# Function approximation and numerical integration

Construct approximations from the information actually available. Derive their error and observe what changes when the function, data or sampling rule violates an assumption.

## Contents


1. **Approximating a resonance with a rational function**
   A smooth response on a measured interval can be governed by a nearby pole. Derive a low-order Padé fit, construct a barycentric rational approximant from samples and check whether a predicted pole survives independent data.
2. **Integrating cosine interpolants on nested grids**
   Chebyshev points turn polynomial interpolation into a cosine series. Derive the exact moments of each basis polynomial, explain aliasing and find a case where two consecutive quadrature grids confidently agree on the wrong answer.
3. **Integrate the oscillations analytically and sample the envelope**
   A sensor model asks for many Fourier integrals of a slowly changing envelope. Direct trapezoids can return a full-scale signal where the true integral nearly cancels. Derive a Filon rule from exact oscillatory moments and find where it needs a different local model.
4. **Spend random samples where they reduce uncertainty**
   A Monte Carlo estimate can be unbiased and still too noisy to be useful. Derive control variates, allocate a fixed budget across strata, and distinguish variance reduction from an underestimated standard error.
5. **How many neighbouring samples belong in a slope estimate?**
   A wider local fit suppresses measurement noise, yet it can smooth away real changes in slope. Derive the competing errors, then use noisy synthetic data to find which differentiator wins at a chosen spacing.
6. **What can smoothness buy at a fixed polynomial degree?**
   Prove density for every continuous signal, derive a faster rate for smooth signals and identify the persistent overshoot when a jump violates the assumptions.
7. **When should the denominator be fitted too?**
   Develop a minimax certificate for rational approximants, find the pole of a measured response and identify the obstruction that appears when a trial denominator vanishes inside the interval.

## 8. Measuring an area integral with finitely many samples

Build deterministic multidimensional rules from moments and geometry. Compare tensor products, symmetric triangle rules and sparse combinations at a fixed number of evaluations.

# Function approximation and numerical integration · 1. Approximating a resonance with a rational function

*A smooth response on a measured interval can be governed by a nearby pole. Derive a low-order Padé fit, construct a barycentric rational approximant from samples and check whether a predicted pole survives independent data.*

## Match series coefficients without requiring a polynomial model

Let $f(x) = c_0 + c_1 x + c_2 x^2 + \cdots$ near zero. A polynomial Taylor approximation fixes coefficients but cannot represent a finite pole. A type $[m/n]$ Padé approximant $P_m(x)/Q_n(x)$ uses polynomials of degree at most $m, n$, and fixes the arbitrary common scaling by $Q_n(0) = 1$ where possible. Require

$$Q_n f - P_m = O(x^{m+n+1}).$$

For powers above $m$ the unknown numerator cannot contribute; they give linear equations $\sum_{j=0}^{n} q_j c_{k-j} = 0,\ k = m+1, \ldots, m+n$ for the denominator, with missing negative-index coefficients taken zero. Lower coefficients then give

$$p_k = \sum_{j=0}^{\min(k,n)} q_j c_{k-j}.$$

Although the desired quotient is nonlinear in its coefficients, series matching has reduced the construction to linear algebra.

For $f = e^{-x}$, matching $1 - x + x^2/2$ gives

$$[1/1] = (1 - x/2)/(1 + x/2).$$

Its expansion agrees through second degree; at $x = 1$ it gives $1/3$, compared with $e^{-1} = 0.3679$ and Taylor degree two $0.5$. But the rational has a pole at $-2$ even though the exponential is entire, and becomes negative for $x > 2$. A good local approximation does not establish that every pole corresponds to physical resonance or every extrapolated value respects a known constraint such as positivity. The model's domain and validation data determine the legitimate interpretation.

## Interpolate at selected support points while fitting other samples

When samples, rather than derivatives at one origin, are available, a barycentric form is convenient:

$$r(z) = \frac{\sum_{j=1}^{m} w_j f_j/(z - z_j)}{\sum_{j=1}^{m} w_j/(z - z_j)}.$$

At a support node use the exact stored $f_j$, interpreting the apparent singularity by its limit when $w_j \neq 0$. Common rescaling of all weights leaves $r$ unchanged. For another data pair $(z_i, f_i)$, a small

cross-multiplied discrepancy requires $\sum_j w_j(f_i - f_j)/(z_i - z_j) \approx 0$. With several unsampled points this is a homogeneous least-squares problem for the weights; impose unit norm and take the right singular vector belonging to the smallest singular value. The adaptive AAA algorithm selects an additional support point where current error is large and rebuilds this fit until a stopping criterion is met.

The cross-multiplied error and true value error differ by the rational denominator. A very small denominator can amplify a small algebraic residual, so validate the fitted quotient directly on independent points and check for poles in the intended domain. Closely paired numerator and denominator roots can almost cancel. Such a weak-residue pole–zero pair can encode numerical noise rather than a well-resolved physical mode. Testing another grid, another noise realisation and the residue is more informative than printing every computed root with equal confidence.

## Compare local and global descriptions on a known response

Take $f(x) = 1/(1 + \beta x^2)$ for $\beta > 0$. Its poles are at $x = \pm i/\sqrt{\beta}$, outside the real line but at a finite distance from the origin. The even Taylor series is geometric in $-\beta x^2$. Through power $x^{2m}$ the relative error magnitude is $(\beta x^2)^{m+1}$. It converges for $|x| < 1/\sqrt{\beta}$, fails to converge outside that radius and can show oscillatory partial values across a perfectly smooth real measurement interval. Its type $[0/2]$ rational reconstruction from $c_0 = 1, c_1 = 0, c_2 = -\beta$ is exact in noiseless arithmetic. That success rewards a model family matching the analytic structure; it does not imply unknown response functions must have just this pair of poles.

## Compare a Taylor and a rational continuation

The plot holds the rational response above as known truth. Vary the pole distance through $\beta$, then change the polynomial truncation. Watch where increasing the expansion order improves agreement and where it worsens it. A noisy curvature estimate changes the fitted rational pole locations; use the readout to compare the pole error with the predicted local change. The experiment intentionally exposes the truth, which is unavailable for a real instrument and must be replaced by out-of-sample checks.

Open this experiment online

### Try the calculation

| | |
|---|---|
| Curvature β | 1.2 |
| Highest even power / 2 | 4 |
| Curvature bias (%) | 5 |

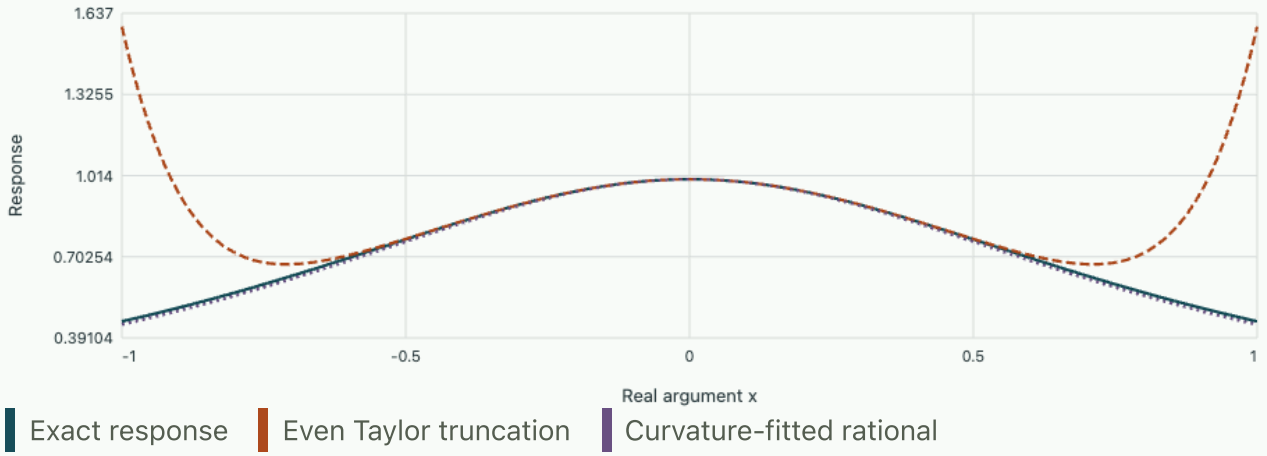


**True pole distance 0.913, fitted pole distance 0.891. The Taylor convergence boundary is |x| = 0.913.**

## Test the assumed response law on an unused reading

A value away from the centre tests curvature without differentiating noisy samples. Under the assumed normalisation $f(0) = 1$, solving $y = (1 + \beta x^2)^{-1}$ gives

$$\beta = (1/y - 1)/x^2.$$

At $x = 1/2$, a reading $y = 10/13$ yields $\beta = 1.2$. Before including any other reading, predict $f(1) = 5/11$. A measured value of $0.50$ differs by about $0.0455$. More Taylor coefficients of the fitted rational function all inherit the same discrepancy. They cannot establish an independent agreement.

The derivative of the inferred parameter with respect to a measured value is $-1/(x^2 y^2)$. At the calibration point its magnitude is $6.76$. An absolute measurement uncertainty of $0.005$ propagates to about $0.034$ in $\beta$, and then about $0.0070$ in the predicted reading at one. A disagreement of $0.0455$ would warrant examining calibration of the unused reading, the assumed normalisation and the model itself. Comparing two computed continuations built from the same reading would miss this test.

# Tutorial · separate a useful fit from a physical pole

## Infer a relaxation constant from two Taylor coefficients

A response starts $f(s) = 4 - 6s + 9s^2 + \cdots$. Construct a type $[0/1]$ fit and test the next coefficient rather than using it to determine the parameter.

