This manuscript is a non–peer-reviewed preprint submitted to EarthArXiv

A prescribed wave probe reveals phase hidden vortex force memory in turbulence

Guoqiang Liu[1, *] and Maryam Al Alshehhi[1]

[1] *Department of Civil, Infrastructural and Environmental Engineering, Khalifa University of Science and Technology, Abu Dhabi, UAE*

(Dated: May 10, 2026)

Phase averaging closes the equations of multi scale flows by removing bilinear couplings with vanishing one time means. Their two time correlations can still survive. We test this distinction for waves passing through turbulence. We study the stochastic vortex force in direct numerical simulation of homogeneous isotropic turbulence. This wave vorticity coupling is removed by classical phase averaging at one time. A prescribed analytic carrier and a modal projection fix the channel before any closure is applied. The resulting two time autocorrelation is finite and integrable. The fixed point measurement tests estimator selectivity. Three lock in controls suppress mismatched alternatives by up to four orders of magnitude. The trajectory measurement tests scale and geometry. In the audited clean support branch, the normalized response is concentrated near the wavelength matched eddy scale. A geometry control replaces the transverse SVF projection with the longitudinal vorticity component. The response peak then moves to a wavelength mismatched shell. In that channel, the mismatched shell is eighteen times stronger than the wavelength matched shell. This rules out a generic vorticity explanation. Sampling along the wave group velocity trajectory gives much faster decorrelation than the matched stationary baseline. The shortening is at least an order of magnitude. This contrast combines scale selection with moving frame sampling.The measurement recovers a real two time object inside a coupling that traditional averaging removes at one time. The same prescribed probe protocol can extract channel resolved memory from other phase suppressed bilinear couplings.

I. INTRODUCTION

Reduced descriptions of multi-scale dynamical systems routinely eliminate bilinear couplings between fast and slow variables when those couplings have vanishing one-time mean. Phase averaging in wave–mean flow theory [1–3], ensemble averaging in Reynolds-stress closures, and projection-based elimination of fast variables in Mori–Zwanzig and related formalisms [4, 5] all share this structure. A named coupling between a slow mode and a fast bath is set aside because its one-time average vanishes. The closed reduced description retains only the surviving mean-field terms.

A vanishing one-time mean does not settle the two-time problem. The eliminated coupling is the bilinear product of carrier and bath. Its two-time autocorrelation is not the product of the one-time means. Schematically,

$$\begin{aligned}&\big\langle \boldsymbol{u}_\phi(t)\,\zeta(t)\,\boldsymbol{u}_\phi(t+\tau)\,\zeta(t+\tau)\big\rangle \\ &\quad\neq \big\langle \boldsymbol{u}_\phi(t)\big\rangle\big\langle\zeta(t)\big\rangle\big\langle \boldsymbol{u}_\phi(t+\tau)\big\rangle\big\langle\zeta(t+\tau)\big\rangle,\end{aligned} \tag{1}$$

where ζ is a representative bath quantity and $\boldsymbol{u}_\phi$ the carrier. The right-hand side vanishes whenever any one-time mean vanishes. The left-hand side need not. The joint two-time correlation between carrier and bath at the bilinear level survives the elimination of the one-time mean. The Mori–Zwanzig formalism reaches the same conclusion in the abstract. Projected-out degrees of freedom can re-enter the reduced dynamics through memory kernels [4, 5]. Whether a *particular* eliminated coupling carries transport-relevant two-time content in a *specific projection* of a real system is a different question. The reduced description cannot answer it. The question has to be settled by direct measurement of the mode-projected force autocorrelation. Standard turbulence diagnostics return undifferentiated bath statistics and cannot, by themselves, isolate a specific bilinear channel's two-time content from generic Eulerian decorrelation.

We address this by introducing a *prescribed-probe* protocol. An externally prescribed analytic carrier defines the phase information of the eliminated coupling by construction. The bilinear product of carrier and bath is then projected onto a modal channel. Fixed-point estimator controls test whether the measured two-time signal requires temporal ordering and matched demodulation rather than only one-point bath statistics. The controls are a low-frequency baseline, time shuffling, and frequency detuning. Geometric channel specificity is tested separately by replacing the transverse SVF projection with alternative vorticity projections in the trajectory branch. The construction returns a channel-resolved force autocorrelation as a measurable object before any closure is imposed. A trajectory-side, shell-filtered extension then exposes how that response is organized in scale and along the moving slow mode. Throughout this work, "response" denotes the Green–Kubo/linear-response force autocorrelation associated with the prescribed probe. It does not mean that the DNS bath is dynamically forced by a resolved wave or that wave back-

* guoqiang.liu@ku.ac.ae

reaction is included.

As a non-trivial instantiation we apply the protocol to the wave–vorticity channel in homogeneous isotropic turbulence. When a surface wave passes through a turbulent flow, its orbital velocity $\boldsymbol{u}_\phi$ acts on the fluctuating vorticity $\boldsymbol{\omega}'$ through the bilinear *stochastic vortex force* (SVF) $\boldsymbol{f}_{\mathrm{svf}} = \boldsymbol{u}_\phi \times \boldsymbol{\omega}'$. This is the genuinely fluctuating piece of the wave–current coupling that arises naturally in a Helmholtz decomposition of the Navier–Stokes equations [6]. It must be distinguished from the Craik–Leibovich vortex force $\boldsymbol{f}_{\mathrm{CL}} = \boldsymbol{u}_{\mathrm{S}} \times \langle\boldsymbol{\omega}\rangle$, which couples the wave-averaged Stokes drift $\boldsymbol{u}_{\mathrm{S}}$ to the mean vorticity and survives phase averaging as a closed mean-field term [1]. The SVF is annihilated at the one-time level by literal phase averaging because $\langle\boldsymbol{u}_\phi\rangle_\phi = 0$. Classical wave-averaged descriptions therefore do not retain it in the closed mean equations [1–3, 7, 8]. The SVF is a concrete, named instantiation of a phase-suppressed bilinear coupling between a wave mode and a turbulent bath. It is the exemplar of the structural question above on which the prescribed-probe protocol can be tested.

The present work delivers three specific results. First, we apply the prescribed wave probe to Navier–Stokes turbulence data, measure the mode-projected force autocorrelation in the SVF channel, and use fixed-point estimator controls to rule out one-point and mismatched-demodulation alternatives. Second, we sample the same channel along the wave group-velocity trajectory through a log-Gaussian shell filter and resolve a scale-selected normalized response time with a clear maximum at the nominal wavelength-matched shell in the audited clean-support window. A trajectory-side channel-geometry control replaces the SVF projection by longitudinal vorticity on the same trajectory and shell filter, obtains a qualitatively different scale ordering, and so isolates the geometric specificity of the SVF response. Third, an encounter-time test advances the probe at the wave group velocity rather than holding it fixed. The shell-filtered moving response has memory at least 9.8 times shorter than the matched fixed-point record in the adjacent audited window. The full construction is summarized schematically in Fig. 1.

