FOUNDATIONS OF THE REALIZATION CALCULUS

Komorebi Space

Sixteen visual investigations in coherent realization

Visual guide

Komorebi P.

Publication and reuse

Komorebi Space: Coherent Realizations, Obstructions, and Descent
Komorebi P.

The original writing, mathematical figures, and editable writing sources in this publication are licensed under the Creative Commons Attribution 4.0 International license. You may share and adapt that material under its terms, with attribution to Komorebi P. and identification of changes. The complete license text accompanies the source package.

`https://creativecommons.org/licenses/by/4.0/`

The author's original Lean proofs, scripts, and other software in the accompanying release use the MIT license. Third-party software, fonts, referenced works, and other external material retain their own rights and notices. The grants above do not replace those terms. Font notices and the software license accompany the source package.

This publication preserves the accepted mathematical statements and proofs. Its source map distinguishes the formal counterparts from explanatory text, illustrations, and written proof routes. The current release report records the checks actually performed and the accepted scope of the formal library.

Contents

How to use this guide 3

01 Read the overlaps 4

02 Slice the extension fiber 5

03 Retain the comparison that matters 6

04 Find the missing equation 7

05 Reverse the order 8

06 Keep the discarded detail 9

07 Follow the remainder 10

08 Repair within the exact plane 11

09 How much can a preservation rule cost? 12

10 Small operators can have widely separated kernels 13

11 Can one residual history conceal two different outcomes? 14

12 Where does a family of data stop being realizable? 15

13 When does more correction stop improving the fit? 16

14 Can a better measured fit give a worse answer? 17

15 What happens when basis vectors nearly coincide? 18

16 Which stopping schedules preserve convergence? 19

How to use this guide

Start with a picture and change one ingredient. Keep a boundary value fixed. Remove an observation. Rotate a correction direction. Tighten a solver tolerance. The following investigations show what moves, what survives, and what can fail.

Each page supplies a small declared model, a figure to inspect, and a worked change to that model. The monograph references lead to the definitions and proofs. You can follow the pictures in sequence or choose a question from the contents.

Three things to identify

1. **The allowed objects.** Integer points, real vectors, matrices, and infinite sequences have different completion rules. A line drawn through integer markers does not create new integer solutions.
2. **The information held fixed.** An exact answer may exist but fail to preserve the prescribed data. Two pictures of exactness can describe different correction problems.
3. **The measured quantity.** A residual, a correction norm, a movement budget, and a coordinate norm are different numbers. Axis labels and metrics are part of the example.

Reading the graphics

Green marks the principal admissible object or successful route. Ochre and red distinguish comparisons, discrepancies, or excluded cases. Labels and legends identify what each color and line pattern measures. The spatial figures use genuine three-dimensional coordinates. Curves are evaluated from the formulas stated in the examples; they are not fitted to observations.

Finite windows of infinite sequences illustrate explicit formulas. An exact tail norm or a convergence condition comes from the accompanying mathematical argument, not from the end of a plotted line. Numerical drawing coordinates are rounded for display.

The companions

The monograph develops the general results and their hypotheses. The Lean guide identifies exact formal declarations and their scope. These visual investigations include concrete instances and illustrative variations; they are not individually asserted to be Lean formalizations. The source package contains the plotting scripts and the recorded checks for the worked values.

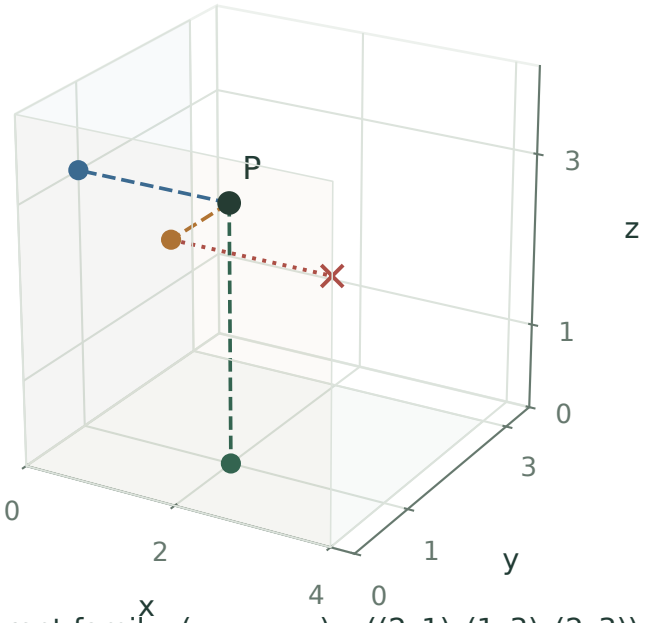
VISUAL INVESTIGATION 01

Read the overlaps

Take three views of a point in $\mathbb{R}^3$: its xy, yz, and xz coordinates. The picture follows $P = (2, 1, 3)$ into the three coordinate planes.

Which appearances belong to one point?

Shared coordinates are the overlap checks.



view	coherent	changed
xy	(2, 1)	(2, 1)
yz	(1, 3)	(1, 3)
xz	(2, 3)	(4, 3)

One changed x-value breaks the overlap.

Coherent family: $(xy, yz, xz) = ((2, 1), (1, 3), (2, 3))$.

Read the picture. Each pairwise view omits one coordinate. The three displayed pairs agree wherever a coordinate is shared, so they describe the same point. Replacing $(2, 3)$ in the xz view by $(4, 3)$ creates a conflict: the xy view still requires $x = 2$.

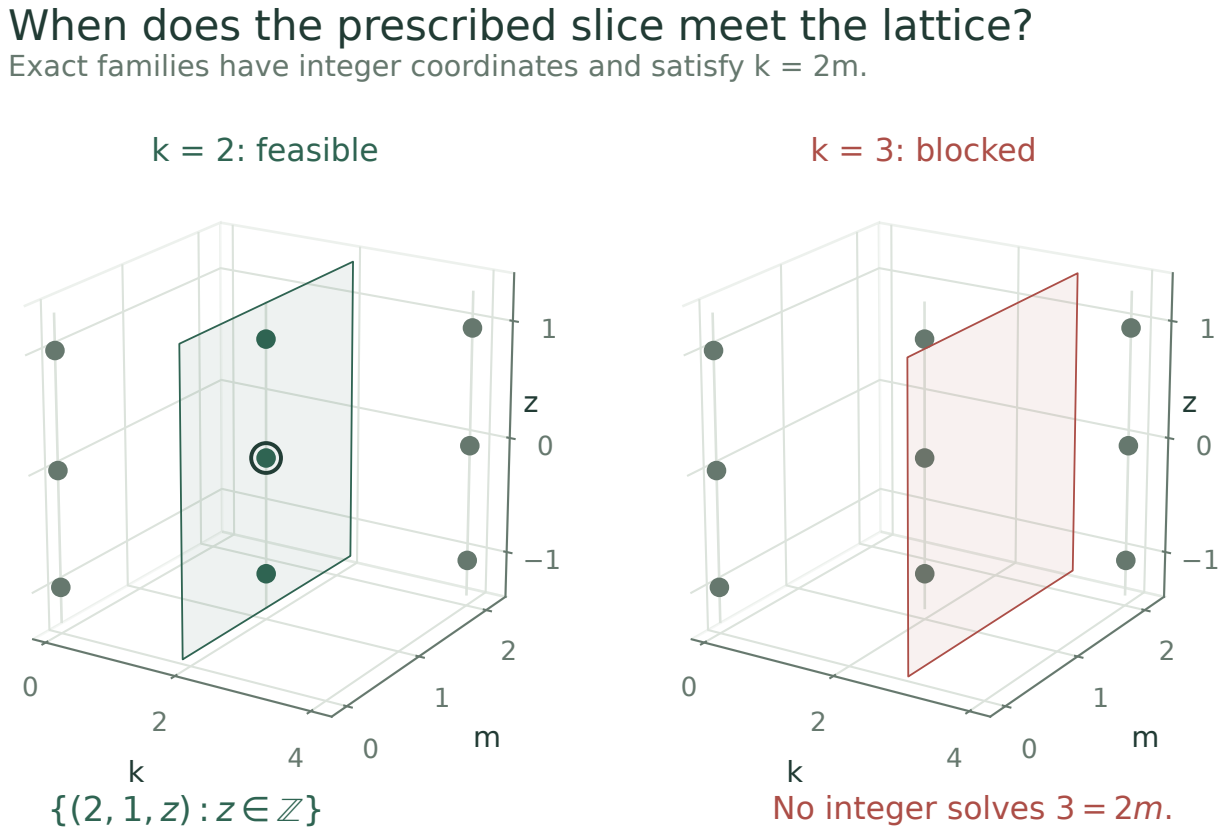
Change one thing. Keep only the xy appearance $(2, 1)$. The compatible points are the whole line $(2, 1, z)$. Restoring the yz appearance $(1, 3)$ selects $z = 3$.

