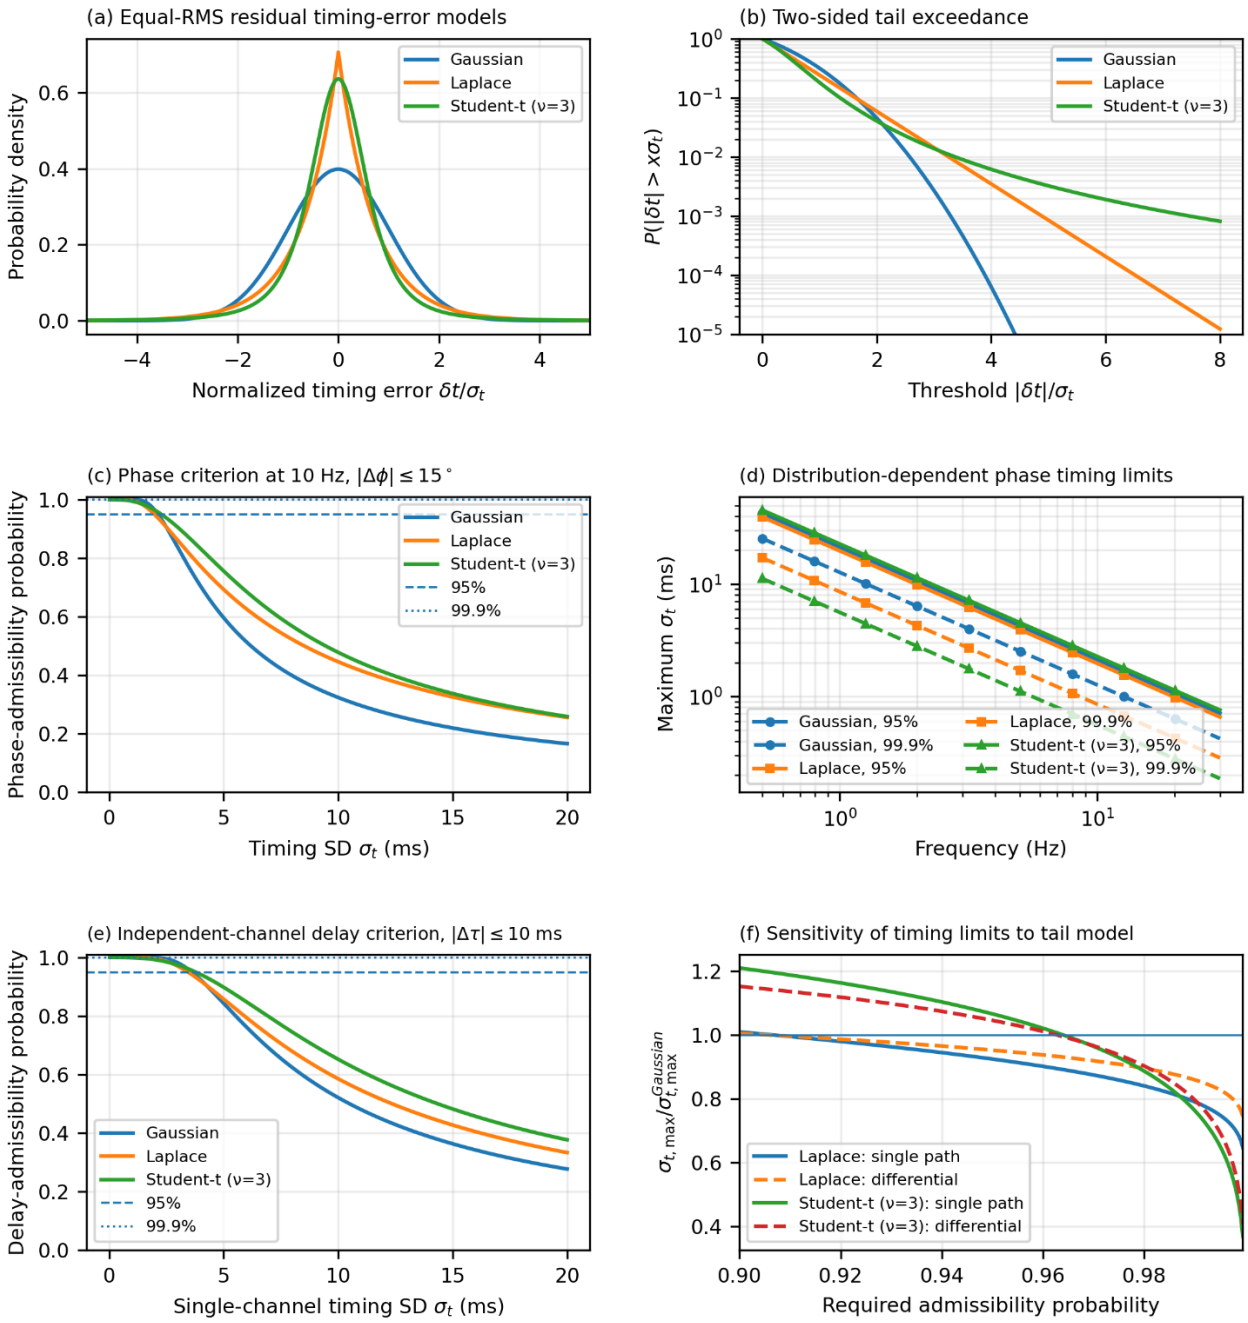


Figure S1. Distributional robustness of task-dependent timing admissibility for controlled marginal residual timing-error models. Gaussian, Laplace, and Student-t distributions with $\nu = 3$ are centred at zero and scaled to the same root-mean-square (RMS) timing error σ_t, so differences reflect distributional shape rather than variance. The Student-t_3 model is used as a deliberately heavy-tailed stress-test distribution rather than as a universal empirical law for clock errors. (a) Unit-RMS probability-density functions. (b) Corresponding two-sided exceedance probabilities, showing the progressively heavier far tails of the Laplace and Student-t models. (c) Probability that the absolute spectral-phase error remains within 15°at 10 Hz as a function of σ_t; horizontal lines mark illustrative 95 % and 99.9 % admissibility levels. (d) Maximum σ_t compatible with the same phase tolerance from 0.5 to 30 Hz for 95 % and 99.9 % target probabilities. All curves retain the inverse-frequency scaling, whereas their amplitudes depend on the relevant distribution quantile. (e) Probability that the absolute differential timing error between two independent, identically distributed channels remains within 10 ms as a function of the single-channel σ_t. (f) Ratio of the admissible $\sigma_t$under the Laplace and Student-t models to the Gaussian value as the required probability increases. Solid curves refer to a single-path timing criterion and dashed curves to independent-channel differential timing. The change in ordering with confidence level illustrates that equal-RMS heavy-tailed models can be more concentrated near zero yet substantially more restrictive in the extreme tails. Independence in panels (e) and (f) is imposed solely to isolate distributional-

shape sensitivity; correlation effects are treated separately in main-text Fig. 3. The Student-t differential distribution is evaluated by numerical convolution on $\pm 200\sigma$ support and renormalised before inversion. The distributions are controlled robustness models and are not asserted to represent a universal empirical clock-error law.