Related wave–turbulence and upper-ocean studies have long established that surface waves organize turbulent structure and mixing without breaking [9–13]. The channel-resolved two-time response measured here is complementary to that body of work. It isolates the SVF coupling as a two-time object that those measurements were not designed to extract. The trajectory-side channel-geometry control adds an independent test of geometric specificity beyond the three fixed-point controls.

The time integral of the mode-projected force autocorrelation is the bath-side ingredient that a Green–Kubo description of wave–bath energy transfer can build on. The Green–Kubo framework rests on stationarity and Bochner's theorem and is independent of the bath spectral form. A separate theoretical analysis [6] develops this connection for remote ocean swell. That work substitutes a Kolmogorov inertial-range spectrum to derive a specific parameter-free attenuation law $\mu_E(\omega) \propto \omega^{8/3}$ tested against 241 trans-oceanic satellite tracks. The Kolmogorov substitution and the oceanographic context both lie outside the present scope. The DNS measurement reported here addresses the existence and structure of the bath-side autocorrelation only. It stands independently of any inertial-range closure.

II. PHYSICAL PLATFORM

All measurements use the `isotropic1024coarse` dataset on the Johns Hopkins Turbulence Database (JHTDB) [14–17]. The flow is statistically stationary homogeneous isotropic turbulence at $\mathrm{Re}_\lambda \approx 433$, $\varepsilon = 0.093$, $\nu = 1.85 \times 10^{-4}$, on a $(2\pi)^3$ periodic domain, a canonical setting for scale-local turbulence and Eulerian decorrelation studies [18–20]. For the long fixed-point records, we retrieve the three-component vorticity $\boldsymbol{\omega}'(\boldsymbol{x}, t)$ at a cadence $\Delta t = 0.002$ for $N_t = 5028$ frames ($T \approx 10$ in DNS time units). Shorter trajectory and shell-filtered products are specified separately below. Further details on the data pipeline and probe construction are provided in the Supplemental Material [21].

The final paired-bootstrap shell-ordering scan and the matched shell-filtered fixed-point check use JHTDB `getData` point queries with chunk size 4096. Server-side cutout data are used only for the explicitly labelled secondary $n = 128$, $\sigma = 0.45$ T/L/W guardrail and are not used to promote the central scale-selection branch.

The fluctuating vorticity field $\boldsymbol{\omega}'(\boldsymbol{x}, t)$ is retrieved from the JHTDB `getData` vorticity service, which evaluates the curl of the stored velocity field using fourth-order centered finite differences on the native simulation grid before spatial interpolation. The documented root-mean-square uncertainty of this service relative to a direct spectral curl on the same grid is approximately 7% of the local vorticity magnitude [14, 16]. Spatial interpolation uses cubic splines and temporal interpolation uses PCHIP, both as supplied by the JHTDB query interface. The same provenance is used for both the fixed-point point time series and the local FFT cubes underlying the trajectory shell-filter measurements. We have not retrieved or stored full velocity volumes; consequently, all spectrum-weighted audit columns in the Supplemental Material [21] use the public isotropic spectrum in `spatial_k56_data.npz` as a radial proxy and are not treated as per-run measurements.

The present measurement is carried out at a single $\mathrm{Re}_\lambda \approx 433$. The channel specificity and memory dependence established here are qualitative properties of the prescribed-probe response protocol. This statement is not itself an inertial-range closure. What requires high Re and a developed inertial interval is the asymptotic scale-selective closure of the memory kernel [18, 19]. The quantitative Re-dependence of the DNS response amplitude is deferred to subsequent multi-Re work.

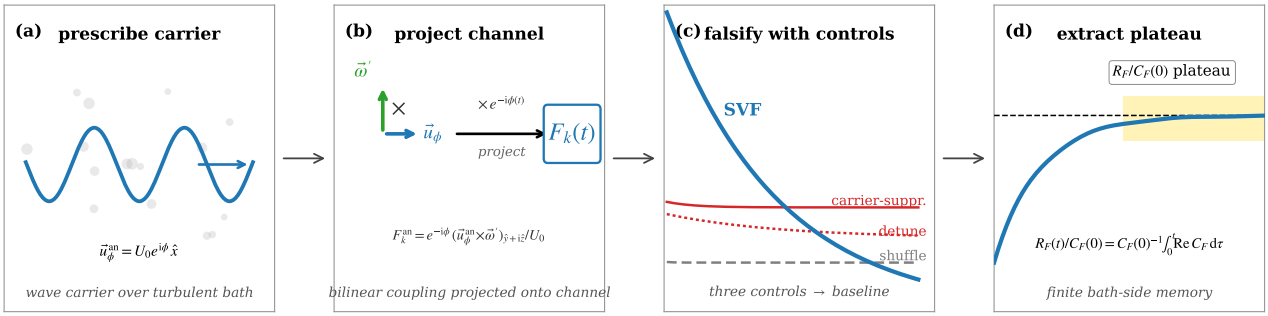

Figure 1. **Methodology. A channel-resolved two-time response protocol.** (a) A prescribed analytic carrier $\boldsymbol{u}_\phi^{\mathrm{an}} = U_0 e^{\mathrm{i}\phi}\hat{\boldsymbol{x}}$ is inserted into a turbulent bath of vorticity fluctuations $\boldsymbol{\omega}'$ (gray). The carrier is equivalent to a cosine–sine quadrature pair. The always-on Fourier mode shown here is the operational basis of the prescribed probe and is the mode-projection representation of the wave–eddy encounter physics developed in [6]. A moving wave packet samples local turbulent vorticity along its trajectory. The time-integrated force autocorrelation at fixed wavenumber k is the Green–Kubo input that the present DNS extracts. This use of an always-on modal carrier is analogous to modal resolved-variable constructions in Mori–Zwanzig theory [4, 5]. (b) The bilinear coupling $\boldsymbol{u}_\phi^{\mathrm{an}} \times \boldsymbol{\omega}'$ is projected onto the target channel by the modal factor $e^{-\mathrm{i}\phi(t)}$, yielding the demodulated force $F_k(t) = -\omega_z' + \mathrm{i}\omega_y'$. The analytic-carrier construction removes the mirror-frequency self-mixing term that would arise from a real-cosine carrier. (c) Three fixed-point controls—a low-frequency baseline, time shuffling, and a frequency-detuned probe—show complementary behavior. The low-frequency baseline collapses to the shuffled-noise floor. Time shuffling suppresses the signal by more than an order of magnitude. The detuned probe suppresses the signal by more than four orders of magnitude in the alias-free case. Together they quantify the selectivity of the fixed-point lock-in estimator: temporal ordering, matched demodulation, and suppression of carrier-independent low-frequency content. (d) The running response integral $R_F(t) = \int_0^t \mathrm{Re}\, C_F(\tau)\, \mathrm{d}\tau$ reaches a finite plateau in the lock-in baseline. The fixed-point branch fixes the bath-side response of the SVF channel and tests the controls. The trajectory branch resolves its scale selection and encounter-time organization.