Definitions, proofs, and general hypotheses: Monograph Sections 1, 2.

VISUAL INVESTIGATION 02

Slice the extension fiber

Use integer variables with $k = 2m$ and an independent coordinate z. Prescribe k, then look for lattice points in that slice. Only three values of z are displayed.



Read the picture. The slice $k = 2$ contains $(2, 1, z)$ for every integer z. The slice $k = 3$ misses the lattice because the required $m = 3/2$ is not an integer. The free vertical coordinate changes the number of answers, while the parity condition still determines whether any answer exists.

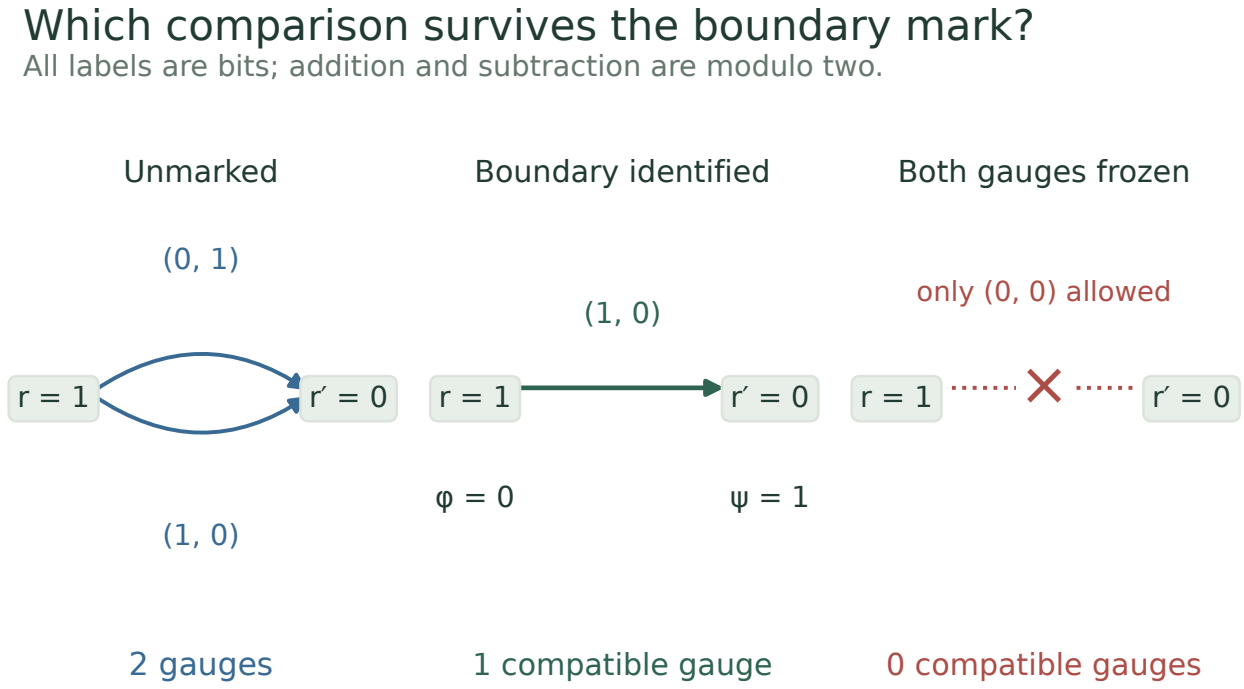
Change one thing. Prescribe $z = 0$ as well as $k = 2$. The circled point $(2, 1, 0)$ is now the unique answer. Prescribing $z = 0$ cannot rescue the blocked slice $k = 3$.

Definitions, proofs, and general hypotheses: Monograph Sections 3, 5.

VISUAL INVESTIGATION 03

Retain the comparison that matters

Duplicate a view and use the coefficient group C_2. A gauge (a_0, a_1) changes the transition label by $r' = a_0 + r - a_1$. Compare $r = 1$ with $r' = 0$. All arithmetic in this example is modulo two.



Read the picture. Two gauges relate the unmarked labels. The boundary marks $\phi = 0$ and $\psi = 1$ impose $\psi + a_0 = \phi$, leaving one compatible comparison. Freezing both representative gauges instead allows only $(0, 0)$, which cannot change the transition label. The three panels retain different data.

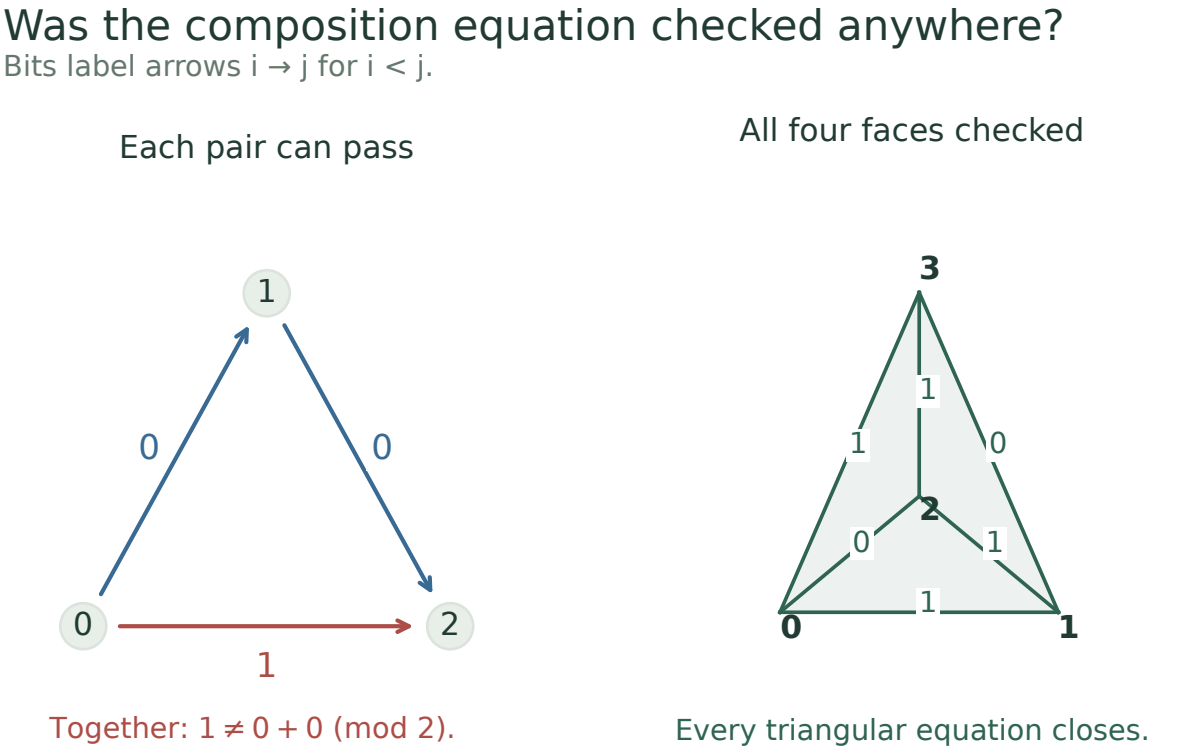
Change one thing. Change the target mark to $\psi = 0$ while keeping $\phi = 0$. The boundary equation then selects $(0, 1)$: a comparison still exists, but the chosen mark determines which one.