▼ **Work through the solution**

The constant numerator is 4 and $4/(1 + qs) = 4 - 4qs + 4q^2s^2 + \cdots$. Matching slope gives $q = 1.5$ and predicted quadratic coefficient 9, consistent with the extra observation. The pole is $-2/3$ in the units of $s$. A further measured coefficient or a separate frequency interval can test the single-pole model. Agreement of three coefficients alone does not preclude additional weak modes.

---

## Calculate a pole from two barycentric nodes

Use nodes $z_1 = 0, z_2 = 1$, samples $1, 1/2$ and weights $1, -2$. Simplify the quotient and locate its pole. What happens at the support nodes?

▼ **Work through the solution**

Multiplying numerator and denominator by $z(z - 1)$ gives numerator $(z - 1) - z = -1$ and denominator $(z - 1) - 2z = -z - 1$. Hence $r = 1/(1 + z)$, with pole at $-1$. The original sums are singular separately at 0 and 1, but their limits give exactly 1 and $1/2$. Direct numerical evaluation should return stored values at support nodes rather than form infinity divided by infinity.

---

## Estimate the sensitivity of an inferred pole distance

The sampled response yields curvature coefficient $-\beta = -4.0 \pm 0.4$. Infer the distance of the conjugate poles and its linearised relative uncertainty. Would this prove the real system has that exact physical singularity?

▼ **Work through the solution**

The fitted distance is $\beta^{-1/2} = 0.50$, with relative uncertainty approximately $\frac{1}{2}(0.4/4.0) = 5\%$, or absolute 0.025. A systematic error in the assumed response shape can dominate this local parameter uncertainty. Test independently sampled values farther from zero, compare alternative rational orders, and examine whether pole positions and residues persist. Noise estimates conditional on one chosen model are not a proof of analytic structure.

---

# Further reading

- Nakatsukasa, Sète and Trefethen · the original AAA rational approximation algorithm
- Course investigation · approximation error across an interval
- Investigate interpolation error in a heat integral

# Function approximation and numerical integration · 2. Integrating cosine interpolants on nested grids

*Chebyshev points turn polynomial interpolation into a cosine series. Derive the exact moments of each basis polynomial, explain aliasing and find a case where two consecutive quadrature grids confidently agree on the wrong answer.*

## Use a change of coordinate to simplify both the grid and the basis

The Chebyshev polynomial $T_k(x)$ is defined by $T_k(\cos\theta) = \cos(k\theta)$. For $N+1$ nodes $x_j = \cos(j\pi/N)$, sampling a degree-$N$ polynomial in Chebyshev coordinates becomes sampling a finite cosine series at equally spaced angles. A discrete cosine transform computes the unique interpolant's coefficients without a Vandermonde solve. Clenshaw–Curtis quadrature integrates that interpolant over $[-1, 1]$. The integral is unweighted $dx$, so in angle coordinates there is a $\sin\theta$ Jacobian; it is not merely the trapezoidal average of the sampled values over angle.

Let $I_k = \int_{-1}^{1} T_k(x)dx = \int_0^{\pi} \cos(k\theta)\sin\theta\, d\theta$. Odd $k$ gives zero by parity. Using a product-to-sum identity for even $k$ gives $I_k = 2/(1-k^2)$, including $I_0 = 2$. With coefficients in the convention $p_N = \sum_{k=0}^{N} a_k T_k$, the quadrature is simply

$$Q_N = \sum a_k I_k.$$

Some DCT implementations store doubled endpoint coefficients. Convert to the polynomial convention before using these moments; otherwise even the constant function can integrate to four instead of two. A quick audit is exact integration of $1, x, x^2$ where the grid supports those degrees.

## Explain why unsampled higher modes can imitate low modes

At sampling angles $\theta_j = j\pi/N$, $\cos[(2N-k)\theta_j] = \cos(k\theta_j)$ for integer $j$. Frequencies folded about multiples of $2N$ are indistinguishable on this grid. A single basis function $T_{2N-k}$ consequently yields the same interpolant and quadrature as $T_k$, even though the exact integrals differ. At $N = 8$, $T_{14}$ aliases to $T_2$, so $Q_8 = -2/3$ while the exact integral is

$$-2/195 \simeq -0.01026.$$

At $N = 16$ the polynomial is fully resolved and integrated exactly. Many practical smooth integrands have rapidly decaying high-mode coefficients, limiting this effect; smoothness without a quantitative resolution check is insufficient.

The error for a convergent Chebyshev expansion can be analysed as a sum of high-mode coefficients times differences between true and aliased moments. The even moments decay roughly as $k^{-2}$, so exact integration of many high modes contributes little, but a high mode folding to $T_0$ can cause a large erroneous contribution. This helps explain how a rule exact for degree only $N$ can often perform competitively with Gaussian quadrature while still suffering carefully designed aliasing failures. Number

of exact polynomial moments and actual performance on an analytic function are different comparisons.

## Reuse data under refinement while looking for unresolved structure

Doubling the number of intervals gives $x_{2j}^{(2N)} = x_j^{(N)}$, so all previous samples can be reused. This is valuable when evaluating one function call involves a time-dependent model or laboratory observation. For a general interval $[a, b]$, map $x$ affinely and multiply every integral by $(b-a)/2$. Compare successive $Q_N$, decay of the latter Chebyshev coefficients and an independent nonnested or local check. If the integrand has an unresolved local spike or frequency intentionally locked to the grid, two nested results can agree spuriously.

For a diagnostic let $f = T_{4N}$. Both the $N$- and $2N$-interval grids sample the constant value one, so both return 2 despite the exact integral $2/(1-16N^2)$ being small and negative. Once the $4N$-interval grid resolves the highest polynomial the answer changes dramatically. This does not imply nested algorithms should always quadruple a mesh unquestioningly. When physical bounds on derivatives, a known bandwidth or local event locations are available, use them to select an initial resolution and targeted check. Without such information no deterministic finite sampling can certify arbitrary unknown integrands.

**Background · Can two consecutive nested rules agree on the wrong integral?**

Use the Chebyshev polynomial $T_{4N}(x)$ on $[-1, 1]$. On the fine grid $x_j = \cos(j\pi/(2N))$, its sampled values are $\cos(2j\pi) = 1$. The coarse grid is a subset and sees exactly the same constant. Both rules consequently return the integral of one, namely 2. But the unweighted integral follows by setting $x = \cos\theta$ and integrating $\sin\theta\cos(4N\theta)$:

$$\int_{-1}^{1} T_{4N}(x)\,dx = \frac{2}{1-(4N)^2}.$$

At $N = 4$ both sampled estimates are 2 while the actual integral is $-2/255$. An independent off-grid evaluation, or a refinement sequence chosen using prior frequency information, can expose the oscillation. This example pinpoints which assumption a difference-based stopping test needs: the current samples must already represent the significant scales.

## Make two nested rules agree on a false integral

Take eight coarse intervals and choose degree 32. Read the two interpolated moments and derive the correct integral analytically. At degree 31 both rules also agree. Explain why they are now exact. At

degree 16 determine which grid returns 2 and which resolves the true moment. What evidence distinguishes symmetry from a spurious agreement due to aliasing?

Open this experiment online

### Try the calculation

Intervals on the coarse cosine grid **8**

Chebyshev polynomial degree to inspect **32**

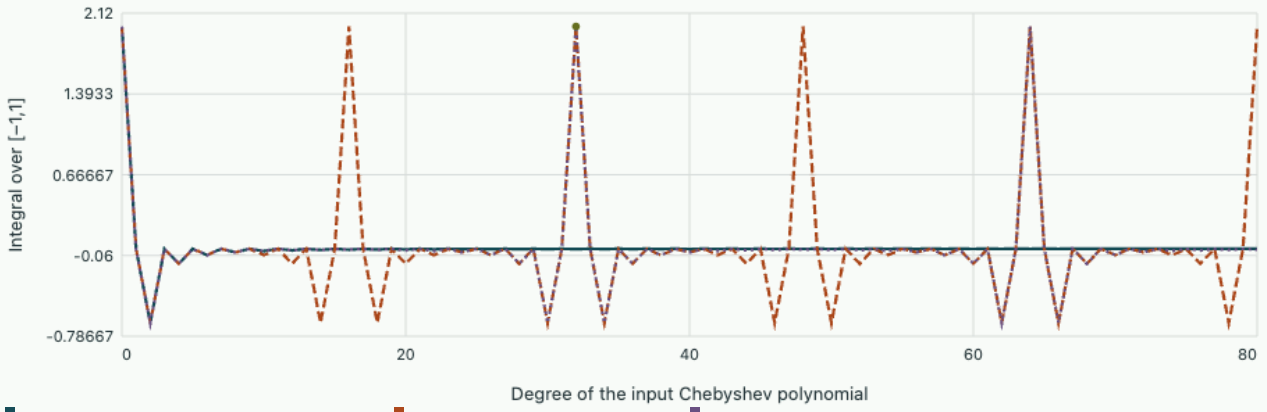


Exact polynomial moment · Coarse cosine grid · Twice as many intervals · Chosen degree on coarse grid

**Degree 32 integrates to -0.001955. The two nested estimates are 2 and 2.**

# Tutorial · inspect coefficients and stopping signals

## Build a three-point rule from moments

With $N = 2$ the nodes are $1, 0, -1$. Use symmetry and moments of $1, x^2$ to obtain the weights. Integrate $x^4$ and explain the error.