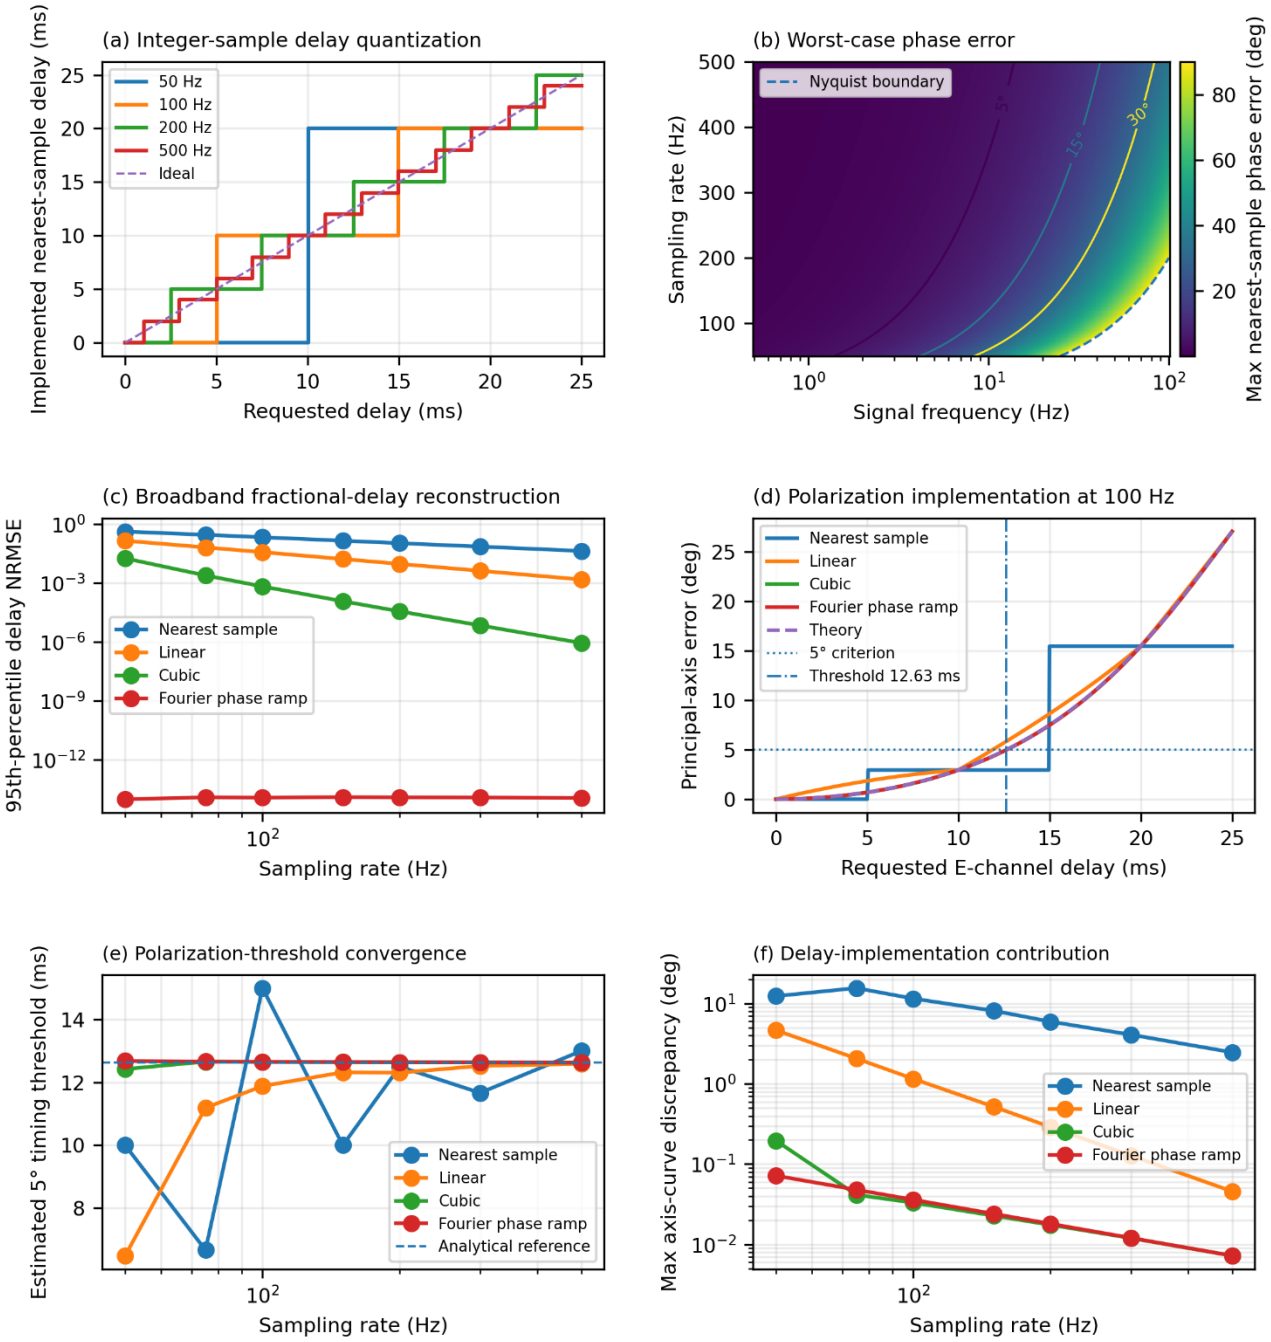


Figure S2. Sensitivity of controlled timing perturbations to sampling rate and fractional-delay implementation. (a) Nearest-sample representation of requested delays from 0 to 25 ms at sampling rates of 50, 100, 200, and 500 Hz; the dashed diagonal denotes an ideal continuously adjustable delay. (b) Analytical worst-case spectral-phase error introduced solely by nearest-sample rounding, $|\epsilon_{\phi,\max}| = 180^\circ f/f_s$. Combinations above the Nyquist frequency are masked, the dashed curve marks $f = f_s/2$, and contours denote 5°, 15°, and 30°. (c) Ninety-fifth-percentile normalised root-mean-square reconstruction error (NRMSE) over the 0–25 ms delay sweep for a deterministic periodic 4–15 Hz multisine benchmark, comparing nearest-sample, linear, cubic, and Fourier phase-ramp fractional-delay implementations. (d) Principal-polarisation-axis error at 100 Hz sampling for the ideal 10 Hz three-component benchmark evaluated in main-text Fig. 4 (azimuth 40°, incidence 45°) when only the E component is delayed. The four numerical implementations are compared with the exact long-window analytical covariance solution; the horizontal and vertical lines mark the illustrative 5°criterion and the corresponding analytical first-crossing threshold, respectively. (e) Convergence of the estimated 5°polarisation timing threshold with sampling rate. For nearest-sample shifting, the threshold is taken from the exact quantisation-cell boundary rather than by interpolation across a staircase discontinuity. (f) Maximum discrepancy between each implemented principal-axis-error curve and the analytical solution over the full delay sweep. This quantity contains both finite-sampling effects in the covariance estimator and delay-

implementation error; therefore, a small residual can remain for the Fourier method even though the Fourier phase-ramp operator itself is exact, up to floating-point precision, for the periodic band-limited controlled signals used here. Integer-sample rounding can dominate at low sampling rates, whereas linear and cubic interpolation converge as temporal sampling becomes finer. The benchmarks isolate numerical timing-implementation effects and do not include analogue anti-alias filtering, Analogue-to-Digital Converter (ADC) aperture effects, response mismatch, or asynchronous hardware clocks.