Definitions, proofs, and general hypotheses: Monograph Sections 15, 14.

VISUAL INVESTIGATION 04

Find the missing equation

On three ordered views, assign bit labels $c_{01} = 0$, $c_{12} = 0$, and $c_{02} = 1$. Each two-view patch contains its own arrow, but no patch contains the full composition.



Read the picture. The left triangle exposes the unchecked equation: the direct label is one, while the two-step label is zero. On the right, use four ordered views and all four three-view patches. Labels $c_{ij} = a_i + a_j$ from $a = (0, 1, 0, 1)$ satisfy every face equation.

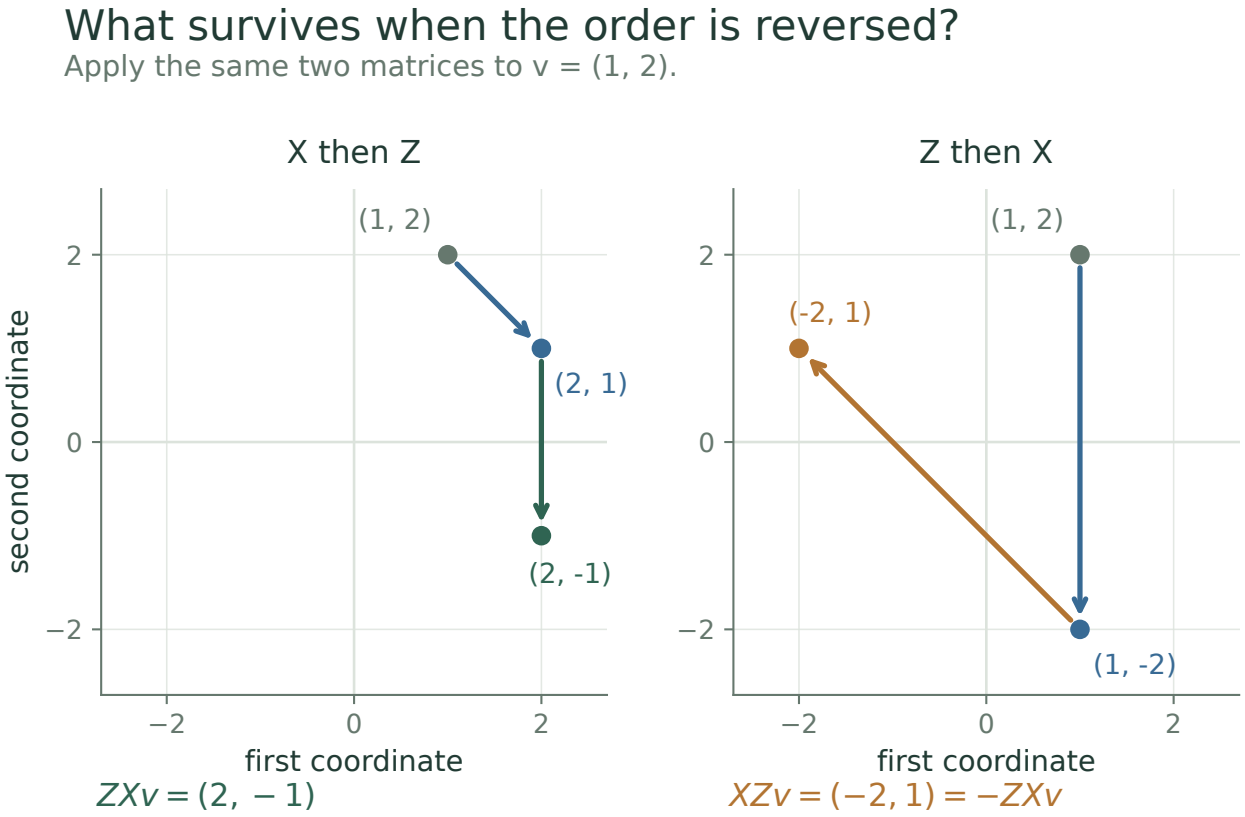
Change one thing. Replace the left direct label by zero. Its composition equation then holds. This repairs that example; the pair-patch cover still provides no general guarantee that composition was checked locally.

Definitions, proofs, and general hypotheses: Monograph Sections 12, 13.

VISUAL INVESTIGATION 05

Reverse the order

Label the coordinate swap X by $(1,0)$ and the second-coordinate reflection Z by $(0,1)$. Start at $v = (1,2)$ and follow both orders. Matrix products act from right to left.



Read the picture. Both routes use the same two-bit labels, whose sum is $(1,1)$. Their final vectors are negatives because $ZX = -XZ$. In this model the phase $c(x,y) = x_2y_1$ records that sign. The composition law remains associative; reversing these two factors is the operation that changes the result.

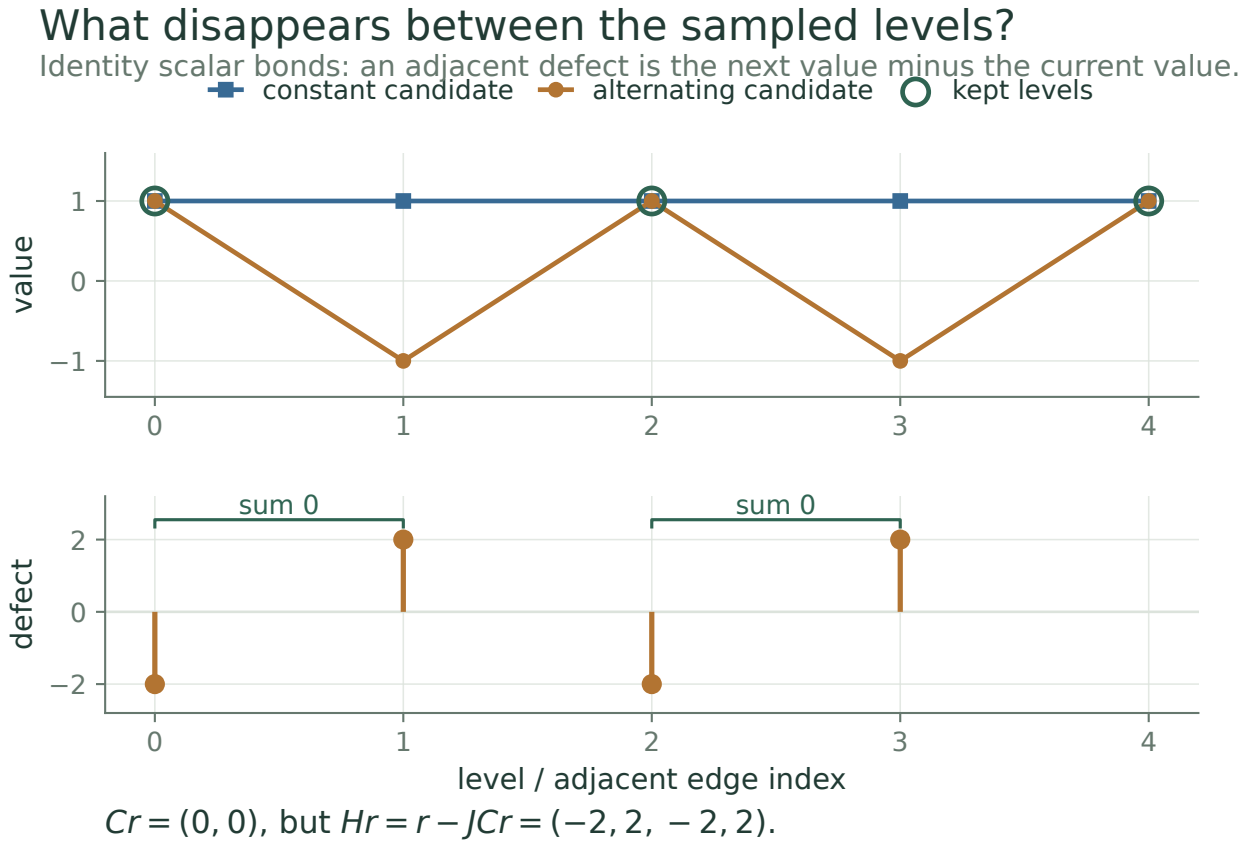
Change one thing. Attach a sign bit to the two-bit label. These two routes then end at different sign bits over the same label $(1,1)$, retaining the distinction that the four labels alone omit.