▼ **Work through the solution**

Symmetry gives endpoint weights $w$ and centre weight $v$. Exactness of the constant requires $2w + v = 2$; for $x^2$ it gives $2w = 2/3$, hence

$$w = 1/3, v = 4/3.$$

The estimate for $x^4$ is $2/3$ whereas the exact integral is $2/5$, an error of $4/15$. Although the rule integrates odd cubics by symmetry, polynomial degree four is not guaranteed. Expanded in Chebyshev modes, its unresolved $T_4$ aliases to the constant.

## Derive the error caused by one small high-frequency contaminant

At $N = 8$ a signal is $1 + 0.01T_{16}$. Compare the sampled integral with the exact one and inspect a second grid with 16 intervals.

▼ **Work through the solution**

Every coarse node samples 1.01, yielding 2.02. The exact integral is

$$2 + 0.01[2/(1 - 256)] = 1.9999216.$$

The error around 0.0201 is large relative to the true high-mode integral because its samples look constant. At 16 intervals the true mode is represented, giving the exact value up to roundoff. A small-amplitude frequency component can contribute disproportionately to a coarse low-mode estimate when it aliases.

---

## Determine whether an adaptive error diagnostic has a rigorous guarantee

Two nested rules agree to $10^{-10}$ and all computed high coefficients are small. Does that certify the integral for an otherwise unconstrained analytic function? Construct an explicit counterexample.

▼ **Work through the solution**

At any fixed finite $N$, the polynomial $T_{4N}$ above is analytic and all coarse coefficients except the constant vanish on both grids. Adding a multiple of $T_{4N} - 1$ to an arbitrary candidate leaves both sampled grids unchanged but shifts its exact integral by nearly minus twice that multiple. Thus agreement is evidence conditional on a smoothness or scale model, not an unconditional bound over every analytic function. A genuine analytic error theorem also needs quantitative information such as a bound on an analytic continuation region and its magnitude.

---

# Further reading

- Trefethen · Is Gauss Quadrature Better than Clenshaw–Curtis?
- Course investigation · Fourier coefficients, sampling and aliasing
- Investigate interpolation error in a heat integral

# Function approximation and numerical integration · 3. Integrate the oscillations analytically and sample the envelope

*A sensor model asks for many Fourier integrals of a slowly changing envelope. Direct trapezoids can return a full-scale signal where the true integral nearly cancels. Derive a Filon rule from exact oscillatory moments and find where it needs a different local model.*

## Estimate an oscillatory integral before choosing a grid

Consider $I(\omega) = \int_a^b f(x)e^{i\omega x}dx$ with smooth envelope $f$ and known constant phase slope. The complex result contains cosine and sine transforms as real and imaginary parts. Direct quadrature can work if the grid resolves both envelope and oscillation. But a uniform grid with spacing $h = 2\pi/\omega$ samples the same complex phase at every point and can report a large entirely spurious response. A calculation accurate for the slow envelope alone need not be accurate for its product with the carrier. The explicit oscillating factor contains information we can integrate analytically rather than sampling it repeatedly.

Integration by parts gives

$$I = [f(x)e^{i\omega x}/(i\omega)]_a^b - (1/(i\omega))\int_a^b f'(x)e^{i\omega x}dx.$$

For bounded derivatives the integral is typically $O(1/\omega)$ at high frequency, and endpoint cancellations or endpoint vanishing can improve the decay. A result as large as $\int |f|dx$ at a very high frequency is therefore suspicious unless the phase has a stationary point, rapid structure exists in the envelope or a special physical resonance changes the setup. Absolute error is a better validation measure near an exact transform zero than relative error, which divides by a possibly vanishing number.

## Derive a two-node Filon panel about its midpoint

On a panel $x = c + s,\ -h \le s \le h$, approximate the envelope by $A + Bs$ through its endpoints, with $A = [f(c+h) + f(c-h)]/2$ and $B = [f(c+h) - f(c-h)]/(2h)$. Retain the carrier exactly. Define $z = \omega h$. Its zeroth moment is

$$M_0 = \int_{-h}^{h} e^{i\omega s}ds = 2h\sin z/z.$$

Differentiating with respect to $\omega$ gives $\partial_\omega M_0 = iM_1$, so

$$M_1 = \int_{-h}^{h} se^{i\omega s}ds = 2ih^2\frac{\sin z - z\cos z}{z^2},$$

$$Q_{\text{panel}} = e^{i\omega c}(AM_0 + BM_1).$$

$M_0$ is real and even in frequency and $M_1$ imaginary and odd, as symmetry requires. The rule is exact for every linear envelope at every frequency, regardless of the number of oscillations in the panel. At $\omega = 0$ it returns the trapezoidal integral of the linear interpolant. Directly subtracting $\sin z - z\cos z$ for tiny $z$ loses precision. Instead use its series $z^3/3 - z^5/30 + \cdots$, and evaluate $\sin z/z$ by a stable sinc routine or a short series near zero.

For $f(x) = 1 + x,\ x \in [0, 1]$ and $\omega = 20\pi$, the constant contribution vanishes over ten cycles. Integration by parts on $x$ gives

$$I = 1/(i\omega) = -i/(20\pi).$$

One linear Filon panel gives the exact $-0.0159155i$. A composite trapezoid with ten subintervals samples $e^{i\omega x} = 1$ at every node, returning 1.5 real. Both methods evaluated the same envelope but represented the carrier quite differently. Refining blindly by factors that preserve a frequency coincidence can produce misleading error estimates on the first few grids.

## Separate envelope approximation from phase modelling

Since the carrier has magnitude one for real phase, $|I - Q| \le \int |f - p| dx$. A piecewise-linear interpolation error bound on a panel of width $H$ is $\max |f''| H^2/8$, giving a global absolute bound $(b - a) \max |f''| H^2/8$ independent of $\omega$. Oscillatory cancellation can make the actual error much smaller, but this bound already identifies a sensible grid criterion based on envelope curvature. Higher-order Filon methods integrate more polynomial moments or use derivatives at endpoints. Stable moment recurrences and consistent treatment near $\omega = 0$ matter as much as theoretical order.

For a nonlinear phase $e^{i\omega g(x)}$, changing variable to $u = g(x)$ can reduce to a linear phase on a monotone interval, with transformed envelope $f(x(u))/g'(x(u))$. At a stationary point $g' = 0$ this substitution becomes singular. In that case oscillations slow locally and the contribution can decay as $\omega^{-1/2}$ for a simple quadratic stationary point, often larger than endpoint terms. Splitting the domain and using a local quadratic-phase or stationary-phase treatment is more informative than declaring the original linear-phase rule universally efficient. Identify the physical phase law before attributing a numerical discrepancy to envelope smoothness.

## Make a fast carrier invisible to the sampling grid

The envelope is exp(x) on the unit interval. Keep six integration panels and select six carrier cycles. Every direct trapezoidal sample sees a cosine of one. Over each full cycle both a constant and a linear envelope have zero cosine integral here. Why does the curved exponential still give 0.0012082? Change to seven panels at the same frequency and compare absolute errors. Finally set the frequency to zero: both methods integrate the same piecewise-linear envelope, so their common value can still differ from the exact integral.

Open this experiment online

### Try the calculation

Uniform integration panels 6

Selected carrier (cycles in the interval) 6

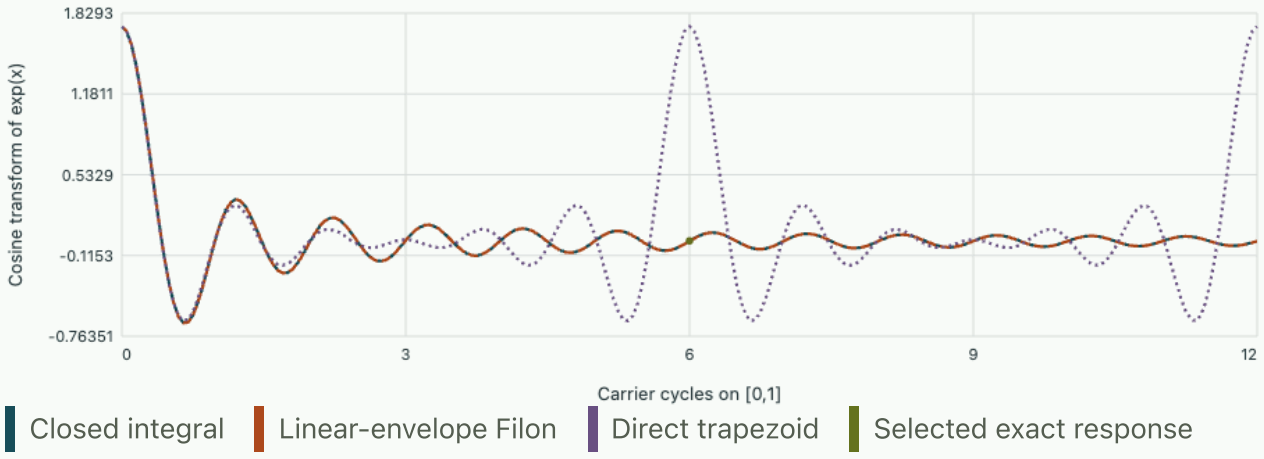


**At the selected frequency the exact response is 0.0012082. Filon gives -1.0674e-16 and trapezoids give 1.7223.**

# Tutorial · check the carrier convention and asymptotic scale

## Integrate a symmetric odd envelope

For $f(x) = x$ on $[-1, 1]$ and $\omega = \pi$, compute $I$ from the first moment. Which part vanishes?