We prescribe an analytic wave-like carrier, implemented as the cosine–sine quadrature pair $\boldsymbol{u}_\phi^{(c)} = U_0 \cos\phi\, \hat{\boldsymbol{x}}$ and $\boldsymbol{u}_\phi^{(s)} = U_0 \sin\phi\, \hat{\boldsymbol{x}}$, with $\phi = \boldsymbol{k}\cdot\boldsymbol{x} - \omega t$, deep-water dispersion $\omega = \sqrt{g_{\mathrm{ref}} k}$, and group velocity $c_g = \omega/(2k)$. The demodulated channel force is

$$F_k(t) = \frac{e^{-\mathrm{i}\phi(t)}}{U_0}\left([\,\boldsymbol{u}_\phi^{(c)} \times \boldsymbol{\omega}'\,] + \mathrm{i}[\,\boldsymbol{u}_\phi^{(s)} \times \boldsymbol{\omega}'\,]\right)\cdot(\hat{\boldsymbol{y}} + \mathrm{i}\hat{\boldsymbol{z}}) = -\omega_z' + \mathrm{i}\omega_y'. \tag{2}$$

This analytic lock-in construction removes the mirror sideband present in a real-cosine carrier. The fixed-point reduction $F_k = -\omega_z' + \mathrm{i}\omega_y'$ is a direct consequence of the always-on analytic-carrier limit and is not generic to the protocol. With a finite wave-packet envelope $A(t)$ replacing the always-on construction, the same projection would give $F_k(t) = A(t)\zeta(t)$ in the same notation $\zeta(t) \equiv -\omega_z'(t) + \mathrm{i}\omega_y'(t)$, and the two-time object would be

$$\langle F_k^*(t)\, F_k(t+\tau)\rangle = A^*(t)\, A(t+\tau)\, \langle \zeta^*(t)\, \zeta(t+\tau)\rangle, \tag{3}$$

with a deterministic envelope gate that does not factor out. The always-on limit $A \equiv 1$ collapses this gate identically and leaves the bath autocorrelation in channel coordinates. The fixed-point branch is therefore the infinite-envelope limit of the prescribed-probe protocol, not a follow-wave-packet measurement. The two are distinct constructions. A finite-packet extension is left to subsequent work. In this limit, the apparent reduction to a vorticity autocorrelation is precisely the projected bath object selected by the SVF vertex. The fixed-point controls below test the selectivity of the lock-in estimator: temporal ordering, matched demodulation, and independence from one-point marginal statistics. They do not by themselves prove geometric channel specificity, which is tested by the trajectory-side channel replacement. The trajectory branch tests how the same projected channel is reorganized by shell selection and wave-speed sampling. It is the construction in which the carrier parameters k and c_g re-enter the measurement. As a diagnostic decomposition, a single real-carrier demodulation in the same notation gives

$$F_k^{\mathrm{raw}}(t) = \frac{1}{2}\zeta(t) + \frac{1}{2}e^{-2\mathrm{i}\phi(t)}\zeta(t), \tag{4}$$

namely a coherent baseband term plus a mirror sideband. The raw signal is used only as a diagnostic decomposition. The quantitative fixed-point response below uses the analytic observable $F_k = \zeta$.

To make the limit structure concrete and to demonstrate that the fixed-point branch is one member of a parametric family of prescribed-probe measurements, we apply finite-envelope versions of the protocol as a post-processing operation on the same fixed-point time series. Gaussian envelopes $A_N(t - t_j) = \exp[-(t-t_j)^2/(2\sigma_N^2)]$ of width $\sigma_N = NT_{\mathrm{wave}}/2$ are applied at multiple centers t_j within the record interior, and the resulting gated signal $F_k^{(N)}(t) = A_N(t - t_j)\, F_k(t)$ is autocorrelated and

center-averaged. The two-time autocorrelation of any time-averaged estimator built this way obeys

$$\frac{C_F^{(N)}(\tau)}{C_F^{(N)}(0)} \approx R_{A,N}(\tau)\,\frac{C_\zeta(\tau)}{C_\zeta(0)}, \qquad R_{A,N}(\tau) = \frac{\sum_t A_N(t)A_N(t+\tau)}{\sum_t A_N^2(t)}, \tag{5}$$

where $R_{A,N}(\tau)$ is the deterministic gate-overlap function and $C_\zeta(\tau)$ is the always-on bath correlation measured directly. The two diagnostics in Fig. 2 demonstrate the protocol family. Panel (a) shows that the measured kernel $C_F^{(N)}(\tau)/C_F^{(N)}(0)$ contracts as N decreases, consistent with the gate-overlap prediction. The $N = 5$ envelope is gate-limited. The $N = 25$ envelope already approaches the always-on bath correlation. The always-on result of the main measurement is the $N \to \infty$ limit of the same family. Panel (b) reduces this to a single crossover diagnostic. The normalized response time $\mathcal{T}_F(N)/\mathcal{T}_F(\infty)$ rises monotonically from the gate-dominated regime at $N = O(1)$ to unity for $N \gtrsim N_{\rm cyc}^E \approx 24$, the bath-memory cycle count read off the always-on measurement. Nine measured values of N collapse onto the analytical prediction $\int R_{A,N}(\tau)\, C_\zeta(\tau)\, \mathrm{d}\tau$ (solid grey) with no free parameters. The protocol family parameterized by N smoothly connects the gate-limited and bath-limited regimes. Numerical verification of the factorization (5) is reported in the Supplemental Material [21].

Two configurations are used: *fixed-point* sampling at $N_p = 200$ stationary positions, and *trajectory* sampling along $\boldsymbol{x}(t) = \boldsymbol{x}_0 + c_g\, t\, \hat{\boldsymbol{x}}$ with $N_{\rm traj} = 200$ trajectories distributed over four values of $\delta = u'(k)/c_g \in \{0.01, 0.03, 0.05, 0.10\}$. In the present fixed-$k$ scans these quoted δ values are kinematic sweep labels realized by varying $g_{\rm ref}$ and therefore c_g. They are not independent re-estimates of a bath velocity diagnostic in each run.