Definitions, proofs, and general hypotheses: Monograph Section 16.

VISUAL INVESTIGATION 06

Keep the discarded detail

Compare constant values $g_n = 1$ with alternating values $f_n = (-1)^n$ on identity scalar towers. Keep only even levels. The plot displays levels zero through four and their four adjacent defects.



Read the picture. The sampled values agree, and each pair of alternating defects sums to zero. The coarse residual therefore cannot distinguish the candidates. The detail does: with $J(c_0, c_1) = (0, c_0, 0, c_1)$, the zero coarse residual lifts to zero, leaving the entire alternating residual in $Hr = r - JCr$.

Change one thing. Restore the omitted odd values. The alternating candidate has defects $-2, 2, -2, 2$ while the constant candidate has zero defects. Exact coherent points are unaffected by this example of cancellation in raw candidates.

Definitions, proofs, and general hypotheses: Monograph Sections 6, 10.

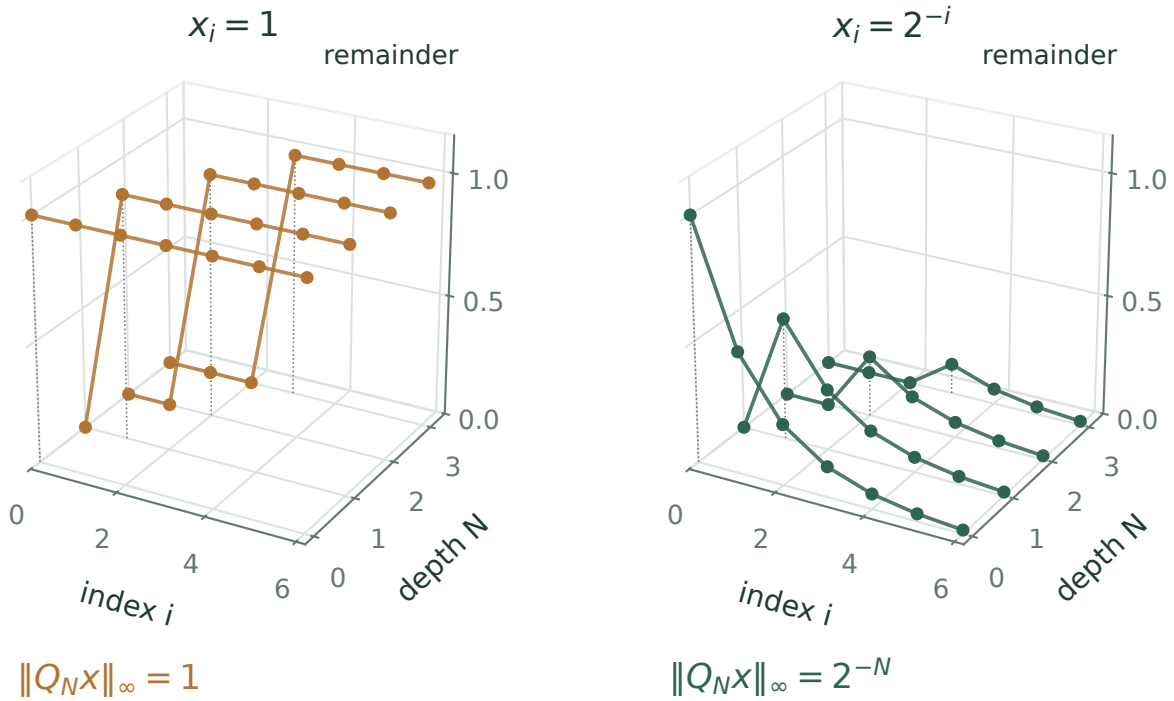
VISUAL INVESTIGATION 07

Follow the remainder

For the shift cascade, $Q_N x$ sets the first N coordinates to zero and leaves the tail in its original positions. Compare the constant sequence with $x_i = 2^{-i}$.

Does recovery move outward, or does the error shrink?

At depth N, the first N coordinates vanish; the rest retain their values.



$\|Q_N x\|_\infty = 1$ $\|Q_N x\|_\infty = 2^{-N}$

The formulas describe infinite tails; the plot shows only indices 0 through 6.

Read the picture. The recovered region grows in both panels: every fixed coordinate eventually becomes zero in the remainder. The constant tail keeps height one, so its uniform error never decreases. The geometric tail starts at 2^{-N} and then decreases further, giving exactly the norm printed beneath the right panel.

Change one thing. At depth $N = 3$, coordinates zero, one, and two are recovered in both examples. Their infinite remainder norms are nevertheless 1 and $1/8$, respectively.

Definitions, proofs, and general hypotheses: Monograph Section 10.

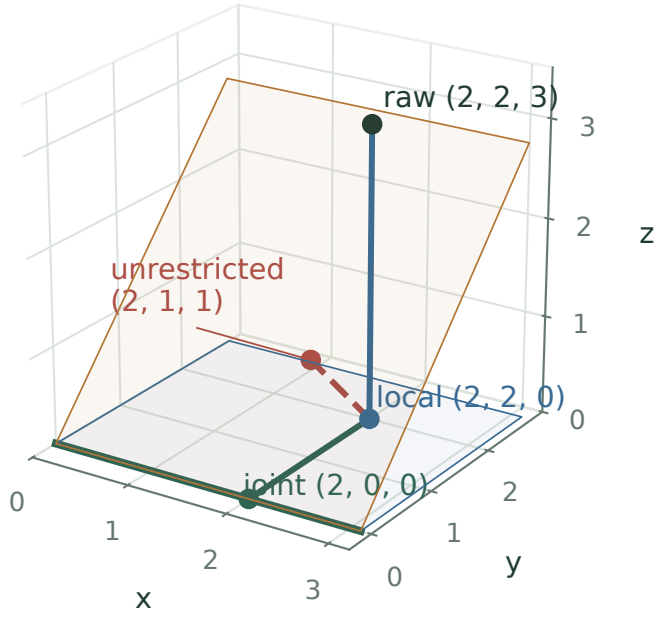
VISUAL INVESTIGATION 08

Repair within the exact plane

In Euclidean $\mathbb{R}^3$, take $D(x,y,z) = z$ and $C(x,y,z) = y - z$. Starting at $(2,2,3)$, first project to the locally exact plane $z = 0$. The two planes encode different constraints on the same raw point.

Does the second repair preserve the first equation?

Local exactness: z = 0. Matching: y = z. Joint exactness: the x-axis.



Inside $z = 0$: squared moves $9 + 4 = 13$. Unrestricted: matching holds, but $z = 1$.

Read the picture. The local correction reaches $(2,2,0)$. Solving the matching problem inside that plane reaches $(2,0,0)$, with perpendicular squared movements 9 and 4. Projecting instead onto the entire plane $y = z$ reaches $(2,1,1)$: matching is repaired, but the local equation has been broken.