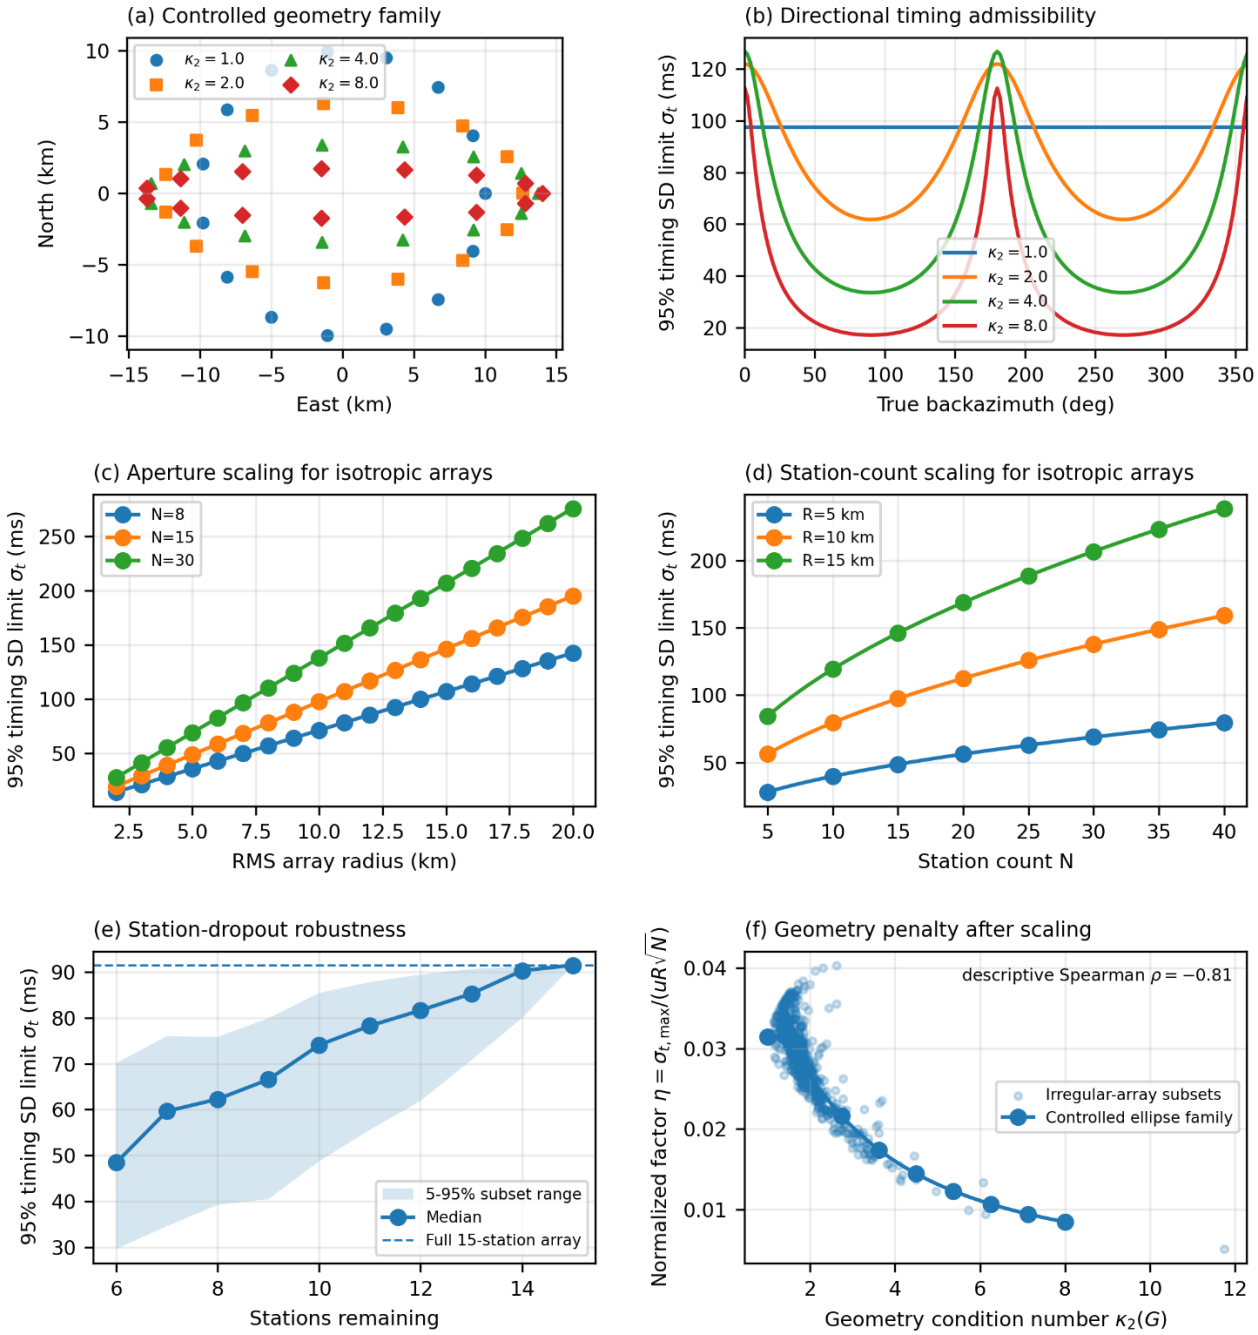


Figure S3. Sensitivity of array-level timing admissibility to geometry, aperture, station count, and station dropout. All panels use the Gaussian angular-sector probability formulation of main-text Fig. 5 with true horizontal slowness $u = 0.08$ s km^{-1}, an illustrative 5° backazimuth tolerance, and a 95 % target probability; probabilities are evaluated using 80-node Gauss–Hermite quadrature. (a) Controlled 15-station ellipse family with fixed 10 km root-mean-square (RMS) radius and increasing 2-norm geometry condition number $\kappa_2(G)$. (b) Maximum independent station timing standard deviation compatible with the backazimuth criterion as a function of true backazimuth for the geometries in panel (a). The isotropic array is direction independent, whereas increasing elongation produces progressively stronger directional sensitivity to timing errors. (c) Timing-admissibility limit versus RMS array radius for isotropic arrays with 8, 15, and 30 stations. (d) Corresponding dependence on station count for isotropic arrays with RMS radii of 5, 10, and 15 km. The controlled sweeps in panels (c) and (d) recover the expected scalings $\sigma_{t,\max} \propto R$ and $\sigma_{t,\max} \propto \sqrt{N}$, respectively. (e) Robustness to station dropout from the 15-station irregular geometry used in main-text Fig. 5. The curve gives the median timing limit for subsets retaining 6–15 stations, and the shaded region spans the 5th–95th percentiles across the subset ensemble; up to 100 unique subsets are used for each station count, all 15 possible subsets are included for $N = 14$, and the dashed line denotes the complete 15-station array. (f) Dimensionless geometry factor $\eta = \sigma_{t,\max}/\left(uR\sqrt{N}\right)$ after removing the leading slowness, aperture, and

station-count scalings. Points show the station-dropout ensemble and the curve shows the controlled ellipse family at the reference backazimuth. The descriptive Spearman association between $\kappa_2(G)$ and η is $\rho = -0.81$. The residual dependence demonstrates that $\kappa_2(G)$ is a useful, but not sufficient, scalar descriptor of the geometry penalty, because the orientation of the principal array axes relative to the arriving wavefront and station-level leverage also contribute. Arrival-phase identification and association are assumed fixed and correct throughout; the limits therefore quantify post-association parameter-estimation admissibility rather than detection or phase-association failure.

