each element. What changes if two finite roots collide?

Solution. Here $g = 4$, with basis dx/y, $x\,dx/y$, $x^2\,dx/y$, $x^3\,dx/y$ and orders $6, 4, 2, 0$. If roots collide, the affine equation becomes singular and the smooth-cover hypotheses fail. One must normalize the curve and recompute its branch data; simply retaining the old four-dimensional basis is not justified.

8.9.3 MG11: deck averaging versus a generic least-squares projection

Let a finite group G act by measure-preserving transformations on (X, μ), and let U_g be the corresponding unitary pullback operators. The operator

$$P = \frac{1}{|G|} \sum_{g \in G} U_g$$

is the orthogonal projection onto invariant functions. Idempotence follows because each group element appears $|G|$ times in the double sum for P^2; self-adjointness follows by inversion in G. The range is invariant because left multiplication permutes the summands. A generic least-squares solve can also project onto a known invariant subspace, but the group representation supplies the projection without first choosing a basis for that subspace.

Metric and caveat. Direct averaging costs $|G|$ operator applications. This is attractive when group actions are cheap and the invariant basis is expensive, but it can be wasteful for a huge group or a low-dimensional already known subspace. Moreover, equal weights rely on the invariant measure and genuine group action. Averaging transforms with unequal physical sampling weights may produce the wrong projection. The correct comparison states the Hilbert inner product, the representation and the output subspace.

Exercise 8.14 (MG11a)**.** On $\mathbb{C}^3$ with the ordinary inner product, average the three cyclic coordinate rotations. Compute its matrix, rank and operator norm.

Solution. The average is $P = \frac{1}{3}\mathbf{1}\mathbf{1}^*$, with every entry $1/3$. It sends a vector to its coordinate mean repeated three times. Its range is the constant-vector line, so rank is one, and its eigenvalues are $1, 0, 0$. Thus its operator norm is one. The zero residual $P^2 - P$ alone would not identify this projection: the intended invariant range must also be specified.

8.9.4 MG12: charts and norms are additional data, not decorations

Suppose f_n converges uniformly to f on a neighborhood of a closed disk of radius r, and $\|f_n - f\|_\infty \le \varepsilon$. The first derivative at the center differs by at most ε/r by Cauchy's formula. The gain from holomorphy is a derivative certificate from a function bound, but the boundary distance r is essential. On a shrinking neighborhood, the same function-error bound may no longer imply derivative convergence.

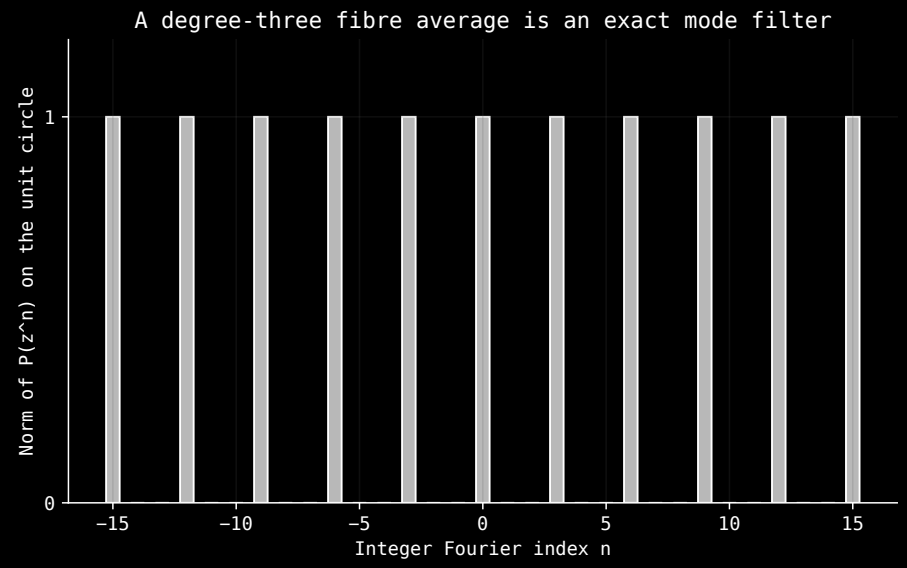


Figure 8.6: The exact norm of the degree-three fibre average on individual Fourier modes. This is an algebraic projection identity under normalized circle measure, not an approximate frequency-response calculation.

8.9.5 MG13: global compatibility cannot be compressed to a local inverse

A covering chart guarantees one inverse branch on one small open set. To transport it around a loop, one must record the induced subgroup or monodromy action. A sequence of individually correct local inverses can return to a different sheet. The correct global data are a basepoint, a fibre label and a homomorphism into a permutation group, with the relation constraints inherited from the base surface.