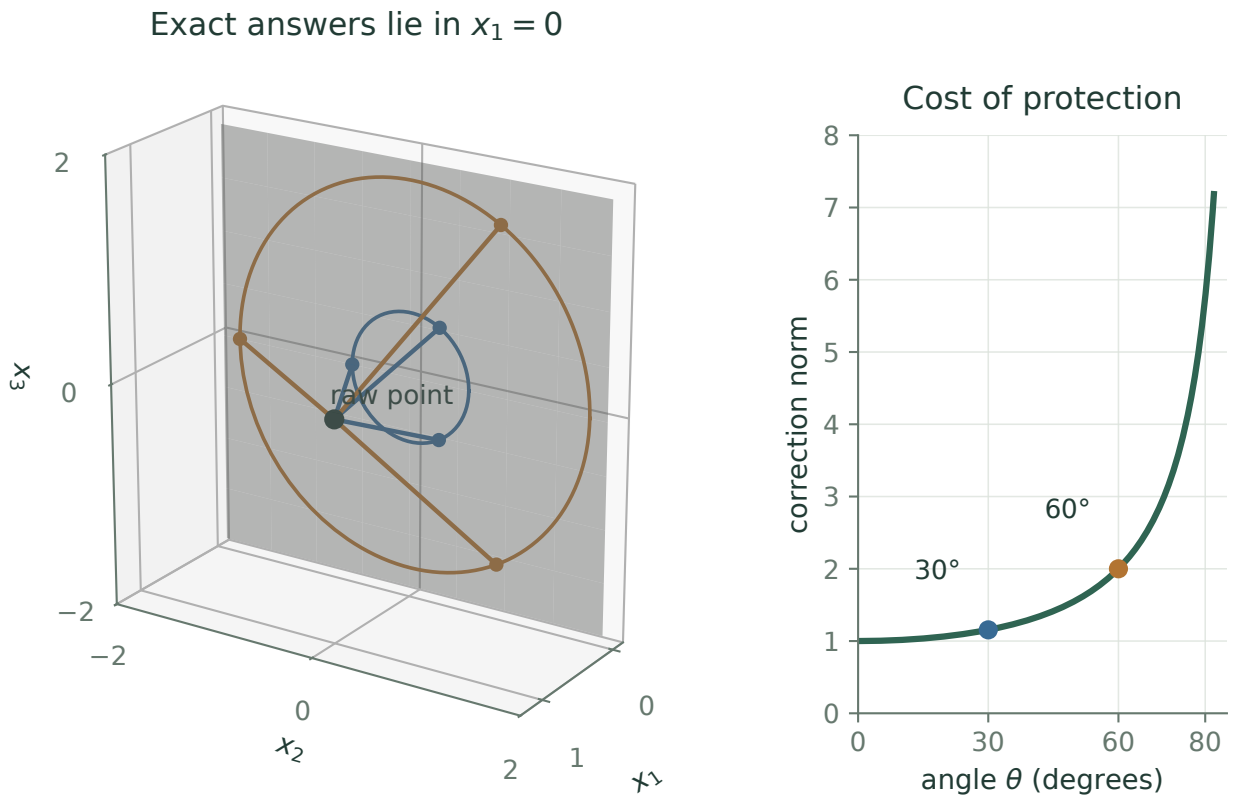
Change one thing. At the unrestricted answer, evaluate both defects: $C(2,1,1) = 0$ and $D(2,1,1) = 1$. The joint answer has both defects zero and total squared correction 13.

Definitions, proofs, and general hypotheses: Monograph Section 20.

VISUAL INVESTIGATION 09

How much can a preservation rule cost?

Keep the raw point $x = (1, 0, 0)$ and equation $D(x) = x_1$ fixed. Vary the allowed unit correction direction $v = (\cos\theta, \sin\theta\cos\phi, \sin\theta\sin\phi)$, with protected observable $p = I - vv^{\mathsf{T}}$. All three coordinates use the Euclidean metric.



Each direction defines its own protected observable.

Read the picture. For $0 \le \theta < \pi/2$, the unique feasible correction is $a = v/\cos\theta$. Its norm is $\sec\theta$; its endpoint $x - a$ lies in the shaded exact plane. Azimuth changes the answer without changing its cost. At $\theta = \pi/2$, allowed changes cannot alter x_1, so this raw point becomes infeasible.

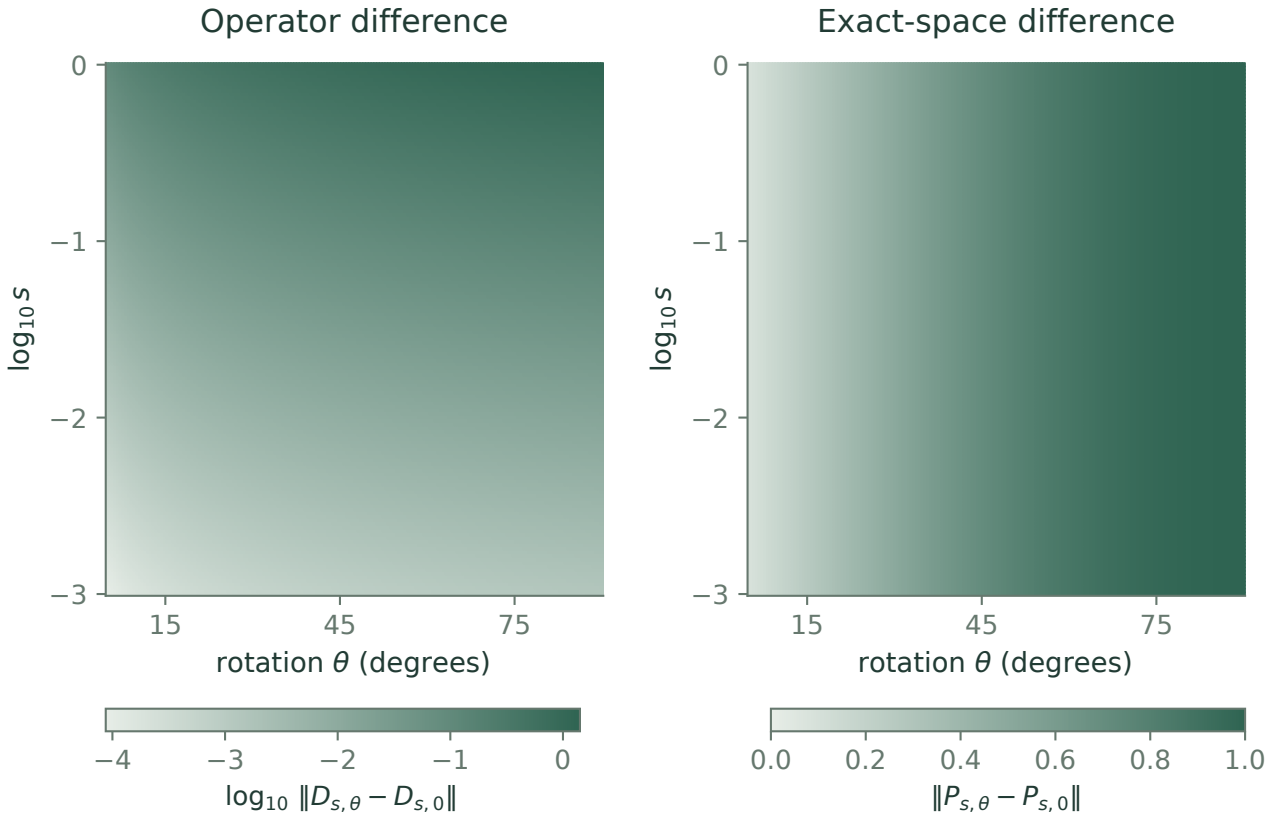
Change one thing. Compare 30° and 60°: the costs are $2/\sqrt{3}$ and 2. Turning the correction toward the exact plane increases its length; every plotted line preserves its own declared observable exactly.

Definitions, proofs, and general hypotheses: Monograph Section 18.

VISUAL INVESTIGATION 10

Small operators can have widely separated kernels

Compare the row operators $D_{s,\theta} = s(\cos\theta \;\; \sin\theta)$ and $D_{s,0} = s(1 \;\; 0)$, with $s > 0$ and $0 \le \theta \le \pi/2$. Their kernels are lines in one Euclidean plane. Sweep both scale and angle; the plot starts at 5°.