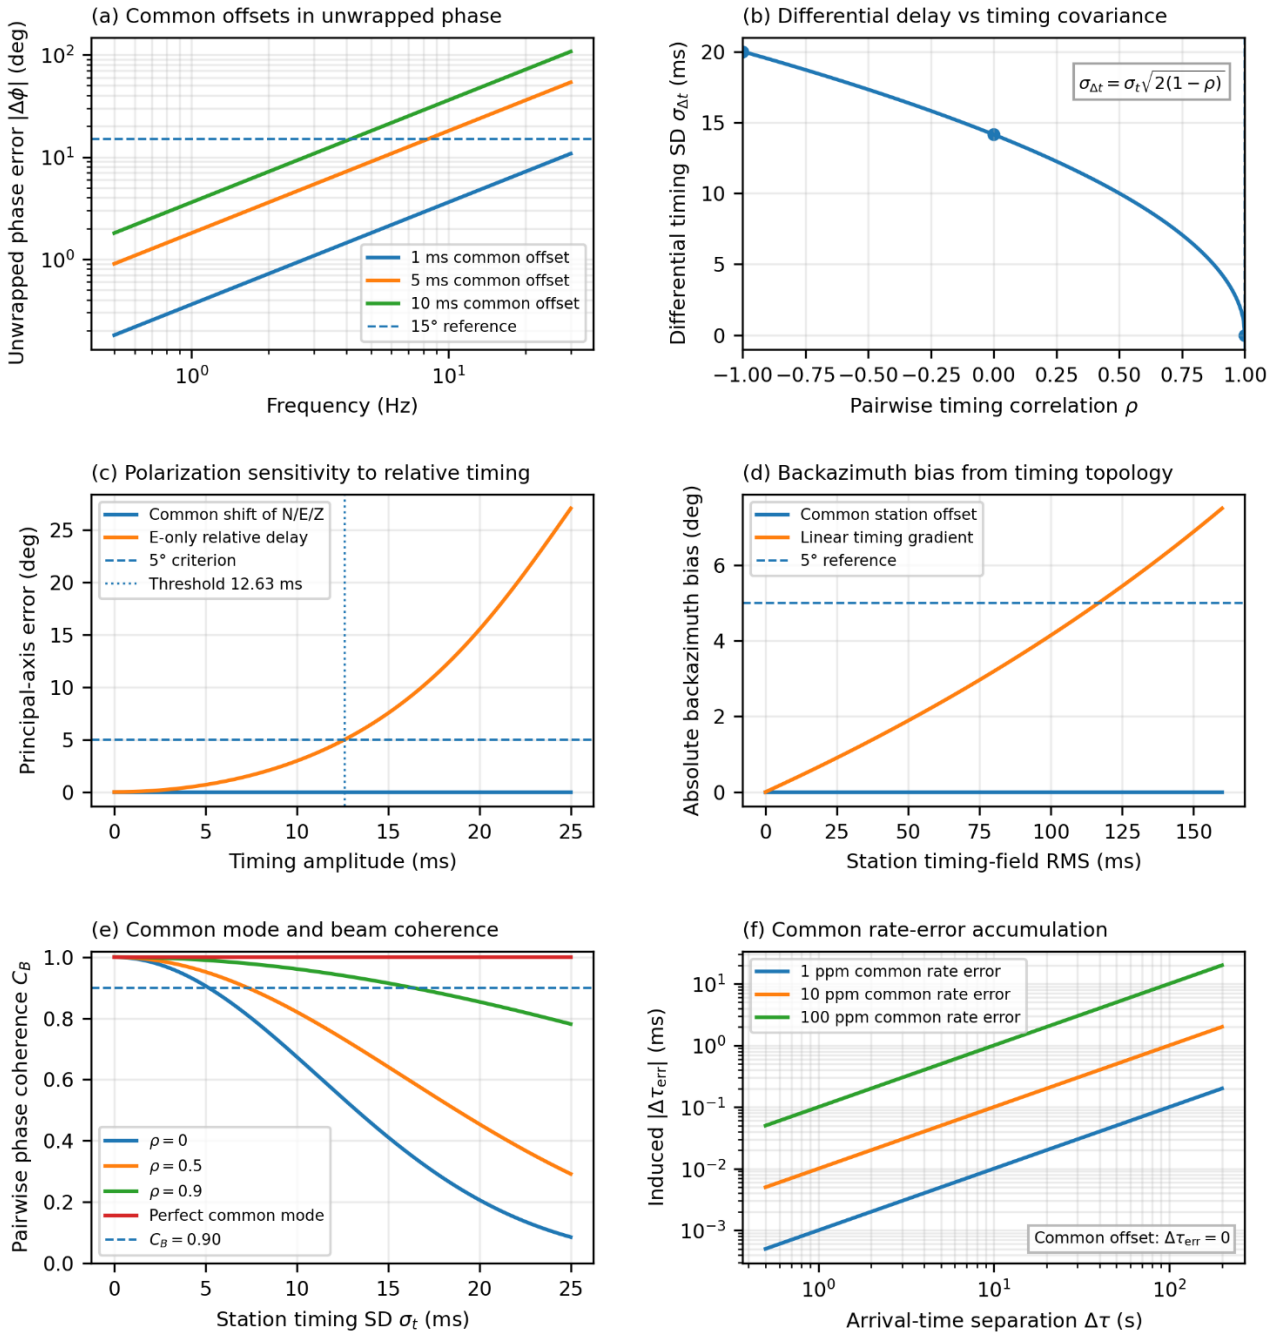


Figure S4. Observable-dependent cancellation and sensitivity of common-mode and differential timing errors. (a) Unwrapped spectral-phase error produced by deterministic constant common-mode offsets of 1, 5, and 10 ms. A common offset remains fully visible to an absolute phase observable even though it may cancel from differential measurements. (b) Differential timing standard deviation for two channels or stations with equal marginal timing standard deviation $\sigma_t = 10$ ms and pairwise correlation ρ, following $\sigma_{\Delta t} = \sigma_t\sqrt{2(1-\rho)}$. Perfectly common timing error ($\rho = 1$) cancels, whereas independent and anticorrelated errors produce progressively larger differential uncertainty. (c) Principal-axis polarisation error for the 10 Hz, 40°-azimuth, 45°-incidence long-window covariance sensitivity case evaluated in main-text Fig. 4. A synchronous shift of the N, E, and Z components leaves the intrinsic covariance polarisation unchanged, whereas an E-only relative delay rotates the estimated principal axis; the dashed and dotted lines mark the illustrative 5° criterion and its corresponding first-crossing threshold of 12.63 ms. This intrinsic invariance excludes indirect changes caused by absolute-window misselection on a non-stationary waveform. (d) Absolute backazimuth bias for a 15-station isotropic array with 10 km Root-Mean-Square (RMS) radius, true horizontal slowness $u = 0.08$ s km^{-1}, and true backazimuth 40°. A constant all-station timing offset is absorbed by the origin-time intercept and produces no slowness or backazimuth bias, whereas a spatially linear timing gradient oriented 60°clockwise from North and normalised to the same station timing-field RMS is formally indistinguishable from a slowness