For the scale-selected trajectory measurement we use the complex transverse vorticity

$$W(\boldsymbol{x}, t) = -\omega_z'(\boldsymbol{x}, t) + \mathrm{i}\omega_y'(\boldsymbol{x}, t), \tag{6}$$

filter it in Fourier space with a log-Gaussian shell window $A_\alpha(q/k)$ satisfying $|A_\alpha|^2 = S_\alpha$, and sample the filtered field along the group-velocity trajectory:

$$W_{k,\alpha}^{\rm sel}(\boldsymbol{x}, t) = \mathcal{F}^{-1}\left[A_\alpha(|\boldsymbol{q}|/k)\,\widehat{W}(\boldsymbol{q}, t)\right], \qquad F_{k,\alpha}^{\rm sel}(t) = W_{k,\alpha}^{\rm sel}(\boldsymbol{x}_0 + c_g(k)t\hat{\boldsymbol{x}}, t). \tag{7}$$

The shell center is $q_c = \alpha k$ with $\alpha \in \{0.5, 1, 2\}$. The shell window is not a resonance ansatz. It localizes the measurement in scale space and tests whether the moving-wave response is preferentially carried by eddies of size comparable to the wavelength.

The analysis below addresses three distinct objects associated with the same SVF channel. **(i)** The fixed-point autocorrelation $\langle F_k^*(t)F_k(t+\tau)\rangle$, which integrates over all bath wavenumbers and answers whether the channel carries a finite, controlled two-time response. **(ii)** The trajectory-side, shell-filtered proxy $F_{k,\alpha}^{\rm sel}(t)$, which adds wave–eddy geometry and encounter timing and answers how that response is organized in scale and along the moving wave. **(iii)** The closed macroscopic kernel $\Gamma_{\rm th}(k)$ derived in [6], which is obtained from the fixed-point autocorrelation by a mode-projected scale-selection window and Kolmogorov inertial-range closure, neither of which is imposed here. Distinguishing these three objects is essential to the present claim, which concerns the bath-side autocorrelation in the SVF channel and its scale and encounter organization, not the closed transport coefficient.

III. ANALYTIC FIXED-POINT LOCK-IN BASELINE

The channel-resolved two-time function $C_F(\tau) = \langle F_k^*(t)F_k(t+\tau)\rangle$ averaged over 200 fixed positions at $k \in \{10, 15, 30\}$ decays on a timescale much longer than the carrier period [Fig. 3]. By construction in the always-on limit, the demodulated observable is the transverse-vorticity autocorrelation in the SVF channel coordinate. This reduction is intentional. The SVF vertex selects the transverse-vorticity bath coordinate whose two-time autocorrelation is the Green–Kubo bath-side ingredient. The fixed-point branch therefore tests the existence, integrability, and estimator selectivity of this projected bath object. The associated (k, ω) independence of the fixed-point plateau is a direct consequence of this projection geometry, not an empirical scaling result; it is the expected behavior of the mode-projected bath autocorrelation and is not a test of any inertial-range closure (this is reiterated in Sec. VII). Geometric channel specificity is tested separately by the moving, shell-filtered branch in Secs. V–VI.

Three controls test the discriminating power of the lock-in hierarchy. A *low-frequency baseline* low-pass filters the carrier-modulated force components at cutoff 0.4ω before modal projection, removing the carrier-locked spectral content that gives rise to SVF coherence and collapsing the envelope to the shuffled-noise floor (suppression factor $\sim 6.9 \times 10^4$; see Supplemental Material [21] for the precise algorithmic ordering). This low-frequency control is not a literal implementation of the classical phase-average operator. Rather, it is a nontrivial baseline that retains carrier-independent bath variability while suppressing carrier-linked oscillatory content (see Supplemental Material [21]). A *time-shuffled* record replaces $F_k(t)$ by a random permutation $F_k(\pi(t))$. Shuffling preserves the one-point distribution of the demodulated samples, including variance and higher one-point moments, while destroying their temporal order. The resulting response integral drops by more than an order of magnitude relative to the SVF channel, ruling out explanations based solely on the marginal statistics of F_k. A *detuned probe* drives the carrier at $\omega_\perp = (1+\sqrt{2})\omega$ while projecting at ω. The cleanest quantitative discrim-

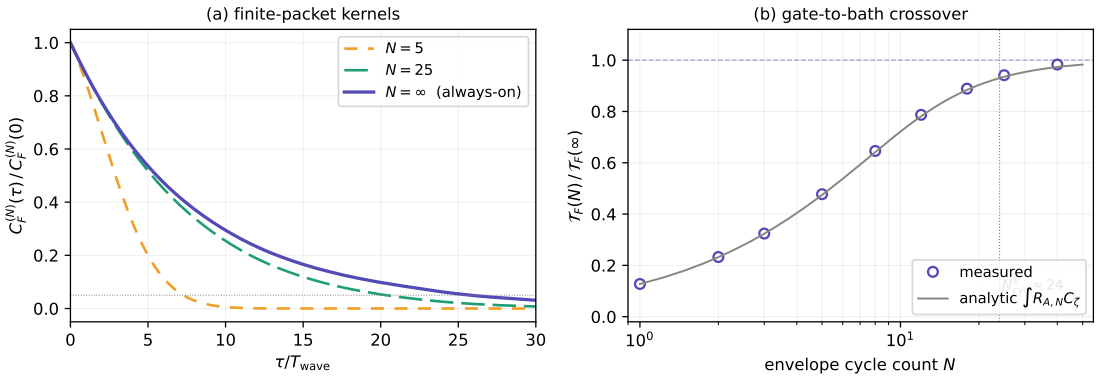

Figure 2. **Finite-envelope demonstration of the protocol family.** Gaussian envelopes of cycle count N are applied as a post-processing step to the same fixed-point time series used for the always-on measurement. **(a)** The measured normalized correlation $C_F^{(N)}(\tau)/C_F^{(N)}(0)$ contracts smoothly as N decreases. The $N = 5$ envelope is gate-limited. The $N = 25$ envelope approaches the always-on bath correlation. An intermediate $N = 10$ curve is shown for reference. The always-on measurement is the $N \to \infty$ limit of the same construction. **(b)** The normalized response time $\mathcal{T}_F(N)/\mathcal{T}_F(\infty)$ undergoes a single crossover from the gate-dominated regime at $N = O(1)$ to unity for $N \gtrsim N_{\mathrm{cyc}}^E$. Nine measured values of N (circles) collapse onto the parameter-free analytical prediction $\int R_{A,N}(\tau)\, C_\zeta(\tau)\, \mathrm{d}\tau$ (solid line). The fixed-point measurement is one member of a parametric family, not an isolated construction.