The traditional path-lifting proof is therefore retained as the baseline theorem. The proposed finite representation is a machine-readable passport and generator action; it is a computational interface to the theorem, not an alternative that eliminates topology. For a source with genus or punctures, omitting a fundamental-group generator silently weakens the global problem. This is the graduate analogue of checking only some edges of a contour certificate.

8.10 Research comparison: theorem strength, optimality and nonuniform limits

This extension supplies a genuine optimality proof for one sharply defined precision model. It also explains why that proof cannot be repurposed into an unrestricted claim of maximality for a mathematical corpus. The distinction is mathematical, not merely terminological: an optimization theorem requires a specified feasible class and objective.

8.10.1 MP8: an exact global integer precision optimum

Let $b_j \in \mathbb{Z}$, $b_j \geq 2$, and define

$$C(b) = b_1 + b_2 + b_3 + b_4, \quad E(b) = \frac{1}{64}2^{-b_1} + \frac{1}{16}2^{-b_2} + \frac{1}{4}2^{-b_3} + 2^{-b_4}.$$

The problem is to minimize C subject to $E \leq 1/64$. The weights arise from an ordered telescope; the lower caps model local admissibility. No measured runtime is equated with this sum-of-bits cost.

Theorem 8.1 (MP8: The optimum is twenty; uniform precision costs twenty-eight)**.** *The global minimum is* $C = 20$*, attained by* $(2, 4, 6, 8)$*. The least uniform choice is* $(7, 7, 7, 7)$ *with cost* 28*.*

Proof. The uniform inequality is $85 \cdot 2^{-b} \le 1$, giving $b \ge 7$. The displayed nonuniform vector has four contributions equal to $1/256$, hence is feasible with cost twenty. For a lower bound, set $a = (b_1, b_2 - 2, b_3 - 4, b_4 - 6)$. Then

$$E(b) = \frac{1}{64} \sum_{j=1}^{4} 2^{-a_j}, \qquad \sum a_j = C(b) - 12.$$

Convexity of 2^{-x} gives

$$\sum_{j=1}^{4} 2^{-a_j} \ge 4\, 2^{-(C(b)-12)/4}.$$

For feasibility this lower bound must not exceed one, so $C(b) \ge 20$. Equality requires all $a_j = 2$, which gives the stated vector. This proves the global result over all admissible integer vectors, without a finite search assumption. □

A second certificate. An exact integer dynamic program and independent enumeration of $13^4 = 28{,}561$ candidates reproduce the result. The enumeration bound is justified: the known feasible cost twenty and the three other lower caps imply that any no-more-expensive coordinate satisfies $b_j \le 14$. The analytic convexity proof and the finite computation check different failure modes. Neither replaces the domain assumptions from which the weights were obtained.

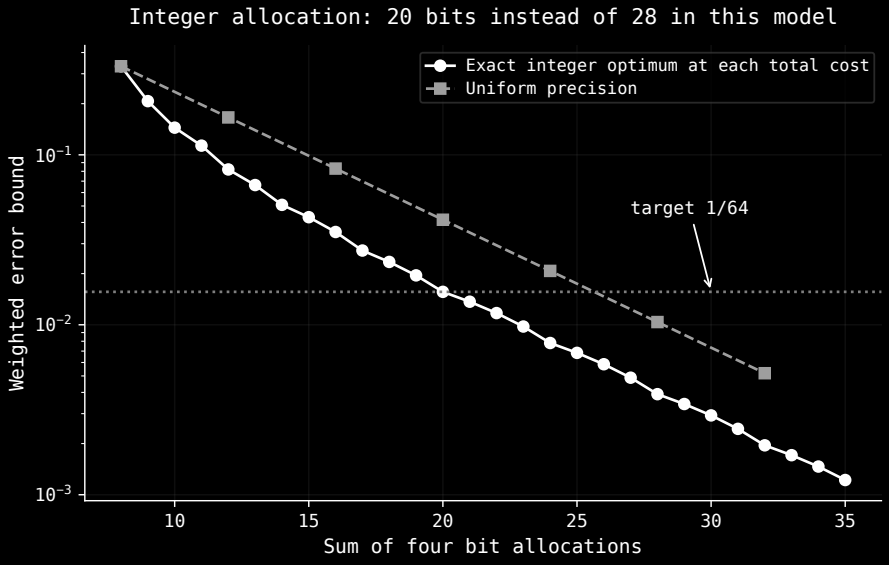


Figure 8.7: Exact minimum error at each integer total cost, compared with uniform precision. The analytic proof establishes the global optimum at the declared target; the finite frontier gives additional design information.