Read the picture. The operator difference is $2s\sin(\theta/2)$, while the projector difference is $\sin\theta$ for the angles shown. Shrinking s changes only the left panel. Each nonzero operator has least residual gain $1/s$: the scale that hides the operator difference also weakens its stability certificate.

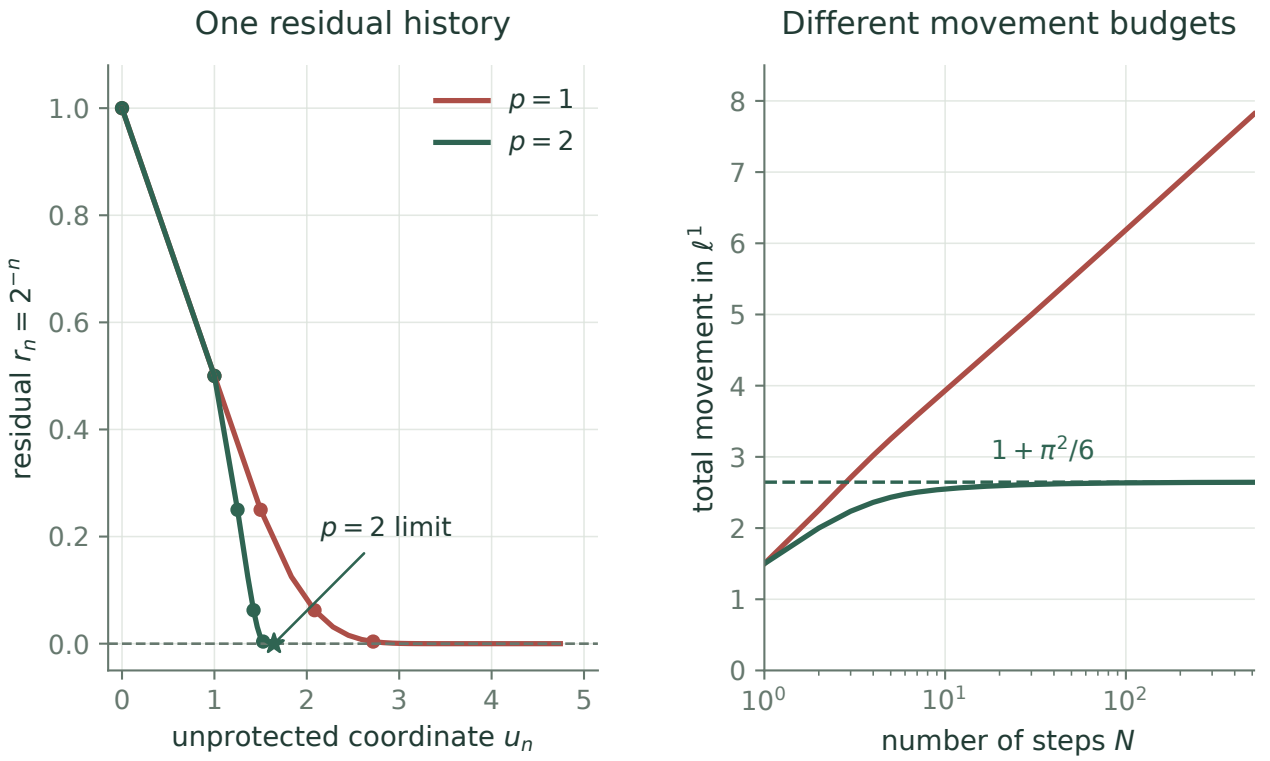
Change one thing. At $\theta = 60°$, reducing s from 1 to 10^{-3} reduces the operator difference by a thousand. The projector gap stays $\sqrt{3}/2$. The zero-scale operator has a different kernel and is excluded.

Definitions, proofs, and general hypotheses: Monograph Sections 18, 21.

VISUAL INVESTIGATION 11

Can one residual history conceal two different outcomes?

Let $u_n^{(p)} = \sum_{k=1}^{n} k^{-p}$ and $f_n^{(p)} = (s, u_n^{(p)}, 2^{-n})$. Protect the coordinate s, and measure candidate movement with the ℓ^1 metric. Both paths start at the same point and have the same residual.



Read the picture. The movement through step N is exactly $\sum_{k=1}^{N} k^{-p} + 1 - 2^{-N}$. For $p = 2$ it converges to $1 + \pi^2/6$, giving an exact limit. For $p = 1$ the unprotected coordinate diverges. These conclusions use the convergent square series and divergent harmonic series, not extrapolation from the finite curves.

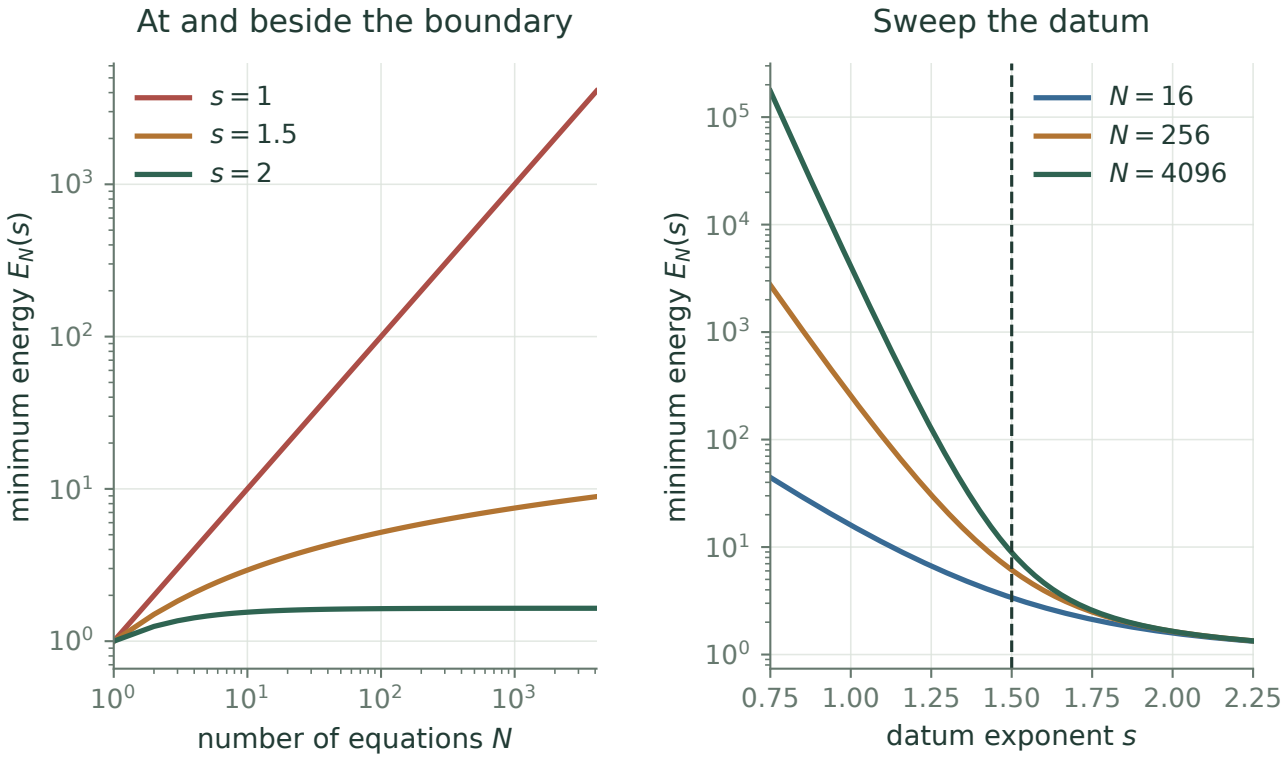
Change one thing. After one step, both paths have residual $1/2$. Nevertheless their eventual behavior differs. The finite residual table alone cannot decide convergence: a summable bound must control movement in the candidate space.

Definitions, proofs, and general hypotheses: Monograph Section 24.