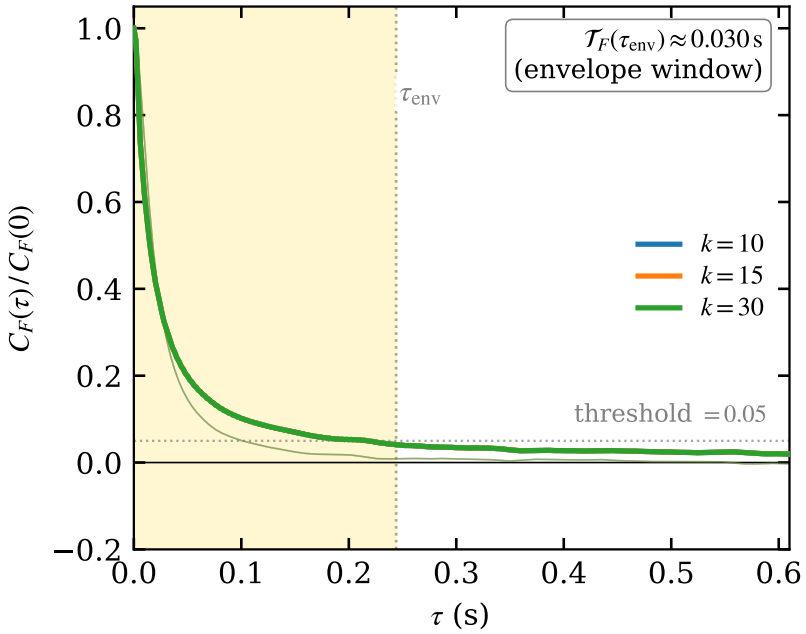

Figure 3. **SVF fixed-point lock-in baseline.** Normalized force autocorrelation $C_F(\tau)/C_F(0)$ averaged over 200 fixed positions for $k \in \{10, 15, 30\}$ at $g_{\mathrm{ref}} = 38{,}269$. All times are JHTDB/DNS time units. Light curves are the raw signal. Solid curves are the Hilbert envelope. The analytic carrier eliminates the real-cosine mirror sideband, so the displayed quantity is the demodulated bath-side autocorrelation in the SVF channel. The envelope defines the operational decorrelation window used for response integration and verifies that the lock-in estimator produces a finite, integrable baseline signal.

ination is in the alias-free case $k = 10$, where the response magnitude is reduced by more than four orders of magnitude. The remaining tested values $k = 15, 30$ show consistent but less clean suppression because of aliasing (discussed in [21]). Together these controls establish the selectivity of the fixed-point lock-in estimator. The signal requires temporal ordering and matched demodulation, and it is not determined by the one-point distribution of the demodulated samples alone. The reported suppression ratios use the envelope-window estimator $\mathcal{T}_{F,\mathrm{env}}$ defined in Eq. (S6) of the Supplemental Material [21], with the matched-SVF $C_F(0)$ as the com-

mon normalization within each control comparison; absolute ratios should be read in this convention. They do not, by themselves, distinguish the transverse SVF geometry from all possible vorticity correlations. That distinction is supplied by the trajectory-side longitudinal-channel control in Sec. V. This estimator-selectivity baseline is compatible with rapid-distortion analyses of wave–turbulence interaction [22, 23] and with broader studies of wave-organized turbulence and Langmuir-type responses [9, 24]. The present measurement extracts a two-time response integral rather than a distortion-induced energy-transfer rate.

IV. THE FIXED-POINT BATH-SIDE RESPONSE IS FINITE AND INTEGRABLE

The running integral

$$R_F(t) = \int_0^t \mathrm{Re}\, C_F(\tau)\, \mathrm{d}\tau, \qquad \mathcal{T}_F(t) = \frac{R_F(t)}{C_F(0)} \tag{8}$$

reaches a finite plateau beyond $t \sim \tau_{\mathrm{env}}$ and remains within the position-block bootstrap band out to $t \gtrsim 5\tau_{\mathrm{env}}$ for all three k. The normalized quantity $\mathcal{T}_F$ has units of time and is the normalized two-time response time of the SVF channel. This is the bath-side input that enters a Green–Kubo description of wave–bath energy transfer through the same force autocorrelation [6, 25–27]. The present DNS measurement establishes the existence and finiteness of this bath-side autocorrelation and tests the estimator selectivity of the fixed-point projection. Geometric channel specificity is tested separately by the trajectory-side longitudinal-channel control. The conversion from this autocorrelation to a macroscopic transport coefficient is not attempted here. It is treated separately in [6], where Bochner's theorem guarantees a non-negative damping rate.

It is useful to express this fixed-point decorrelation window in carrier-cycle units. Define the bath-memory cycle count

$$N^E_{\mathrm{cyc}}(k) \equiv \frac{\tau^E_{\mathrm{env}}(k)}{T_{\mathrm{wave}}(k)} = \frac{\omega(k)\tau^E_{\mathrm{env}}(k)}{2\pi}. \tag{9}$$

For the analytic fixed-point records in Fig. 3, $\tau^E_{\mathrm{env}} = 0.244$ DNS time units, giving $N^E_{\mathrm{cyc}} = 24.0, 29.4, 41.6$ for $k = 10, 15, 30$, respectively. This is the timing content of the lock-in baseline. The all-q bath-memory window spans many carrier periods. The fixed-point observable performs no spatial scale selection. N^E_{cyc} is therefore a statement about temporal separation between the prescribed carrier and the all-q bath memory. It is not a statement about how many wave periods are spent inside an eddy of a particular size. That latter question belongs to the moving, scale-selected trajectory construction below.