▼ **Work through the solution**

Use $h = 1, c = 0, A = 0, B = 1$. Since $\sin \pi = 0, \cos \pi = -1$, $I = 2i/\pi$. The cosine part vanishes by parity and $x \sin(\pi x)$ is nonnegative on both sides, confirming the positive imaginary sign. Using a Fourier convention $e^{-i\omega x}$ instead would conjugate the result. A test on this odd envelope detects a sign error that any even real envelope would miss.

---

## Give a sufficient interpolation-based error budget

On a unit interval a calibrated envelope obeys $|f''| \le 6$. How many equal panels suffice for absolute error $10^{-4}$ using the elementary bound? Does doubling frequency change the stated count?

▼ **Work through the solution**

The bound is $6/(8N^2)$. Requiring it at most $10^{-4}$ gives $N \ge \sqrt{7500}$, so 87 panels suffice. The bound makes no frequency assumption beyond a known real linear phase and does not require more panels

merely because frequency doubles. Floating-point phase reduction for extremely large arguments and uncertainty in the physical frequency are separate sources of error not included in the envelope bound.

---

### Estimate sensitivity to a miscalibrated frequency

Show that a frequency error $\delta\omega$ can perturb the integral by at most $|\delta\omega| \int |x f(x)| dx$. Evaluate for the previous $f = 1 + x$ on $[0, 1]$ and $|\delta\omega| = 0.01$.

▼ **Work through the solution**

Differentiate under the integral to obtain $I'(\omega) = i \int x f(x) e^{i\omega x} dx$, then bound its norm and integrate over the frequency error. Here $\int_0^1 x(1 + x) dx = 5/6$, giving an absolute perturbation at most 0.00833. Relative to the small exact response at $20\pi$, this can be substantial. An exactly integrated carrier on the wrong frequency is therefore not a high-accuracy physical prediction merely because the quadrature test is exact.

---

## Further reading

- Iserles · numerical quadrature of highly oscillating Fourier integrals, full paper
- Course investigation · refinement and quadrature verification
- Investigate interpolation error in a heat integral

# Function approximation and numerical integration · 4. Spend random samples where they reduce uncertainty

*A Monte Carlo estimate can be unbiased and still too noisy to be useful. Derive control variates, allocate a fixed budget across strata, and distinguish variance reduction from an underestimated standard error.*

## Subtract a fluctuation with an independently known mean

To estimate $I = \mathbb{E}[f(X)]$, suppose evaluating $f(X)$ also supplies a cheap correlated quantity $g(X)$ with known exact mean $\mu_g$. For any fixed coefficient $\beta$, $Y = f(X) - \beta[g(X) - \mu_g]$ has the same mean $I$. Its variance is $V_f - 2\beta C_{fg} + \beta^2 V_g$. Differentiating gives optimum $\beta_* = C_{fg}/V_g$ if $V_g > 0$ and minimum $V_f(1 - \rho_{fg}^2)$. For a crude physical surrogate it is accuracy of its fluctuations, not its absolute scale, that matters. A negatively correlated surrogate can also reduce variance by a negative coefficient. A constant control carries no information and should not be divided by zero variance.

As an exact example take $X \sim U[0,1]$, $f = X^2$, $g = X$. We know $\mu_g = 1/2$ and seek $I = 1/3$. Direct moments give

$$V_f = 1/5 - 1/9 = 4/45,\ V_g = 1/12,\ C = 1/4 - 1/6 = 1/12.$$

Hence $\beta_* = 1$ and residual variance $4/45 - 1/12 = 1/180$, sixteen times smaller than the original. At the same sample count the standard deviation falls by four. Each controlled sample $X^2 - X + 1/2$ is not equal to the true answer, but it fluctuates much less. An independent pilot run can estimate $\beta$ for unknown covariances. Reusing the same small sample to optimise $\beta$ and estimate the mean can introduce finite-sample bias and overstate achieved variance reduction, especially when many controls are fitted.

## Partition the input space and allocate observations

Divide the domain into disjoint strata with known probabilities $w_j$, then draw $n_j$ independent observations within each. The stratified estimator is $\hat{I} = \sum_j w_j \bar{f}_j$ with variance $\sum_j w_j^2 \sigma_j^2 / n_j$, where $\sigma_j^2$ is the conditional variance. Correct weighting is crucial: the allocation probabilities need not equal the target distribution's stratum probabilities. At a fixed total sample count $N = \sum n_j$, differentiating with a Lagrange multiplier gives $n_j \propto w_j \sigma_j$, known as Neyman allocation. For unequal sample costs $c_j$ under fixed $\sum c_j n_j$, the corresponding optimum is $n_j \propto w_j \sigma_j / \sqrt{c_j}$. Pilot uncertainty and minimum counts per stratum matter before rounding to feasible integers.

For two equal-probability strata with conditional standard deviations 1 and 4, $N = 100$ calls yields optimal counts 20 and 80. The variance becomes $(1/4)(1/20 + 16/80) = 0.0625$, compared with $0.085$ from 50 each. The difference is material but smaller than one might guess from the 4:1 standard deviation ratio. Overallocating all observations to the difficult stratum would leave the other conditional mean unidentified. A deterministic stratum whose mean is known separately can be removed; one whose variance merely looked zero in a tiny pilot may still need independent sampling. For rare heavy-tail excursions, a naive pilot can miss the very fluctuations it was intended to characterise.

## Choose an antithetic pair with a justified covariance

For symmetric $U \sim U[0,1]$, both $U$ and $1-U$ have the target distribution. The pair average $Z = [f(U) + f(1-U)]/2$ is unbiased. Its variance is $[V_f + \mathrm{cov}(f(U), f(1-U))]/2$. With the same two function evaluations, two independent samples would have variance $V_f/2$. Antithetic pairing therefore helps exactly when the paired function values are negatively correlated, as with a monotone function over this scalar interval. If $f(U) = f(1-U)$ the pair is redundant; the method then doubles cost per independent value. Compute standard errors across independent pairs, not by treating both values inside each pair as independent calls.

For $f = X^2$, $Z = X^2 - X + 1/2$, algebraically the same controlled random variable encountered above. Its variance is $1/180$ per pair and it costs two evaluations. The linear control cost only one expensive $f$ evaluation when $g = X$ was nearly free. The two methods yield identical random fluctuations per independent draw of $X$ here, but different error per unit computation. Comparing only the algorithm's reported sample count without accounting for pairs, pilot runs and surrogate evaluations can reverse a practical recommendation.

## When does the control increase the noise?

Draw X uniformly on the unit interval and estimate the mean of X². The known mean of X is one half. Compare the standard errors at coefficients zero, one and three. Use the graph to find how many direct samples match the uncertainty of 32 optimally controlled samples, then confirm the answer from the exact variance ratio.

Open this experiment online

### Try the calculation

Coefficient β in X² − β(X − ½) **1**

Sample count to inspect **64**

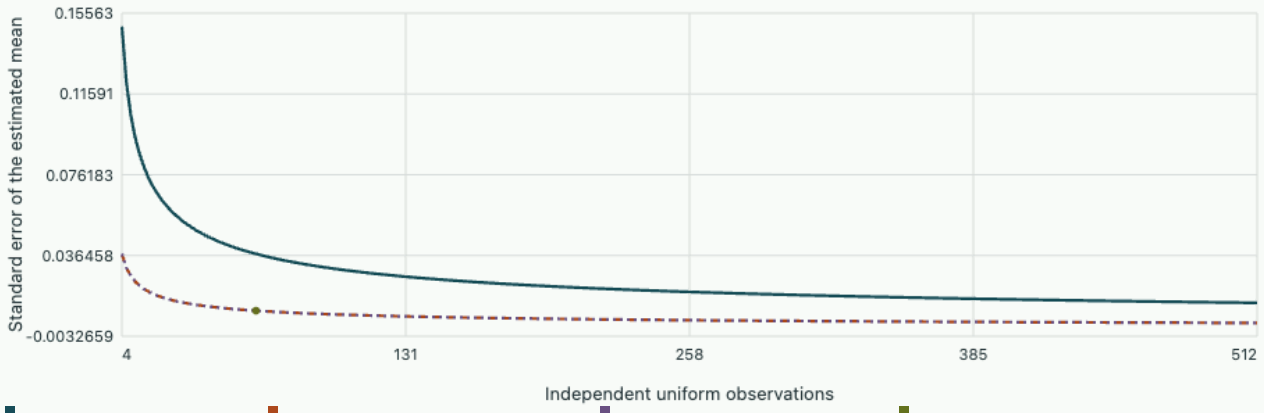


Direct estimator | With selected control | Optimal coefficient | Selected budget

**Variance ratio controlled/direct = 0.0625. With 64 independent samples the controlled standard error is 0.0093169.**

# Tutorial · verify variance and bias separately

## Budget calls with a measured correlation

A model response has standard deviation 6 and a validated fixed-coefficient control has correlation 0.95. Estimate calls for standard error 0.10 with and without control, neglecting control cost.

**▼ Work through the solution**

Plain sampling requires $6^2/0.1^2 = 3600$ independent calls. The controlled variance is $36(1 - 0.95^2) = 3.51$, requiring 351 independent calls. These counts assume the reported correlation and optimal coefficient apply to the production distribution. Validate with held-out paired observations. A mistaken known control mean introduces bias $\beta(\mu_{\text{used}} - \mu_{\text{true}})$, which averaging more calls cannot remove.

---

## Find the cost-aware allocation

Two equal-probability strata have standard deviations 1 and 4. Calls in the second cost four times those in the first. Allocate 180 cost units, measuring cheap calls in units of one.