perturbation. The illustrative 5° backazimuth bias occurs at a timing-field RMS of 116.66 ms for this controlled orientation. (e) Pairwise beam phase coherence at 10 Hz for equicorrelated Gaussian station timing errors. Perfect common mode preserves inter-station coherence, whereas decreasing correlation increases coherence loss for the same marginal timing standard deviation. (f) Differential-delay error induced by a constant common fractional-frequency error over arrival-time separations of 0.5–200 s. Unlike a constant common offset, which cancels from differential delays, a common fractional-frequency error produces $|\Delta\tau_{\mathrm{err}}| = |y|\Delta\tau$. The parts-per-million (ppm) values are controlled illustrative cases rather than empirical clock specifications. Together, the panels demonstrate that timing amplitude alone does not determine measurement degradation: the spatial/channel topology and temporal structure of the timing error must be specified together with the observable.

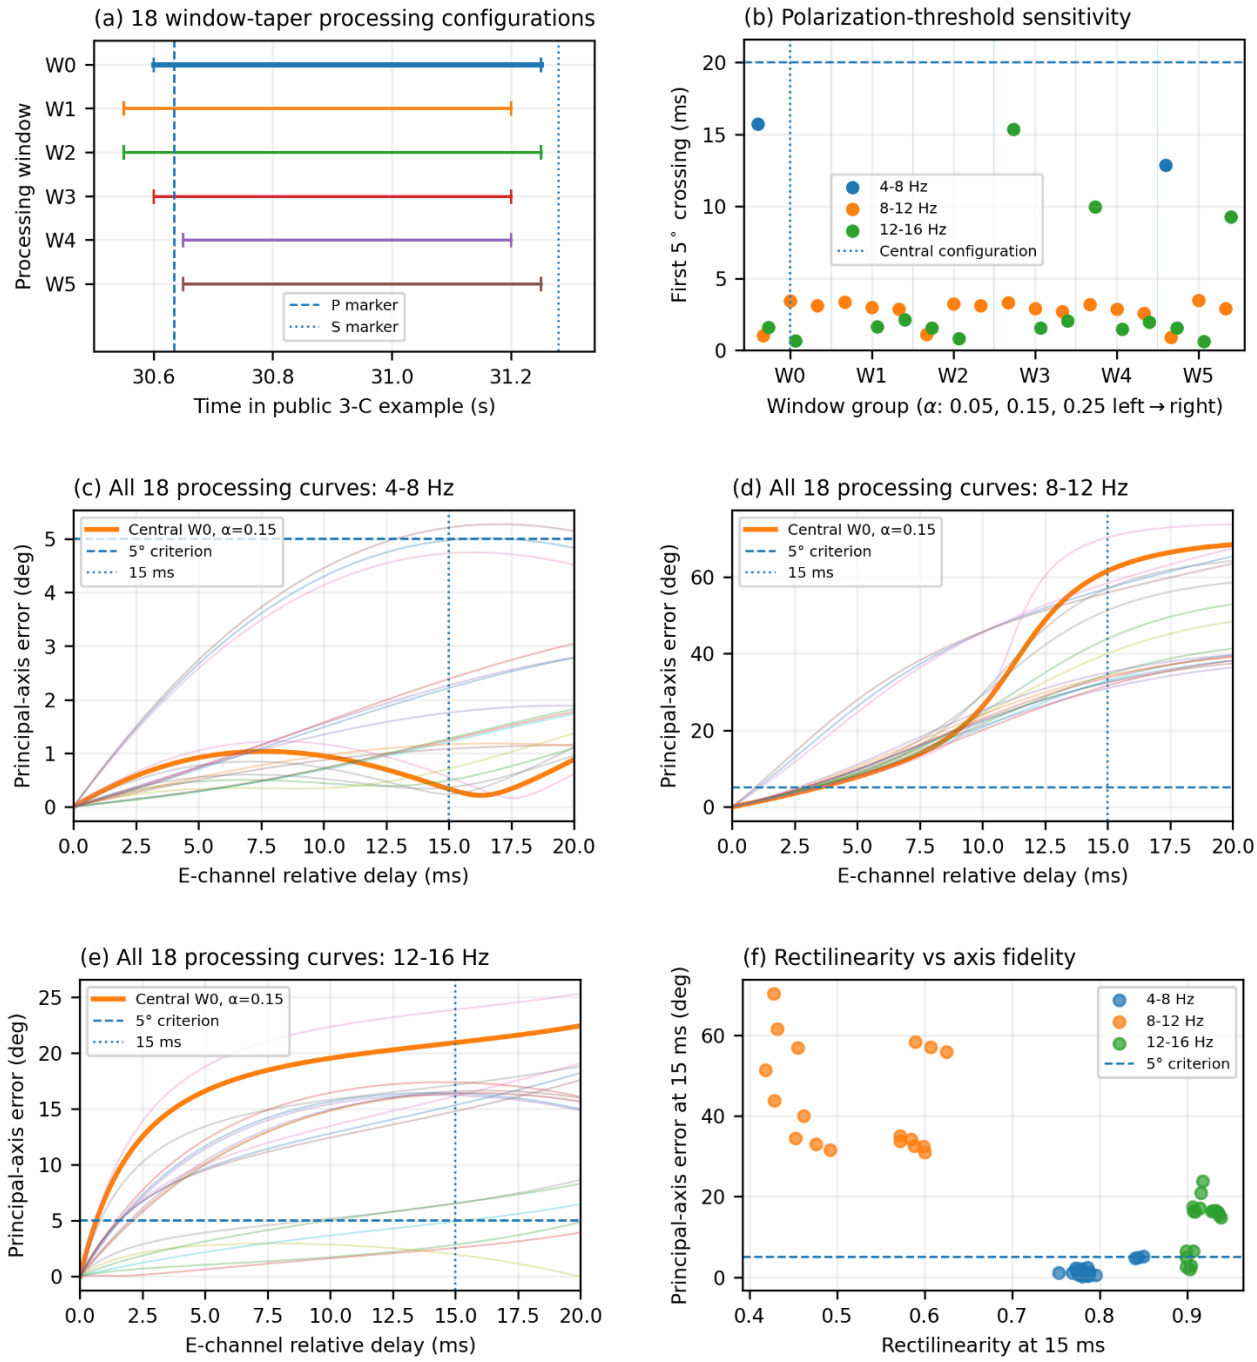


Figure S5. Processing robustness of the controlled relative inter-channel timing experiment applied to the public observational three-component waveform example of main-text Fig. 7. (a) Six pre-S analysis windows combined with Tukey taper parameters $\alpha = 0.05$, 0.15, and 0.25 define 18 processing configurations. The P and S markers are the reproducible temporal markers used in main-text Fig. 7. Across the three frequency bands, the 99.9 % two-sided impulse-response energy radii (0.630–0.685 s) remain smaller than the minimum post-window margin of 0.92 s, so