V. SCALE-SELECTED TRAJECTORY RESPONSE

When the probe advances at the wave group velocity rather than remaining stationary, the same SVF channel exposes two new ingredients absent from the fixed-point measurement. The first is wave–eddy geometry. It asks which scale of bath fluctuations dominates the response along the trajectory. The second is encounter timing. It asks how long the moving wave remains in the same scale-selected eddy. The scale-selected trajectory observable in Eq. (7) samples the bath space–time correlation along $\boldsymbol{r} = c_g\tau\,\hat{\boldsymbol{x}}$ through a log-Gaussian shell filter A_α localized at $q_c = \alpha k$. We evaluate it on a coarse periodic 64^3 grid using $N_p = 64$ trajectories, $k \in \{10, 15, 30\}$, $\alpha = q_c/k \in \{0.5, 1, 2\}$, and log-Gaussian shell widths $\sigma \in \{0.30, 0.45\}$ over $N_t = 128$ snapshots ($t_{\mathrm{end}} = 0.254$). Because this trajectory measurement uses a coarse local FFT grid, we first audit each shell window before using it in the central claim. The class labels are defined from unweighted shell-window diagnostics on the resolved grid. Let η_{supp} be the resolved support fraction and η_{edge} the fraction of resolved window weight with $\max_i |q_i| \geq 30$, a conservative two-grid-point buffer below the component Nyquist value $q_N = 32$. Class C denotes an out-of-support shell center with $q_c > \sqrt{3}\, q_N$. Class B denotes $\eta_{\mathrm{supp}} < 0.75$ or $\eta_{\mathrm{edge}} \geq 0.25$. All remaining filters are class A. Spectrum-weighted audit columns are reported only as supplementary diagnostics using the isotropic spectrum in `spatial_k56_data.npz` as a proxy, not the unavailable per-run $\widehat{W}$ fields, and are not used to assign the class labels. Because the finite log-Gaussian window and the coarse grid shift the actual support, the label α denotes the nominal shell center $q_c = \alpha k$, not an exact effective center. For the clean $k = 10$, $\sigma = 0.30$ branch, the audit gives $q_{\mathrm{eff}}/q_c = 1.27$–$1.37$ for the unweighted window and 1.15–1.20 with the isotropic spectrum proxy. Thus the nominal $q_c = k$ filter samples an effective support $q_{\mathrm{eff}}/k \simeq 1.2$–$1.4$. The central claim is therefore a finite-band wave-matched eddy-scale result, not an exact on-shell or phase-velocity resonance.

At 128^3 resolution the paired-bootstrap scan sharpens, but does not broaden, the central claim. We treat $\sigma = 0.30$ and $\sigma = 0.45$ as two analysis windows and use the normalized two-time response time $\mathcal{T}_F = \int \mathrm{Re}\, C_F\, d\tau / C_F(0)$ as the ranking metric. This is a per-variance, coherence-only view of the shell-detuning response. The unnormalized Green–Kubo integral $\int \mathrm{Re}\, C_F\, d\tau$ would additionally weight each shell by $C_F(0)$. The central claim of this section is therefore the shape of the shell-detuning response under fixed $C_F(0)$ normalization. Multi-shell comparison of the unnormalized integral across resolutions is left to future work. For the $k = 10$ branch, both windows pass the paired position-bootstrap gate. The narrower $\sigma = 0.30$ window gives $P_{\mathrm{boot}}(\mathrm{on} \geq \mathrm{off}_{\mathrm{best}}) = 0.993$ over 5000 paired resamples. The wider $\sigma = 0.45$ window gives no failures in 5000 paired resamples [$P_{\mathrm{boot}}(\mathrm{on} \geq \mathrm{off}_{\mathrm{best}}) = 1.000$].

The adjacent $k = 11$ branch is bandwidth-sensitive. Its $\sigma = 0.30$ window is suggestive but does not pass the gate ($P_{\rm boot} = 0.685$), while $\sigma = 0.45$ gives no failures in the trajectory bootstrap. We nevertheless retain this case as a secondary check rather than promoting it to a second central clean branch, because the matched T/L/W channel-control file at the same $k = 11, \sigma = 0.45$ setting shows a transverse plateau ordering biased toward the sub-shell branch. Changing the shell center therefore reorganizes the finite-time response and the normalized two-time response time, but the robust central branch remains $k = 10$, especially the audited $\sigma = 0.30$ case whose effective support lies at $q_{\rm eff}/k \simeq 1.2$–1.4 [Fig. 4(a,b)]. Higher-wavenumber branches do not pass the same paired-bootstrap gate at the resolutions tested here. At $k \geq 12$ the off-shell $\alpha = 0.5$ branch overtakes $\alpha = 1$ in the resampled comparison. The robust conclusion is therefore structured scale-selected trajectory memory in the clean support region, not a universal exact on-shell maximum across all coarse-grid branches.

We next ask whether the clean shell-detuning signature is tied to the SVF transverse geometry, or whether any scale-selected vorticity sampled along the same moving path would give the same result. To test this, we rerun the trajectory calculation with identical bath, trajectory, and shell filter, and replace the channel field by two controls. The first is the longitudinal component $W_L = \omega'_x$. The second is the wrong-handed transverse field $W_W = +\omega'_z + \mathrm{i}\omega'_y$. Here $\hat{\boldsymbol{x}}$ is both the carrier-propagation direction and the trajectory direction. ω'_x is therefore parallel to that direction, while ω'_y and ω'_z are transverse components.

The clean $k = 10, \sigma = 0.30$ control gives the central geometry result. The SVF transverse field $W_T = -\omega'_z + \mathrm{i}\omega'_y$ peaks at $\alpha = 1$. The longitudinal channel peaks at $\alpha = 2$ with $\mathcal{T}_F(1)/\mathcal{T}_F(2) = 0.056$. The wrong-handed transverse field gives the same real two-time response time as W_T, as expected because $W_W = -W_T^*$ under the same real shell filter. It is therefore a pipeline invariance check rather than an independent positive channel. A separate $n = 128$, $\sigma = 0.45$, $k = 10, 11$ T/L/W check preserves this channel-geometry separation but also shows bandwidth-sensitive sub-shell ordering in the transverse plateau metric. The decisive geometry rejection is T versus L. Sampling the vorticity component parallel to the trajectory and carrier direction reorganizes the finite-band shell response toward the higher shell instead of generically reproducing the transverse ordering. This separation is not explained by a one-point isotropic vorticity-variance argument. In an isotropic incompressible bath, the transverse and longitudinal components share the same radial vorticity-spectrum scale factor while differing through their angular projection relative to the sampling direction. The observed T/L reversal therefore tests the combined effect of SVF projection geometry, shell filtering, and group-velocity sampling, rather than a generic increase of vorticity variance with wavenumber. This is why the central shell-detuning claim is not expanded beyond the $k = 10$, $\sigma = 0.30$ branch [Fig. 5, see also Supplemental Material [21]].

VI. ENCOUNTER-TIME DECORRELATION

When the probe advances at c_g rather than remaining fixed, the decorrelation time shortens. The cleanest comparison is the matched shell-filtered fixed-point check on the adjacent audited window used above. At $n = 128$, $N_t = 128$, $\sigma = 0.45$, and $\alpha = 1$, the fixed records for $k = 10$ and $k = 11$ remain unresolved through $t_{\rm end} = 0.254$, whereas the corresponding transverse trajectory channel has $\tau_{\rm env} = 0.026$ for both wavenumbers. The fixed-to-moving shortening is therefore bounded below by $0.254/0.026 \approx 9.8$ in the matched shell-filtered channel. Because the T/L/W control at the same bandwidth makes the $k = 11$ transverse ordering sub-shell biased, this comparison is used as an encounter-sampling control rather than as a promotion of $k = 11$ to the central scale-selection branch.