**▼ Work through the solution**

The optimum ratio $n_2/n_1 = (4/\sqrt{4})/(1/\sqrt{1}) = 2$. With $n_1 + 4n_2 = 180$ this gives

$$n_1 = 20, n_2 = 40.$$

Their variance is

$$(1/4)(1/20 + 16/40) = 0.1125.$$

Assigning half the calls to each would instead permit 36 and 36 with variance about 0.118. Assigning half the budget to each gives a different allocation again. Specify the constrained physical resource before quoting an optimum.

---

## Identify an antithetic failure and repair the error estimate

For $f(x) = (x - 1/2)^2$, an analyst generates 100 independent antithetic pairs but computes an ordinary sample error as though all 200 values were independent. What is wrong?

▼ **Work through the solution**

Each pair consists of identical function values. There are only 100 independent observations of this response; the variance of the paired estimator is $V_f/100$, whereas the false independent count suggests $V_f/200$. Estimate uncertainty from the variance of the 100 pair averages or use 200 genuinely independent calls if possible. Here the pairing has no cancellation because the function is symmetric under precisely the transformation used to generate the second member.

---

# Further reading

- ETH Zürich · numerical methods for CSE, quadrature and statistical estimation
- Course investigation · statistical error in random quadrature
- Investigate interpolation error in a heat integral

# Function approximation and numerical integration · 5. How many neighbouring samples belong in a slope estimate?

*A wider local fit suppresses measurement noise, yet it can smooth away real changes in slope. Derive the competing errors, then use noisy synthetic data to find which differentiator wins at a chosen spacing.*

## Fit a polynomial in local coordinates

Suppose we observe $y_j = f(jh) + \epsilon_j$ at equally spaced times around zero. The errors are independent with mean zero and common standard deviation $\sigma$. Fit $a + bx + cx^2$ by minimizing the sum of squared differences on five points with $j = -2, -1, 0, 1, 2$. Differentiate that sum with respect to $b$. Because $\sum j = \sum j^3 = 0$, the odd column of the design matrix is orthogonal to both even columns. The fitted slope is therefore

$$\hat{b} = \frac{\sum j y_j}{h \sum j^2} = \frac{-2y_{-2} - y_{-1} + y_1 + 2y_2}{10h}.$$

Its weights have zero sum and give exactly one on a unit linear ramp. A quadratic term cancels by symmetry. A cubic term $\beta x^3$ contributes $\beta h^2 \sum j^4 / \sum j^2 = 3.4 \beta h^2$ to the estimate. Meanwhile the independent noise variances add after weighting. Their sum is $\sigma^2 / (10h^2)$. For a three-point fit the same calculation gives a cubic bias of $\beta h^2$ and a variance of $\sigma^2 / (2h^2)$.

Write the leading bias as $Bh^2$ and the noise standard deviation as $S/h$. The mean square error around the true derivative is $B^2 h^4 + S^2 / h^2$. Its minimum is at $h = (S^2 / (2B^2))^{1/6}$ when both contributions are present. Repeated noisy measurements could reduce the noise term too. Correlated measurement errors require the covariance matrix between samples in the variance calculation.

## Diagnose a surprising result at two spacings

Use a signal $f(x) = 1 + x + 0.4x^2 + x^3$ and noise standard deviation 0.02. The target derivative at zero is one. At $h = 0.20$ the five-point bias is 0.136 and its noise standard deviation is about 0.0316, giving a root mean square error of 0.1396. The three-point bias is 0.040, its standard deviation is 0.0707 and its total error is 0.0812. The shorter fit wins despite its greater sensitivity to random noise.

At $h = 0.08$ the five-point values change to 0.02176 and 0.07906, giving a total of 0.0820. The three-point result is about 0.1769. Changing the spacing has exchanged the dominant source of error. A single realization can put its estimate on either side of the derivative. Change the random seed while keeping the spacing fixed to distinguish this variability from the average error in the second plot.

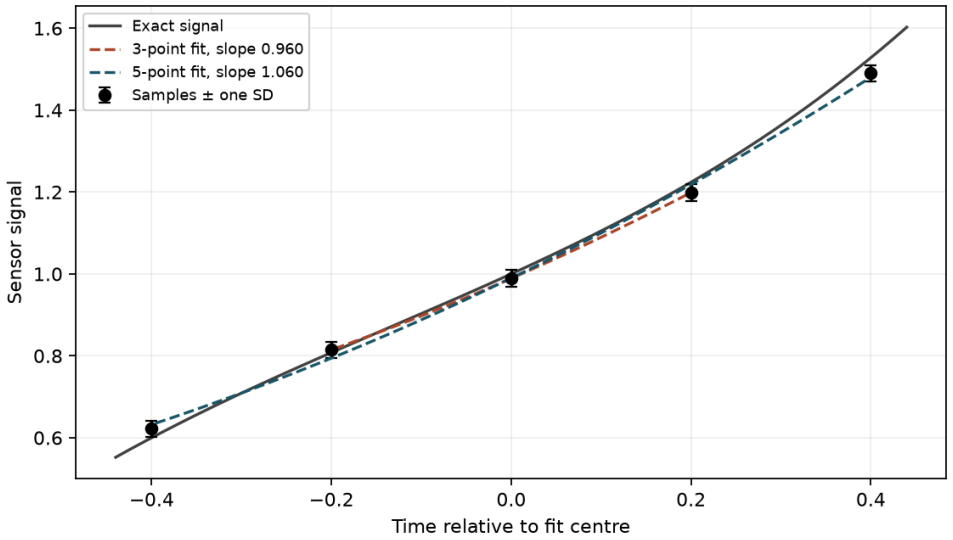


The three-point and five-point local fits use nested samples. Inspect the tangent at the centre, where the derivative of the underlying cubic signal equals one.

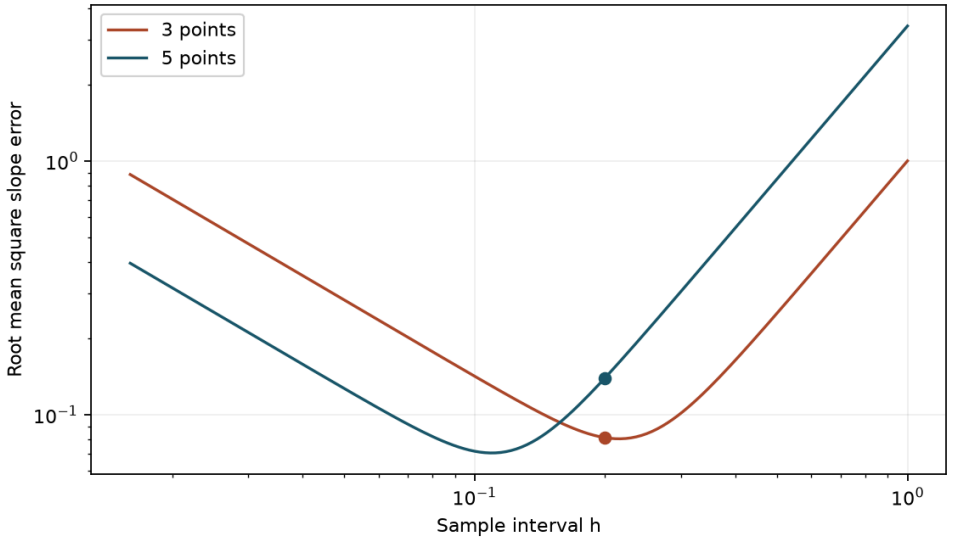


For a cubic signal the bias-plus-variance formula is exact. The location of the minimum is an experimental design choice controlled by signal curvature change and noise.

## Exercise · predict a single perturbation

A five-point experiment at spacing 0.2 measures the rightmost sample 0.03 above its true value. By how much does that change the slope? What happens if the same offset appears on all five samples?

▼ **Work through the solution**

The rightmost derivative weight is $2/(10h) = 1$ in these units. The slope rises by 0.03. A common offset contributes its magnitude times the sum of all weights, which is zero. This cancellation also explains why an uncertain but constant instrument zero does not affect this derivative.

# Exercise · choose a spacing when noise increases

If an otherwise identical experiment has eight times the noise standard deviation, how should the optimum spacing change? Will this scaling hold for a perfectly quadratic signal?

▼ **Work through the solution**

With fixed cubic coefficient $B$, the optimum is proportional to $\sigma^{1/3}$. Eight times the standard deviation therefore doubles the best spacing. For a quadratic the cubic bias vanishes identically. The stated model then predicts decreasing noise error as the spacing grows. In a real record the usable window can be limited by unmodelled higher derivatives, nearby boundaries or changing noise statistics.

---

### Change parameters · Explore this notebook model

Change one input at a time and compare the recalculated curves with the argument above. The browser evaluates this small model locally; the accompanying numerical values state its assumptions and reference checks. Reset restores the reference inputs, and the data can be downloaded as CSV.