Exercise 8.15 (MP8a)**.** For m stages with weights $w_j = 2^{c_j}$, use convexity to derive a lower bound on $\sum b_j$ under $\sum_j w_j 2^{-b_j} \le E$. State when equality is compatible with integer and lower-cap constraints.

Solution. Set $a_j = b_j - c_j$. Convexity gives $\sum_j 2^{-a_j} \ge m 2^{-\sum a_j/m}$. Hence $\sum b_j \ge \sum c_j + m \log_2(m/E)$. Equality requires all $a_j = \log_2(m/E)$, and therefore all $b_j = c_j + \log_2(m/E)$ must be admissible integers above their caps. Otherwise the lower bound remains valid but need not be attained. A capped convex or integer optimization must determine the actual optimum.

8.10.2 MP9: error budgets must be matched to an observable

If an observable J is L_J-Lipschitz on the certified state tube, the state bound gives $|J(x) - J(\hat{x})| \leq L_J e$. For analytic derivatives the constant depends on boundary clearance. For a root branch it depends on separation from the discriminant. For a period near a degenerating neck it may be logarithmic or worse. A single "accuracy" field therefore loses essential information.

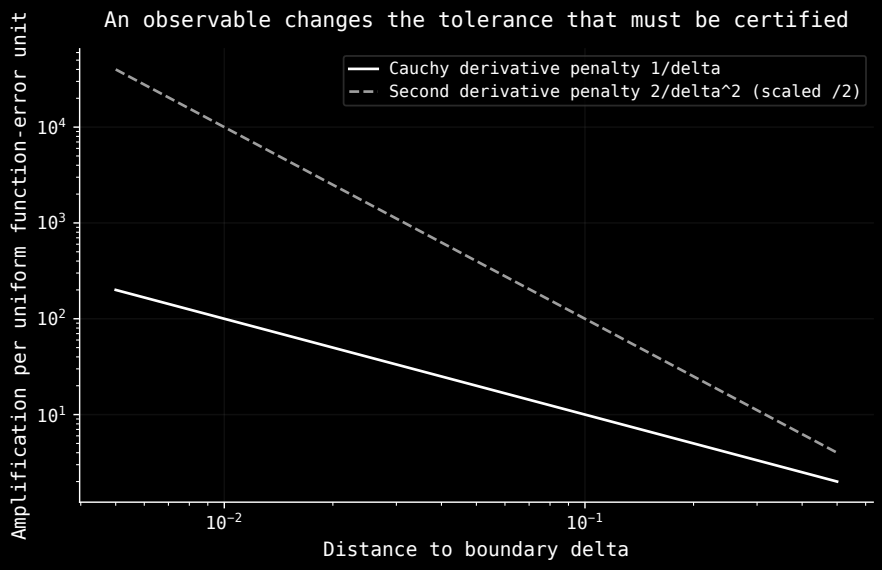


Figure 8.8: The cost of turning a function bound into a derivative bound depends on the distance to the boundary. The curves are theorem-derived amplification factors, not an empirical fit.

A quantitative design question. Suppose the final output is $f'(z_0)$, not $f(z_0)$. If the analytic neighborhood has radius r, a uniform function tolerance $r\varepsilon$ suffices for derivative tolerance ε. If r is reduced by a factor of sixteen, four additional binary precision units are required in this simple bound. A method that reports an unchanged function tolerance while claiming unchanged derivative accuracy has changed the contract silently.

Failure case. Consider $f_n(z) = z^n/n$ on the unit disk. Its uniform function norm tends to zero, but on the boundary $f_n'(1) = 1$. Interior compact subsets behave differently: at $|z| \leq \rho < 1$, $|f_n'| \leq \rho^{n-1} \to 0$. The distinction is exactly the positive boundary-distance requirement in the Cauchy estimate.

8.10.3 MP10: operator topology controls whether an index argument applies

For $0 < a < 1$, let P_a be the rank-one projection onto the normalized Hardy kernel k_a. The coefficient of z^n in k_a is $\sqrt{1-a^2}a^n$. For every fixed polynomial p, $\langle p, k_a \rangle \to 0$. Density of polynomials and $\|k_a\| = 1$ give weak convergence $k_a \rightharpoonup 0$. Consequently $P_a f = \langle f, k_a \rangle k_a \to 0$ for every fixed f: strong operator convergence. However $\|P_a\| = 1$ identically.

Comparison boundary. A finite collection of test vectors can diagnose strong behavior but cannot prove norm convergence of this moving projection. Thus a numerical claim that an index is preserved by a limit must specify the topology that the Fredholm theorem requires. More tests of fixed low-degree vectors do not repair the missing uniform quantifier. The operator-theoretic representation adds a counterexample to an inference that a root or boundary picture alone cannot settle.