For continuity with the original encounter dataset, we also report the shorter paired trajectory record with $N_t = 64$ snapshots, separate from the $N_t = 128$ shell-detuning dataset used in Fig. 4. In the matched short-record fixed-point data, the envelope remains unresolved. It has not fallen below the threshold by $t = 0.126$ for any of the three tested wavenumbers. The long fixed-point record gives the all-q bath-memory time $\tau_{\rm env}^{E,\,\text{all-}q} = 0.244$ DNS time units, while the corresponding trajectory record at $\alpha = 1$ has $\tau_{\rm env} = 0.028, 0.026, 0.028$ for $k = 10, 15, 30$, respectively [Fig. 6]. Relative to the long-record all-q fixed-point baseline, the shell-filtered moving response is shorter by factors of approximately 8.7, 9.4, and 8.7. Within the matched short-record comparison, the conservative lower bounds are 4.5, 4.85, and 4.5.

These ratios should not be interpreted as a pure encounter-only coefficient. They combine two effects. The first is the scale selection imposed by the shell filter, which by itself shortens fixed-point bath memory because shell-restricted Eulerian times scale as $\tau_E(q) \sim 1/(q\, u_{\rm rms})$ for $q \sim k$. The second is the effect of moving the probe at c_g, which adds a sweep contribution $\sim c_g/\ell_{\rm eddy}$ to the decorrelation rate. The matched shell-filtered fixed-point baseline isolates this issue more cleanly than the all-q comparison, but broader bandwidth and wavenumber coverage are still needed to separate (i) and (ii) quantitatively.

What the present comparison does establish is that the shell-filtered moving response is bounded above by an encounter-cycle count

$$N_{\rm cyc}^{\rm enc}(k, \alpha) \equiv \frac{\tau_{\rm enc}(k, \alpha)}{T_{\rm wave}(k)}, \tag{10}$$

which for the $\alpha = 1$ trajectory records in Fig. 6 (using $\tau_{\rm enc} \equiv \tau_{\rm env}$) is $N_{\rm cyc}^{\rm enc} = 2.76, 3.14, 4.77$ for $k = 10, 15, 30$, in contrast to the many-cycle all-q fixed-point bath memory in Eq. (9). The relevant trajectory memory is there-

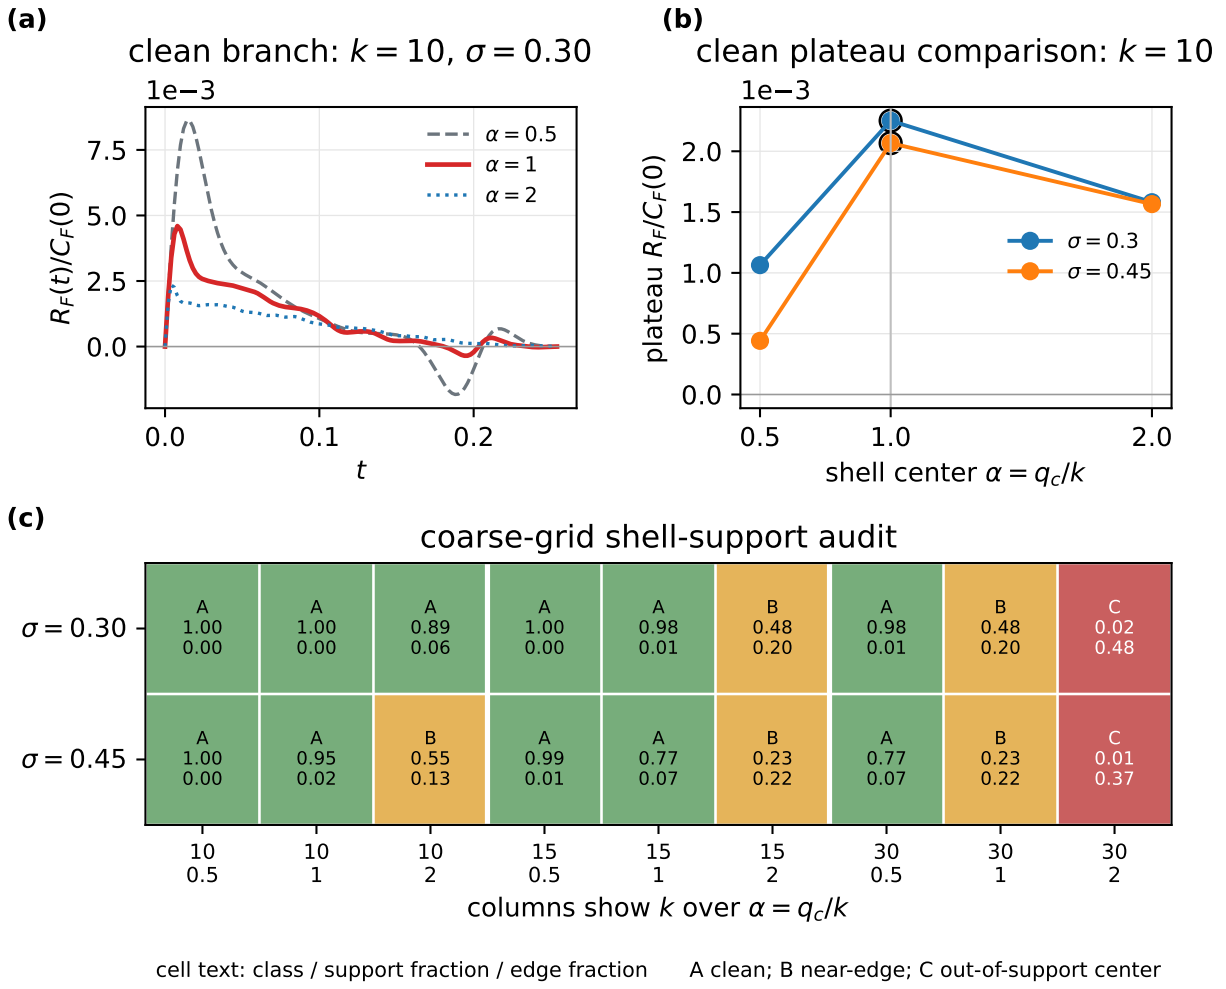

Figure 4. **Scale-selected trajectory response.** (a) Normalized running response $R_F(t)/C_F(0)$ for the clean $k = 10$, $\sigma = 0.30$ branch. All three shell centers in this panel are class-A filters in the support audit. (b) Normalized two-time response time $R_F/C_F(0)$ versus shell center for $k = 10$. Open circles mark the largest branch at each width. The nominal $q_c = k$ filter is largest for both widths $\sigma = 0.30$ and $\sigma = 0.45$, both of which pass the paired-bootstrap gate. Effective support of the nominal filter lies at $q_{\mathrm{eff}}/k \simeq 1.2$–$1.4$. (c) Coarse-grid support audit for the full shell-detuning matrix. Each cell gives class, resolved support fraction, and edge fraction with $\max_i |q_i| \geq 30$. Class A is used for central claims, class B is near-edge by the $\eta_{\mathrm{supp}} < 0.75$ or $\eta_{\mathrm{edge}} \geq 0.25$ criterion, and class C has an out-of-support shell center.