Open this calculation with adjustable inputs

#### Try the calculation

| | |
|---|---|
| Noise standard deviation | **0.02** |
| Cubic coefficient | **1** |
| Selected sample spacing | **0.2** |

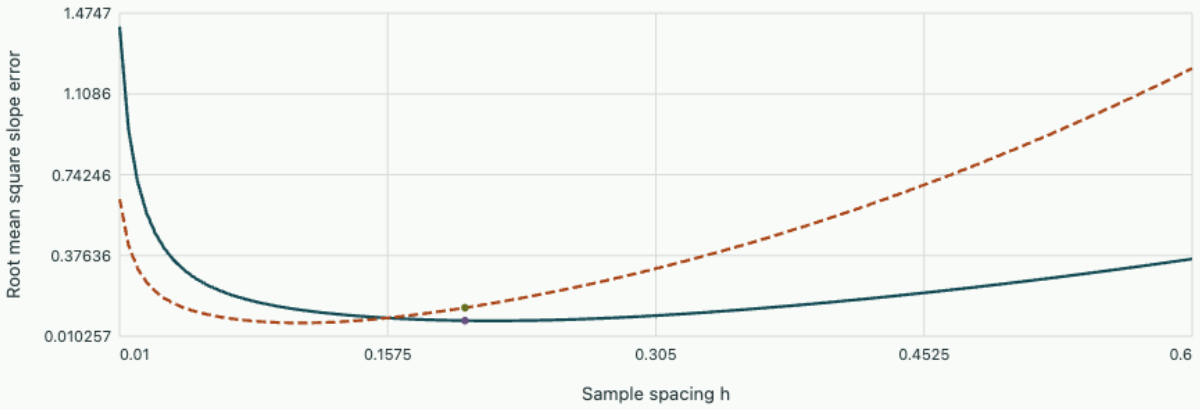


**At h = 0.2, slope RMSE is 0.08124 for three samples and 0.13963 for five. The wider fit reduces noise but admits more cubic bias.**

Open editable notebook  Download .ipynb

# Function approximation and numerical integration · 6. What can smoothness buy at a fixed polynomial degree?

*Prove density for every continuous signal, derive a faster rate for smooth signals and identify the persistent overshoot when a jump violates the assumptions.*

## Construct a convergent polynomial without differentiating the data

For continuous $f$ on $[0,1]$, let $K$ count successes in $n$ independent Bernoulli trials of success probability $x$. The expected sample $B_n f(x) = E[f(K/n)]$ is a polynomial of degree at most $n$. The weights sum to one and $E[K/n] = x,\ \mathrm{Var}(K/n) = x(1-x)/n$. Let $\omega_f(\delta)$ be the largest change in $f$ between points at distance at most $\delta$.

$$|B_n f(x) - f(x)| \le \omega_f(\delta) + 2\|f\|_\infty P(|K/n - x| > \delta) \le \omega_f(\delta) + \frac{\|f\|_\infty}{2n\delta^2}.$$

Given a target error, choose $\delta$ using uniform continuity first, then $n$ to control the second term. The bound holds simultaneously for every $x$. For an $L$-Lipschitz signal, Cauchy–Schwarz improves the argument directly to $\|B_n f - f\|_\infty \le L/(2\sqrt{n})$. For $f(x) = x^2$, taking the second moment gives the exact excess $x(1-x)/n$. Existence of a uniformly good polynomial gives no guarantee that interpolating arbitrary samples will find it.

## Relate interpolation error to best achievable error

For distinct nodes let $I_n$ interpolate and $\ell_j$ be the cardinal polynomials. Its operator norm in the maximum norm is $\Lambda_n = \max_x \sum_j |\ell_j(x)|$. At a maximising point, data with the signs of the cardinal functions attain the bound and can be extended continuously between nodes. For any polynomial $p$ of degree at most $n$, $I_n p = p$, so subtraction gives

$$\|f - I_n f\|_\infty \le (1 + \Lambda_n) \inf_{\deg p \le n} \|f - p\|_\infty.$$

For Chebyshev extrema the angular representation $x = \cos\theta$ makes the cardinal functions shifted trigonometric kernels. Their absolute contributions away from the nearest angular node are bounded by a constant times $1/j$, using $\sin u \ge 2u/\pi$ on $[0, \pi/2]$. The nearest term is bounded. Summing yields $O(\log n)$ growth. At equally spaced nodes the barycentric weights are proportional to $(-1)^j \binom{n}{j}$, giving very large alternating interior weights near an endpoint as degree increases. To obtain an explicit lower bound, measure distance in grid spacings and evaluate at $1/2$. Since $\prod_{i=0}^{n} |1/2 - i|/n! = \binom{2n}{n}/(2\,4^n)$, the absolute value of the $j$th cardinal polynomial there is this factor times $\binom{n}{j}/|j - 1/2|$. For $j = \lfloor n/2 \rfloor$, the central-binomial estimate gives growth at least proportional to $2^n/n^2$. A single cardinal function is already enough to prove exponential growth of the norm. This explains why the same best error can coexist with very different interpolation errors.

# Compare an analytic continuation, a corner and a jump

Put $g(\theta) = f(\cos\theta)$. Chebyshev coefficients are cosine coefficients of an even periodic function. If the periodic $g$ has $m$ absolutely continuous derivatives and the next derivative is integrable, integration by parts $m+1$ times bounds its coefficients by $Ck^{-(m+1)}$. For $m \geq 1$, summing the tail gives $O(n^{-m})$. Piecewise smooth functions can be treated by integrating each smooth interval and retaining internal jump terms. For $|\cos\theta|$, continuity removes the first boundary term while the derivative jumps, giving coefficients of order $k^{-2}$ and a tail of order $1/n$. Analytic extension to an ellipse instead gives the geometric tail developed in the course's interpolation chapter.

At a jump of height $J$, continuous approximants cannot have a uniform error below $|J|/2$ on a domain containing both sides. For a periodic step of height 2, its sine series involves $(4/\pi)\sin((2j+1)x)/(2j+1)$. At $x = z/(2N)$ the first $N$ terms form a midpoint Riemann sum tending to $(2/\pi)\int_0^z \sin(t)/t\,dt$. Its first maximum, at $z = \pi$, is about 1.179. The overshoot persists while its width shrinks. A mean-square error can still decrease.

**How should a convergence plot treat the jump itself?**

Report the maximum over the whole interval, a second maximum excluding a fixed physical neighbourhood of the jump and an integrated squared error. Sampling at a fixed small set of points can miss the narrowing peak entirely. Increase the sampling resolution with the polynomial degree.

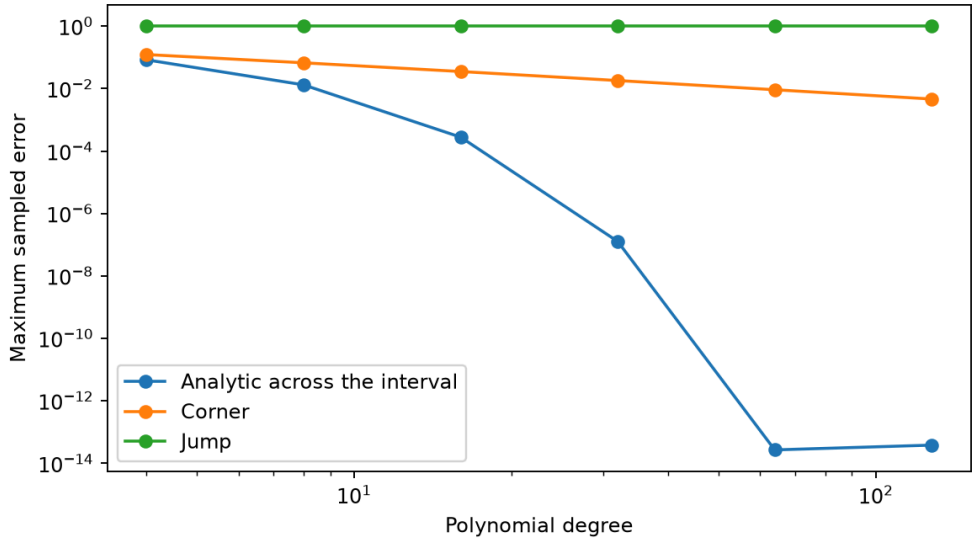


The same Chebyshev sampling rule resolves an analytic profile quickly, converges more slowly at a corner and retains a large maximum error near a jump. The last analytic points reveal finite-precision effects.

# Tutorial on changing the order of limits

## Obtain an explicit degree requirement for a bounded slope

▼ **Work through the solution**

For a target maximum error $0.01$ and $L = 2$, the Lipschitz estimate guarantees $n \geq 10^4$. This is a sufficient degree. For the specific quadratic the exact worst error is $1/(4n)$, requiring only $n \geq 25$.

---

## Predict the effect of measurement error on both meshes

▼ **Work through the solution**

Bounded nodal errors of size $\epsilon$ perturb the interpolant by at most $\Lambda_n \epsilon$. Even if the noiseless tail decreases geometrically, at fixed instrument accuracy the noise term eventually dominates. On an equally spaced mesh increasing the degree can make the prediction worse.

---

Trefethen · approximation, regularity and interpolation

### Change parameters · What can smoothness buy at a fixed polynomial degree?

Change one input at a time and compare the recalculated curves with the argument above. The browser evaluates this small model locally; the accompanying numerical values state its assumptions and reference checks. Reset restores the reference inputs, and the data can be downloaded as CSV.