all filter-edge checks pass. (b) First E-channel relative delay at which the principal-axis error reaches the illustrative 5°criterion for each processing configuration. Configurations for which the criterion is not reached within the tested 0–20 ms delay interval are treated as right-censored and are not assigned a finite crossing value in the panel. The vertical dotted line identifies the central processing configuration used in main-text Fig. 7 ($W0$, $\alpha = 0.15$). Because the polarisation response can be non-monotonic with delay, reported values represent the first crossing only. (c–e) Principal-axis error versus imposed E-channel delay for all 18 processing configurations in the 4–8, 8–12, and 12–16 Hz bands, respectively. Thin curves show the processing ensemble, whereas the emphasised curve denotes the central configuration $W0, \alpha = 0.15$; horizontal and vertical reference lines mark 5°and 15 ms, respectively. The 4–8 Hz band remains comparatively insensitive over most processing choices, whereas the two higher-frequency

bands show substantially larger and more configuration-dependent rotations. (f) Principal-axis error versus rectilinearity at 15 ms for all processing configurations. The coexistence of large axis rotations with high rectilinearity, particularly in the 12–16 Hz band, demonstrates that rectilinearity alone is not a sufficient diagnostic of polarisation-direction fidelity. These tests quantify sensitivity to window and taper choices only; the synchronised observational waveform is treated as the reference, and no absolute source-polarisation ground truth is inferred.