VISUAL INVESTIGATION 12

Where does a family of data stop being realizable?

Keep $Ta = (a_n/n)$ on ℓ^2 fixed, and vary the datum $r_n^{(s)} = n^{-s}$ for $s > 1/2$. The least correction satisfying its first N equations has energy $E_N(s) = \sum_{n=1}^{N} n^{2-2s}$.



Read the picture. An exact preimage must have coordinates $a_n = n^{1-s}$. Thus the square-series test gives realizability precisely for $s > 3/2$. At $s = 3/2$, the energy is harmonic and still diverges. Every finite problem remains solvable on both sides; only a uniform energy bound can pass to the Hilbert realization.

Change one thing. Predict the middle curve: at $s = 3/2$ its energy lies between $\log(N+1)$ and $1 + \log N$. Slow growth is still unbounded growth. The plot evaluates finite sums through $N = 4096$.

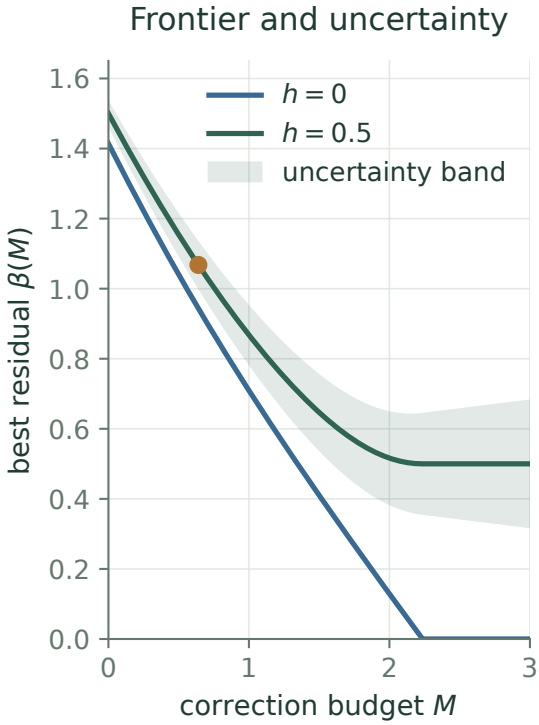
Definitions, proofs, and general hypotheses: Monograph Sections 26, 27, 30.

VISUAL INVESTIGATION 13

When does more correction stop improving the fit?

Let $T(a,b) = (a, b/2, 0)$ and $r_h = (1, 1, h)$, with Euclidean norms. The correction ball maps to a flat ellipse. Vary its radius M and compare targets at heights $h = 0$ and $h = 1/2$.





Read the picture. The regularized corrections $a_\alpha = (1/(1+\alpha), 2/(1+4\alpha))$ trace the exact frontier at their own budgets. At $M = \sqrt{5}$, the horizontal mismatch disappears; the residual floor is $|h|$. The shaded band allows data error $\delta = 0.03$ and operator error $\eta = 0.05$, using width $\delta + \eta M$ on one fixed correction space.

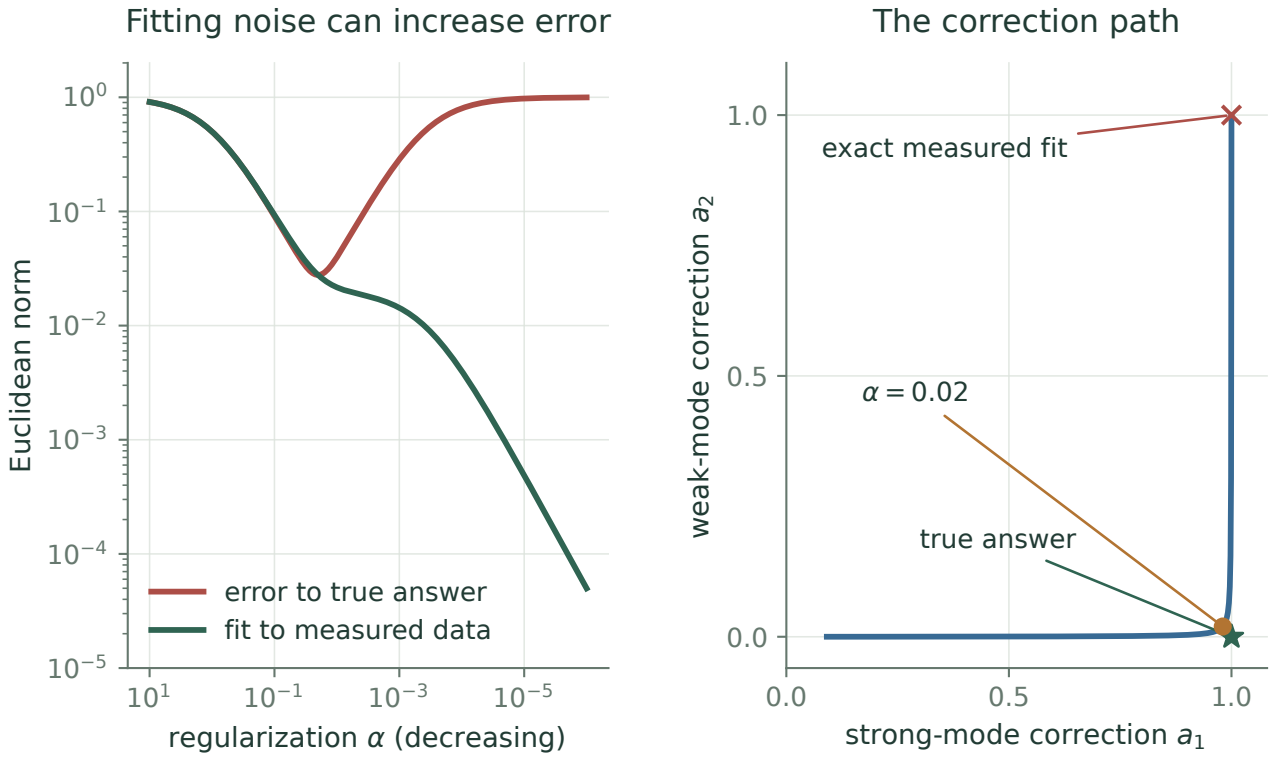
Change one thing. At the marked point, $\alpha = 1$, $a = (1/2, 2/5)$ and $M = \sqrt{41}/10$. For $h = 1/2$, the residual is $(1/2, 4/5, 1/2)$; its normalized direction certifies the entire budget ball.

Definitions, proofs, and general hypotheses: Monograph Sections 28, 29.

VISUAL INVESTIGATION 14

Can a better measured fit give a worse answer?

Use $T = \mathrm{diag}(1, 0.02)$ and true correction $a^\dagger = (1, 0)$. Add the deterministic data error $(0, 0.02)$, entirely in the weak mode. The measured datum is $(1, 0.02)$; all norms are Euclidean.



Read the picture. The exact regularized answer is $a_\alpha = (1/(1+\alpha), 0.0004/(0.0004+\alpha))$. As α decreases, measured fit improves throughout, but the weak-mode error eventually grows. At zero-parameter limit the measured equations are fitted exactly by $(1, 1)$, a distance one from the true answer. This is a finite two-mode calculation.