Open this calculation with adjustable inputs

## Try the calculation

Largest even Chebyshev degree **20**

Corner exponent p **1**

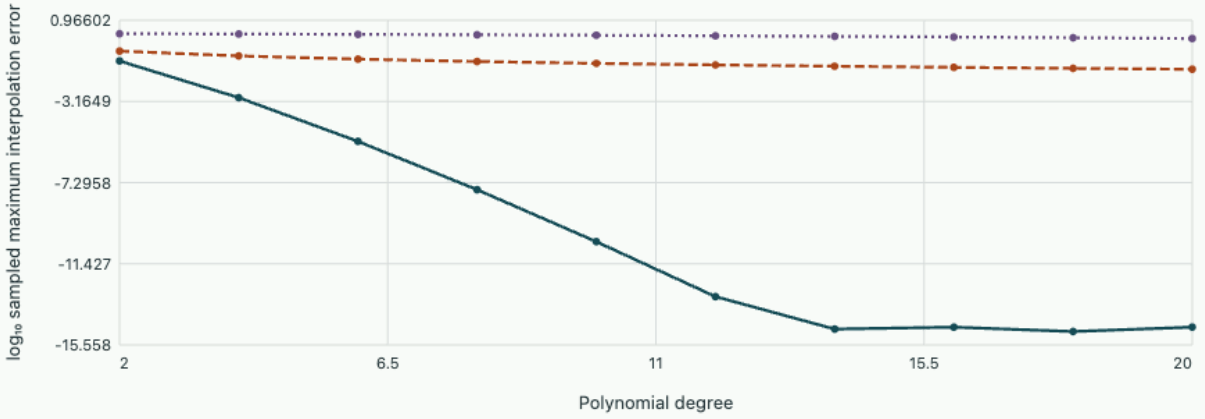


**Errors are sampled on 402 fixed test locations, not certified supremum bounds. Smoothness changes the attainable rate; the jump remains incompatible with uniform convergence by continuous polynomials.**

# Function approximation and numerical integration · 7. When should the denominator be fitted too?

*Develop a minimax certificate for rational approximants, find the pole of a measured response and identify the obstruction that appears when a trial denominator vanishes inside the interval.*

## Count independently adjustable coefficients

Write $r(x) = p_m(x)/q_n(x)$ on a compact real interval. Scaling numerator and denominator together leaves the response unchanged. Require a pole-free denominator of fixed sign and a normalisation, such as $q(0) = 1$ when admissible. For $r, s$ of the same type,

$$r - s = \frac{p_m \tilde{q}_n - \tilde{p}_m q_n}{q_n \tilde{q}_n}.$$

The numerator has degree at most $m + n$. If the error of a candidate takes alternating signs with magnitude $E$ at $m + n + 2$ ordered points and never exceeds $E$, any strictly better pole-free approximant would force $m + n + 1$ sign changes of the numerator. That is impossible unless the difference vanishes identically. This proves sufficiency. Common factors or unused degrees cause defects and can reduce the necessary alternation count. The certificate requires the global maximum between tested points to be checked.

## Recompute the extreme points after changing the pole

For $m = n = 1$, $q = 1 + cx$, impose $p(x_i) = [f(x_i) - (-1)^i E](1 + cx_i)$ at four proposed points. Products $Ec$ make this nonlinear. One Remez-type iteration solves those four equations for the coefficients and $E$, searches between the old points for extrema of $f - r$, then exchanges points until magnitudes agree and the denominator retains its sign. Locating roots of $(f - r)'$, plus testing endpoints, is essential. A dense sample alone gives only a lower bound on the supremum.

For an explicit example approximate $f(x) = 1/(2 - x)$ on $[-1, 1]$. A linear best approximant must have the endpoint-error line as a starting shape. The endpoint secant is $s(x) = 2/3 + x/3$. Convexity makes $s - f \geq 0$, maximised where $1/(2 - x)^2 = 1/3$, hence $x = 2 - \sqrt{3}$. Set $d = s(x) - f(x) = (4 - 2\sqrt{3})/3$ there. Moving the line down by $d/2$ produces three equal alternating errors and proves optimality for degree one. The type $(0, 1)$ rational function $p = 1/2,\ q = 1 - x/2$ gives zero error. A comparison at equal parameter count has a physical interpretation when a material response genuinely has a nearby pole.

**Why does matching a series at the origin give a different optimisation problem?**

Padé chooses derivatives at one point. Uniform best approximation chooses the largest error over an interval. For a response with a small additional smooth background the Padé coefficients still reproduce the local series, but its largest positive and negative errors generally have unequal magnitudes. Use held-out values near the interval boundaries.

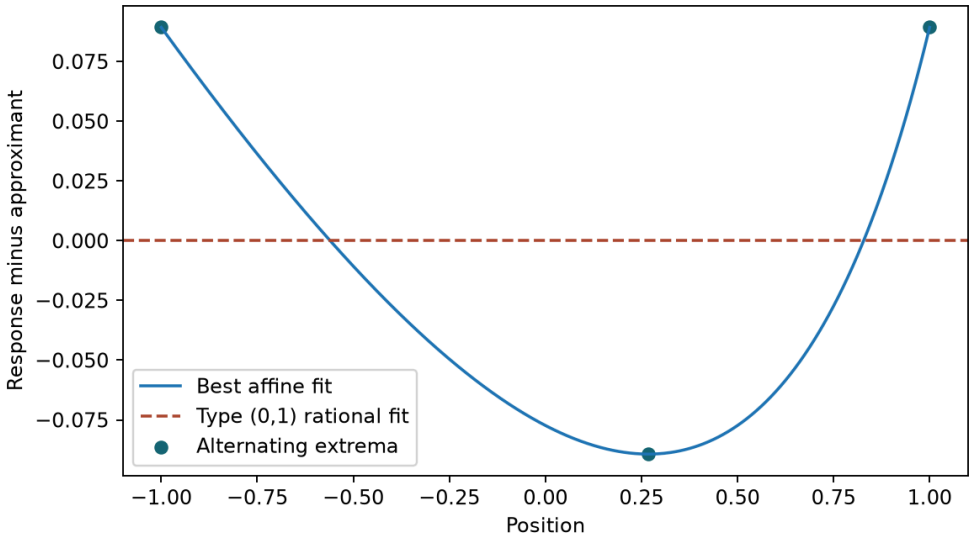


The three signed extrema certify the best affine fit. Allowing one denominator parameter puts a pole at the physical location and reproduces this response exactly.

# Tutorial on poles and certificates

## Give an interval-wide error bound when the fitted pole is uncertain

▼ **Work through the solution**

For $f_a = 1/(a - x)$ with $a > 1 + \epsilon$ and $|\Delta a| \le \epsilon$, subtraction gives $|f_{a+\Delta a} - f_a| \le \epsilon/[(a-1)(a-1-\epsilon)]$. A tightly fitted pole near the endpoint can still create a large uncertainty in the response.

---

## Construct a counterexample to using a few residuals as the certificate

▼ **Work through the solution**

Samples at $-1, 0, 1$ all avoid $x = 0.3$. The function $\eta/(x - 0.3)$ has arbitrarily small values at those samples when $\eta$ is small, yet unbounded interior magnitude. Verify the denominator is nonzero over the entire interval before interpreting alternation at measured points.

---

Trefethen · rational approximation and minimax theory

### Change parameters · When should the denominator be fitted too?

Change one input at a time and compare the recalculated curves with the argument above. The browser evaluates this small model locally; the accompanying numerical values state its assumptions and reference checks. Reset restores the reference inputs, and the data can be downloaded as CSV.

Open this calculation with adjustable inputs

## Try the calculation

Pole location a outside interval **2**

Fitted pole error **0**

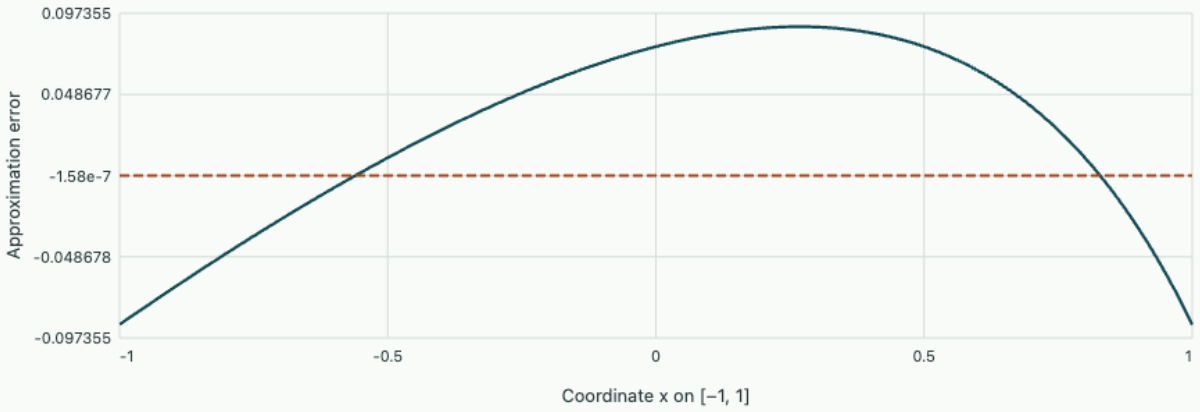


Best affine approximation  Rational with chosen pole

**Best affine error magnitude 0.089316 with interior extremum x = 0.26795. A correctly placed single pole reproduces this target exactly; pole error introduces a structured bias.**

# Function approximation and numerical integration · 8. Measuring an area integral with finitely many samples

*Build deterministic multidimensional rules from moments and geometry. Compare tensor products, symmetric triangle rules and sparse combinations at a fixed number of evaluations.*

## Which polynomials does a product rule integrate?

On a rectangle Fubini's theorem factors integrals of $x^i y^j$. Applying a one-dimensional rule in each direction forms weights $w_a w_b$ at pairs $(x_a, y_b)$. If each factor integrates through degree $2n-1$, the tensor rule integrates every monomial with each exponent at most $2n-1$, including some of total degree above that number. For $d$ axes the price is $n^d$ values. Unequal smoothness can justify different orders on different axes.