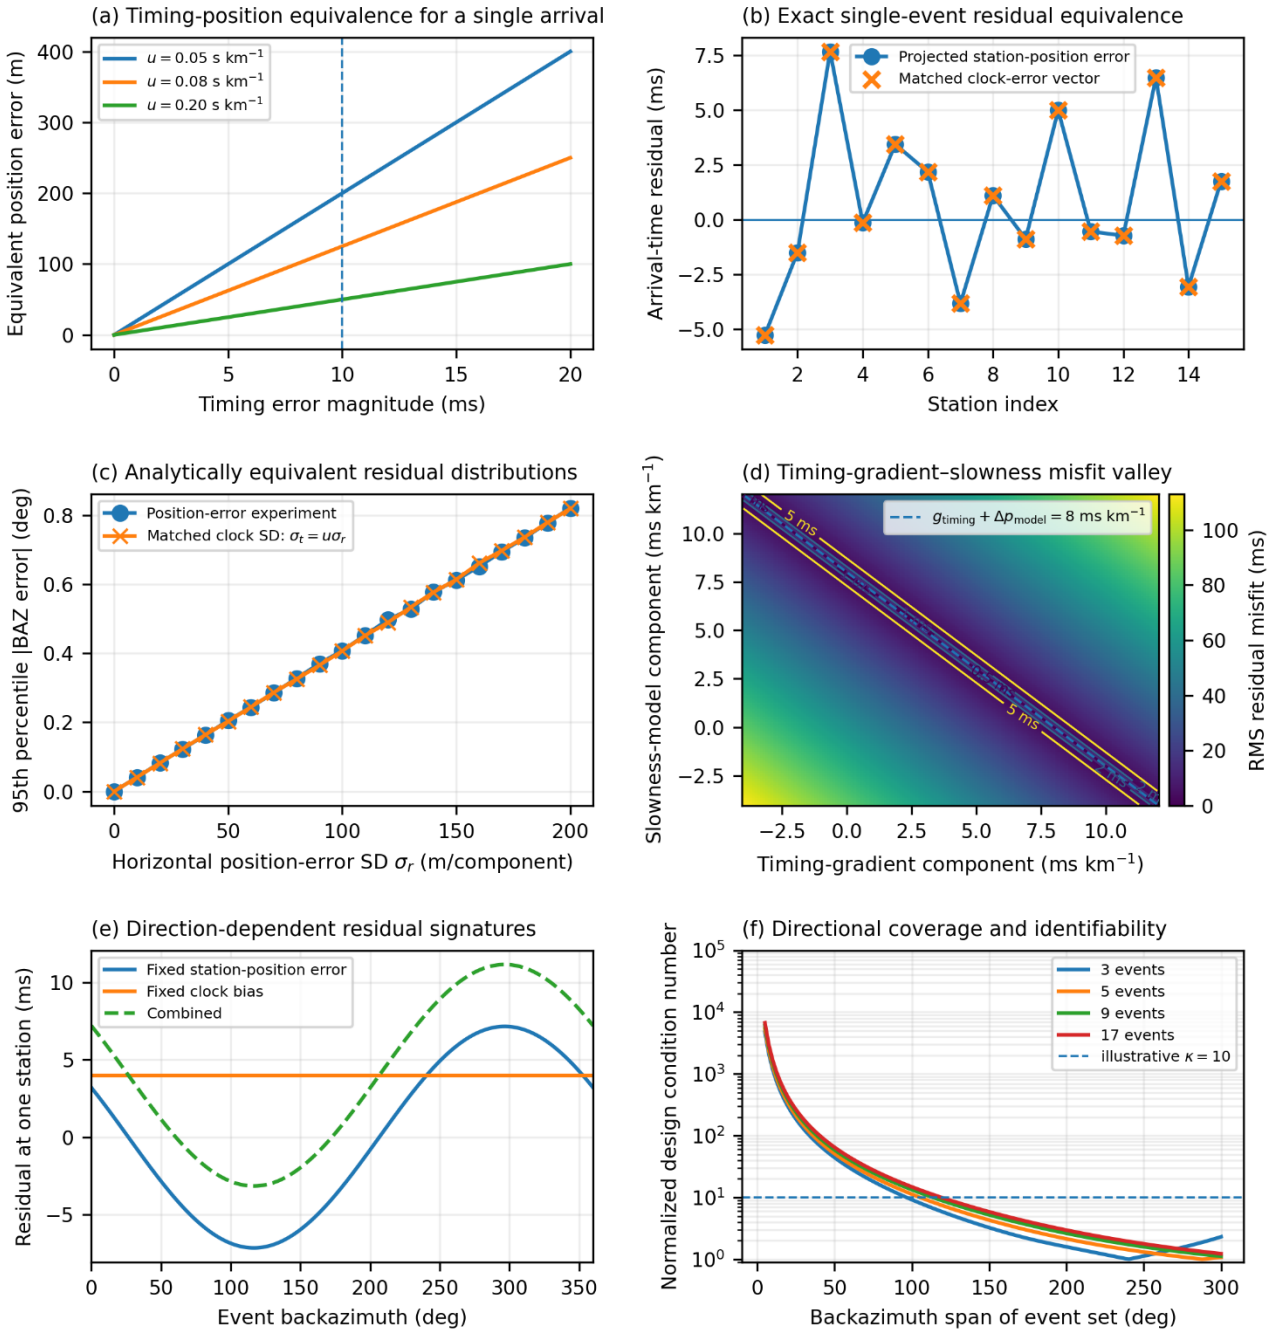


Figure S6. Controlled first-order identifiability tests for timing, station-position, and linear propagation-model/slowness errors in array measurements. All array calculations use the 15-station isotropic geometry of main-text Fig. 5 and, unless otherwise stated, a true horizontal slowness $u = 0.08$s km^{-1} and backazimuth 40°. The station-position error convention is $\delta\mathbf{r} = \mathbf{r}_{\mathrm{true}} - \mathbf{r}_{\mathrm{used}}$, so that the induced first-order arrival-time residual is $\delta t = \mathbf{p}^{\mathrm{T}}\delta\mathbf{r}$. (a) Equivalent station-position error parallel to the slowness vector for a prescribed single-arrival timing error, following $|\,\delta r_{\parallel}\,| = |\,\delta t\,|/u$, shown for three illustrative slownesses. (b) One controlled realisation with independent 50 m horizontal coordinate-error standard deviation per component. Projection of the position errors onto the true slowness vector produces an arrival-time residual vector that is exactly reproduced by the matched station clock-error vector, illustrating the single-event timing–position degeneracy. (c) Ninety-fifth-percentile absolute backazimuth error from Monte Carlo station-position perturbations compared with independent station timing perturbations having matched standard deviation $\sigma_t = u\sigma_r$. Under the stated isotropic Gaussian assumptions, the two perturbation models are analytically distribution-equivalent; the small separation between the curves reflects finite-sample Monte Carlo variability. (d) Root-mean-square arrival-residual misfit