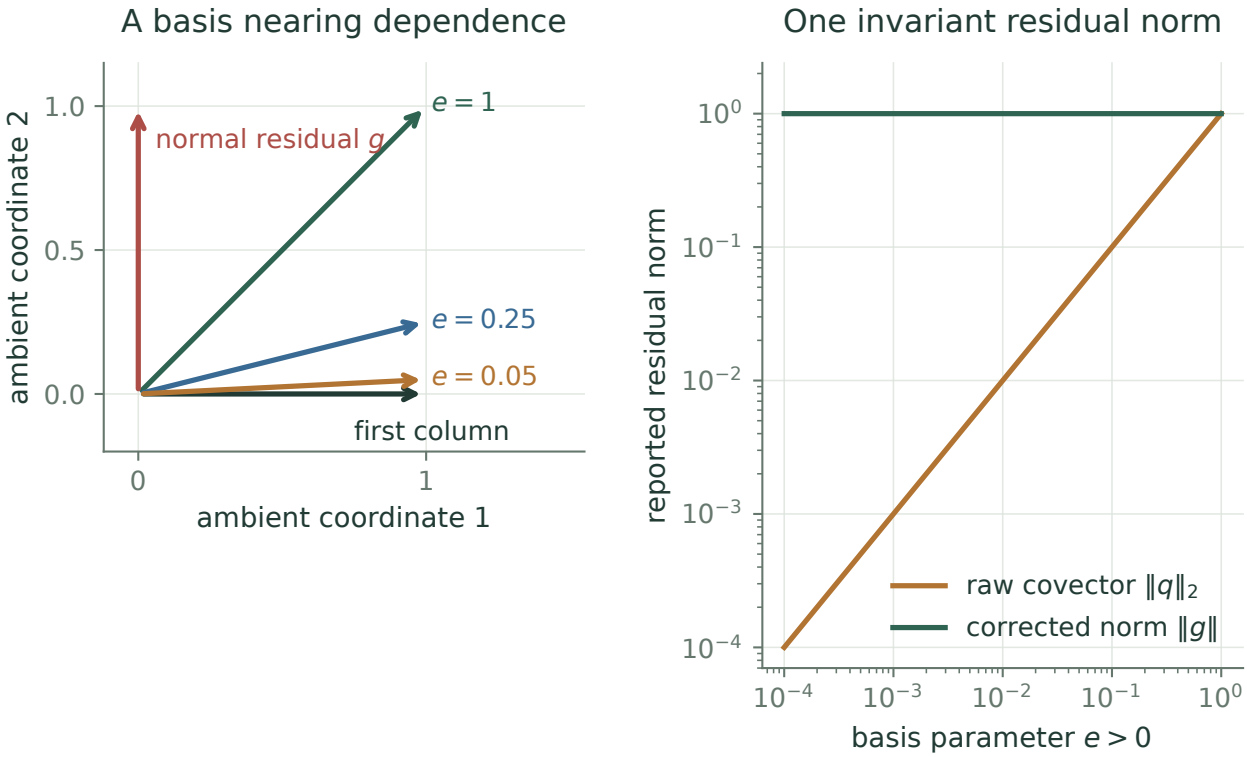
Change one thing. At $\alpha = 0.02$, the answer is $(50/51, 1/51)$ and its true error is $\sqrt{2}/51$. Smaller measured residual alone therefore does not select the most accurate correction for this noisy datum.

Definitions, proofs, and general hypotheses: Monograph Sections 29, 32.

VISUAL INVESTIGATION 15

What happens when basis vectors nearly coincide?

Keep $T = I$, $\alpha = 1$, datum $r = (0, -1)$ and proposal $a = 0$ fixed in $\mathbb{R}^2$. Use the full basis $B_e = \left(\begin{smallmatrix} 1 & 1 \\ 0 & e \end{smallmatrix}\right)$ for $e > 0$. Its columns approach the same direction.



Read the picture. The actual normal residual is always $g = (0, 1)$, but its coordinate covector is $q = B_e^{\mathsf{T}} g = (0, e)$. With $G = B_e^{\mathsf{T}} B_e$, the invariant square is $q^{\mathsf{T}} G^{-1} q = 1$. At $e = 0$ the columns no longer span the permitted space, so the full-space certificate formula no longer applies.

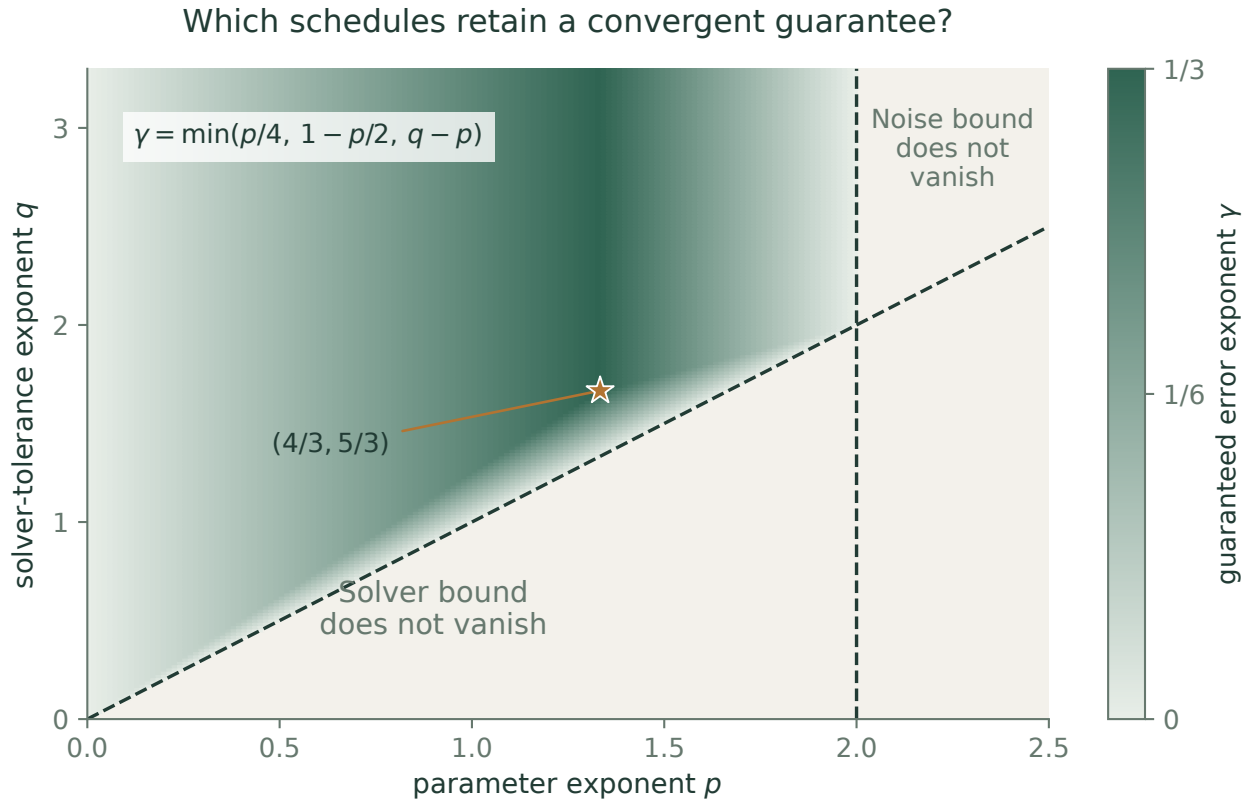
Change one thing. For $e = 0.01$, the raw coordinate norm is 0.01, although the invariant norm is one. The optimizer remains $(0, -1/2)$, and the unchanged proposal still has optimizer error $1/2$.

Definitions, proofs, and general hypotheses: Monograph Sections 31, 32.

VISUAL INVESTIGATION 16

Which stopping schedules preserve convergence?

For the attainable diagonal datum $r_n = 1/n^2$, choose $m = \lceil \delta^{-p/2} \rceil$, $\alpha = m^{-2}$, and solver tolerance $\varepsilon = \delta^q$, where $0 < \delta \leq 1$. Vary the positive exponents p, q instead of fixing one noise schedule.



Read the picture. The proved bias, noise and solver bounds have orders $\delta^{p/4}$, $\delta^{1-p/2}$ and δ^{q-p}. All vanish in the colored region $0 < p < 2, q > p$. Its color records their slowest exponent. The uncolored region loses this sufficient guarantee; it does not assert failure of every actual solver.

Change one thing. The marked choice $(p, q) = (4/3, 5/3)$ balances all three exponents at $1/3$. Choosing $q = p$ instead leaves a nonvanishing solver bound. The ceiling is retained in the underlying parameter rule.

Definitions, proofs, and general hypotheses: Monograph Sections 30, 32.