On the reference triangle $x, y \geq 0,\ x + y \leq 1$, integration of the inner monomial followed by the beta integral gives

$$\int_T x^i y^j \, dA = \frac{i!\,j!}{(i+j+2)!}.$$

Seek a symmetric three-node rule at permutations of barycentric coordinates $(a, b, b)$. Symmetry integrates affine functions once the weights each equal $1/6$. Matching $x^2$ gives $(a^2 + 2b^2)/6 = 1/12$ with $a + 2b = 1$. One solution uses $a = 2/3, b = 1/6$. The quadratic cross moment follows from barycentric symmetry and $(x + y + z)^2 = 1$. For a physical triangle multiply the reference weights by the absolute Jacobian determinant. For a curved mapping the determinant may itself vary across nodes.

## Move the singularity into the Jacobian

Suppose a surface flux density on the same triangle is $1/\sqrt{x+y}$. With $s = x + y,\ v = y/(x+y)$, the map $x = s(1-v), y = sv$ has determinant $s$. Hence the total is $\int_0^1 \int_0^1 s^{1/2}\, dv\, ds = 2/3$. Substituting $s = u^2$ gives integrand $2u^2$. Two Gauss nodes along $u$ and one along $v$ now suffice in exact arithmetic. A Cartesian mapping $x = u, y = (1-u)v$ leaves an endpoint singularity even after multiplication by its Jacobian.

**OneSimulation experiment · Integrate a singular field over a triangle**

Vary the tensor order. Predict which transformed integrand is exactly integrated first.

Download the executable notebook · Python source

**Adjust the experiment in the browser.** Generalise the flux to $(x+y)^{-p}$, with $0 \leq p < 2$. Writing $s = x + y$ gives $\int_0^1 s^{1-p}\, ds = 1/(2-p)$. The squared radial map has transformed integrand $2u^{3-2p}$. Change the exponent and quadrature order, then predict when a polynomial integrand permits exact Gauss integration. In particular, increasing $p$ past $3/2$ leaves an integrable

singularity even in the radial coordinate. The browser constructs the Gauss nodes by Legendre root iteration; the notebook obtains the same rules from OneSimulation for the $p = 1/2$ reference case.

Open the live experiment

## Try the calculation

Corner exponent p **0.5**

Maximum Gauss order **6**

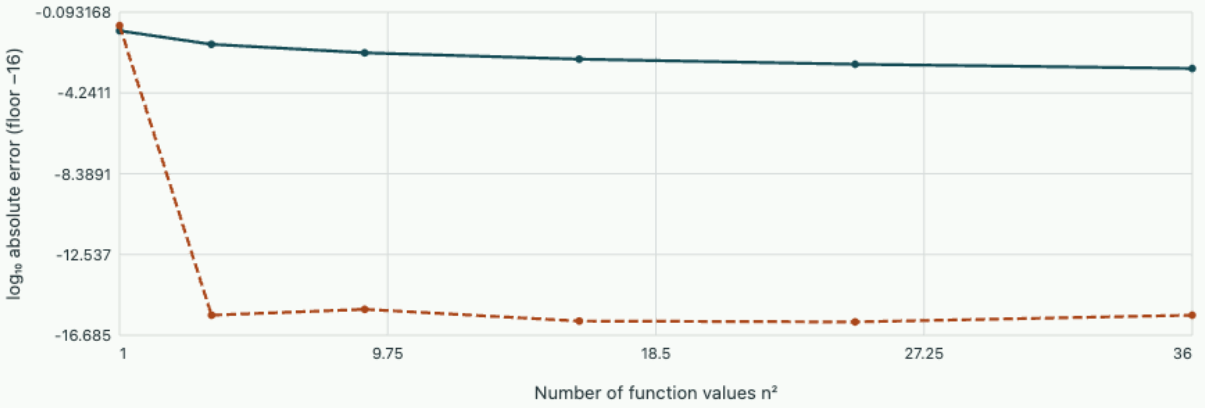


Cartesian triangle map  Squared radial map

**Exact integral 1/(2 − p) = 0.66667. With 36 evaluations, errors are 0.0010391 (Cartesian) and 2.2204e-16 (radial). For p = 1/2 the radial integrand is 2u², so order 2 already integrates it to rounding accuracy.**

```
"""Integrate a corner singularity over the unit right triangle by two maps."""
import matplotlib.pyplot as plt
import numpy as np
from onesim.numerics.discretization.geometry.carriers.kernels.quadrature.rules.reference_quadrature
import gauss_legendre

orders = [1, 2, 3, 4]
errors = [[], []]
for order in orders:
    points, weights = np.asarray(gauss_legendre(order)).T
    nodes, weights = (points+1)/2, weights/2
    u, v = np.meshgrid(nodes, nodes, indexing="ij")
    # First map x=u, y=(1-u)v. Second map x=u²(1-v), y=u²v.
    values = [(1-u) / np.sqrt(u+(1-u)*v), 2*u**2 * np.ones_like(v)]
    for sequence, integrand in zip(errors, values):
        sequence.append(abs(weights @ integrand @ weights - 2/3))
print(f"Order {orders[-1]} errors: Cartesian map {errors[0][-1]:.6g}, radial map {errors[1]
[-1]:.3g}")
assert max(errors[1][1:]) < 1e-13
fig, ax = plt.subplots(figsize=(7, 4), layout="constrained")
for err, label in zip(errors, ["Cartesian map", "Radial map"]):
    ax.semilogy(np.square(orders), np.maximum(err, 1e-16), "o-", label=label)
ax.set(xlabel="Number of function values", ylabel="Absolute error (plot floor 1e-16)")
```

```
ax.legend()
fig.savefig("numerics-triangle-cubature.png", dpi=170)
```

Output from the reference run

```
Order 4 errors: Cartesian map 0.00311717, radial map 1.11e-16
```

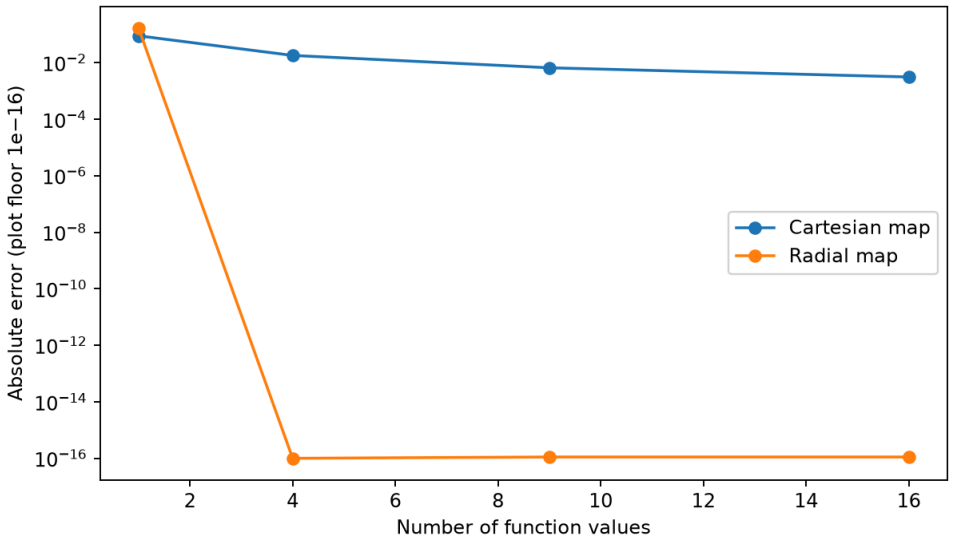


Two transformations of the same triangle lead to different quadrature errors.

# Buy mixed resolution only where it matters

Let $Q_l$ be nested one-dimensional rules of growing order, $Q_0 = 0, \Delta_l = Q_l - Q_{l-1}$. Telescoping the full tensor rule gives a sum of products of increments. In two dimensions retain terms with $l_1 + l_2 \le q$, $l_i \ge 1$. This construction spends few evaluations at points fine in both directions. If mixed derivatives control the product increments, many omitted terms are small. A sharp oblique ridge can violate that premise. Refining only along individual axes then misses the ridge.

For $q = 3$, the combination is $Q_2 \otimes Q_1 + Q_1 \otimes Q_2 - Q_1 \otimes Q_1$. Take midpoint and Simpson as nested levels on $[0, 1]$. Five distinct points replace the nine in the full tensor Simpson grid. A test on $x^2 y^2$ gives $1/12 + 1/12 - 1/16 = 5/48$, compared with exact $1/9$. The difference is exactly the mixed increment omitted from this level.

**How can two deterministic estimates agree for the wrong reason?**

A compact nonnegative bump placed between all sampled nodes returns zero on both grids. To certify an error, use an available derivative or interval bound on each region, or change the sampling geometry independently. An observed difference is an estimator whose reliability depends on resolving the integrand.

# Tutorial on domains and evaluation budgets

## Transport the triangle rule to vertices (0,0), (4,0), (1,2)

▼ **Work through the solution**

The matrix mapping coordinate differences has columns $(4, 0), (1, 2)$ and determinant 8. The three weights are $8/6$. Their sum is the area 4. Compute each physical node as the same barycentric combination of the physical vertices.

---

## Decide which quadrature to buy for three dimensions

▼ **Work through the solution**

With five levels of points per coordinate, the tensor rule costs $5^3 = 125$. If the integrand is a sum $g(x) + h(y) + k(z)$, product increments involving refinement in two directions vanish because differences annihilate constants. A sparse rule can recover the separate one-dimensional accuracy. A dependence on $x + y + z$ generally reintroduces mixed terms.

---

Davis and Rabinowitz · Methods of Numerical Integration