when a controlled 8 ms km^{-1} spatial residual gradient is partitioned between a timing-gradient component and a linear propagation-model/slowness component along the same direction. The diagonal zero-misfit valley represents a one-dimensional slice of the more general vector degeneracy $\mathbf{g}_{\mathrm{timing}} + \Delta\mathbf{p}_{\mathrm{model}}$ and

shows that only their sum is identifiable from a single plane-wave arrival under this linearised model. (e) Residual at one station as event backazimuth varies for a fixed horizontal station-position error, a fixed clock bias, and their sum. The clock contribution is direction independent, whereas the position-induced contribution varies with incident slowness direction. (f) Column-normalised condition number of the multi-event design matrix for jointly estimating one stable station clock bias and two horizontal station-position-error components as a function of event backazimuth span for sets of 3, 5, 9, and 17 events. Event reference-time terms are assumed externally constrained or removed, and event slowness vectors are assumed known. Broader and suitably distributed directional coverage improves conditioning and can break the single-event clock–position degeneracy; the dashed $\kappa = 10$line is an illustrative geometry-conditioning reference, not a physical admissibility threshold or a parameter-uncertainty bound. These controlled benchmarks isolate first-order identifiability and do not include wavefront curvature, general three-dimensional heterogeneity, phase misassociation, surveying systematics, event-model errors, or clock-telemetry constraints.