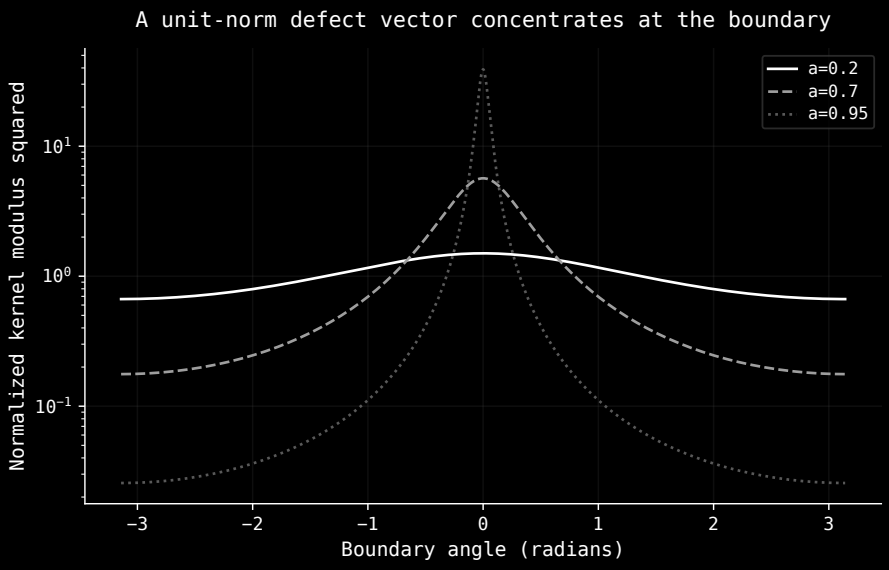


Figure 8.9: Unit-norm defect vectors concentrate at a boundary point. A disappearing interior picture is compatible with an unchanged operator norm.

Exercise 8.16 (MP10a)**.** Why is $\sup_{\|f\|=1} \|P_a f\| = 1$ consistent with $\|P_a f\| \to 0$ for every fixed f? Identify a maximizing vector depending on a.

Solution. The maximizing vector is $f = k_a$, which moves with a. Strong convergence fixes f before taking the limit; operator-norm convergence takes the supremum over f first. Interchanging these quantifiers is invalid. This is a direct functional-analytic instance of the difference between a finite test portfolio and a uniform theorem.

8.10.4 MP11: what a maximality theorem would have to quantify over

Let $\mathcal{U}$ be a finite, versioned set of required source obligations. A corpus D has a relation $W \subseteq \mathcal{U} \times \mathcal{A}$ linking obligations to inspectable artifacts. If each obligation has a qualifying witness, the coverage ratio is exactly one. The elementary inequality $|\{u : \exists a, (u, a) \in W\}| \leq |\mathcal{U}|$ then proves that the ratio cannot be increased *for that universe.* It does not prove that $\mathcal{U}$ contains every valuable mathematical extension, that a witness is correct, or that the exposition is pedagogically optimal.

The framework therefore exports separate fields for obligation coverage, written proof, executable witness, candidate formalization, actual kernel acceptance, nonredundancy diagnostics and measured comparison outcomes. A Boolean "done" flag would conceal these independent dimensions. Expanding the universe creates a new version instead of retroactively declaring the old completeness theorem false.

A rigorous but modest extremal claim. Fixed-universe coverage saturation is demonstrable. The precision problem MP8 has a genuine global optimum under an explicit objective. The statement "this is the largest possible useful mathematical document" has neither a finite feasible class nor an agreed objective. It is not asserted. The maximality audit is hardened by making that boundary part of the theorem, not by attaching a larger adjective to an unmeasured quantity.

8.10.5 MP12: research novelty needs a separate literature and replication gate

The new witness generator and examples are independently implemented combinations of classical convexity, Bernstein geometry and winding. The scope of originality asserted here is the document's derivation, the concrete integration and the reproducible artifact. A research-priority claim would require comparison against established certified root-counting, validated continuation and polynomial optimization methods beyond this focused corpus.

References

[1] J. E. Hutchinson. *Fractals and self similarity.* Volume 30, pp. 713–747. The general invariant-set framework is background; the special constructions are derived in the text.

[2] A. Hatcher. *Algebraic Topology.* Singular homology, cohomology and homotopy conventions.

[3] M. Lyubich. *Conformal Geometry and Dynamics of Quadratic Polynomials*, volumes I–II. Reference manuscript; §1.1.1, §3 and §4, with prerequisites cited locally.