fore limited by both bath persistence and the time over which the moving wave remains in the same scale-selected eddy, with the relative weight of scale selection and trajectory sampling to be quantified by the matched-control extension. Dimensionally, if the selected eddy length is $\ell_{\mathrm{eddy}} \sim 2\pi/q_c$, then $\tau_{\mathrm{enc}} \sim \ell_{\mathrm{eddy}}/c_g$ with the deep-water group velocity $c_g = \omega/(2k)$ gives

$$N_{\mathrm{cyc}}^{\mathrm{enc}} \sim \frac{\ell_{\mathrm{eddy}}/c_g}{2\pi/\omega} = \frac{2k}{q_c} = \frac{2}{\alpha}, \tag{11}$$

up to the finite-band, eddy-lifetime, and envelope-threshold corrections measured here. The measured $\alpha = 1$ values are therefore order-one longer than this simplest crossing estimate, rather than parametrically inconsistent with it. This is consistent with the classical distinction between Eulerian, Lagrangian, and sweeping decorrelation times in turbulence [18, 20, 28].

VII. SCOPE OF THE PRESENT CLAIM

Throughout this work, "response" denotes the Green–Kubo/linear-response force autocorrelation associated with the prescribed SVF probe. It does not mean that the DNS bath is dynamically forced by a resolved wave or that wave back-reaction is included.

The claim is deliberately limited in four ways. First, the DNS is performed at a single $\mathrm{Re}_\lambda \approx 433$. The fixed-point estimator selectivity, trajectory-side channel specificity, and memory dependence established here are qualitative properties of the prescribed-probe response

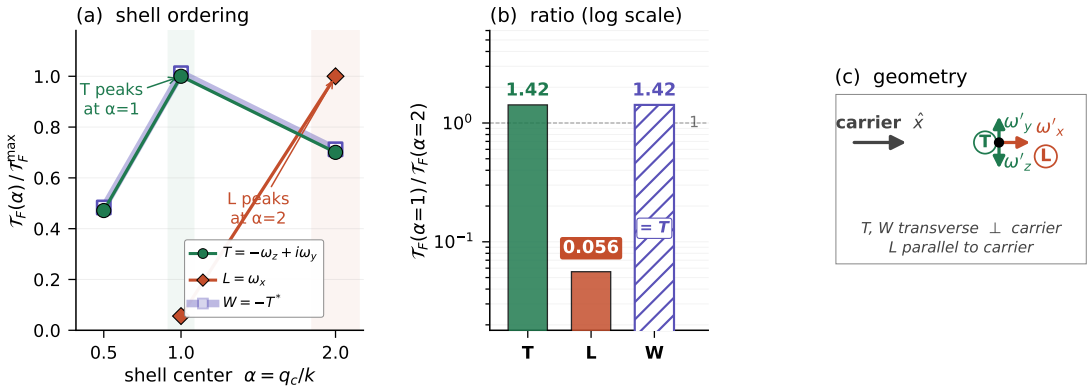

Figure 5. **Trajectory-side channel-geometry control.** The clean trajectory branch is recomputed with the same bath, trajectory, shell filter, $k = 10$, and $\sigma = 0.30$, changing only the channel field. The SVF transverse channel $T = -\omega'_z + \mathrm{i}\omega'_y$ peaks at the nominal $q_c = k$ filter, while the longitudinal channel $L = \omega'_x$ peaks at $q_c = 2k$. The wrong-handed transverse check $W = +\omega'_z + \mathrm{i}\omega'_y$ reproduces the same real response-time ordering as T, as expected for the same real shell filter. Panel (a) shows the channel-normalized shell ordering. Panel (b) shows ordering ratios, including the decisive longitudinal value $\mathcal{T}_F(1)/\mathcal{T}_F(2) = 0.056$.

protocol. Quantitative Re-dependence of the response amplitude, and the asymptotic scale-selective closure of the kernel [18, 19], are deferred to subsequent multi-Re work. Second, the central shell-detuning evidence at the present grid resolution is the clean-support configuration $(k, \sigma) = (10, 0.30)$. Within this configuration, transverse-SVF specificity is established by the longitudinal-channel control. The wrong-handed transverse check serves only as an invariance test of the pipeline.

A 128^3 paired-bootstrap survey at $\sigma \in \{0.30, 0.45\}$ across $k \in \{10, \ldots, 15\}$ resolves the wavelength-matched maximum as a finite-k window phenomenon. The $\alpha = 1$ peak passes the paired-bootstrap gate robustly at $k = 10$ for both windows. The adjacent $k = 11$, $\sigma = 0.45$ case passes the trajectory-bootstrap gate but is kept as a secondary bandwidth-sensitive check, because its matched T/L/W control is sub-shell-biased in the transverse plateau ordering. For $k \geq 12$ the gate fails at the present resolution. The off-shell $\alpha = 0.5$ branch overtakes $\alpha = 1$ in the resampled comparison, consistent with reduced separation between neighboring shells when k approaches the dissipation-affected end of the resolved range.

At the lowest tested wavenumbers ($k = 6, 8$) the peak shifts to $\alpha = 2$, consistent with a crossover to large-eddy infrared sweep dominance when the wavelength becomes comparable to the integral scale of the domain. T/L/W controls at $\sigma = 0.45$ replicate the longitudinal flip at both $k = 10$ and $k = 11$. They also expose the bandwidth sensitivity that prevents promoting the adjacent $k = 11$ case to a second central branch. Shell-filtered fixed-point matched controls are reported in the Supplemental Material [21].

Third, the maximum in the clean branch is a finite-band eddy-scale statement. The nominal $q_c = k$ filter has effective support $q_{\mathrm{eff}}/k \simeq 1.2$–$1.4$, not an exact resonance condition.

Fourth, the present DNS measurement does not impose inertial-range closure of the bath-side response. It does not, therefore, test the macroscopic $\omega^{8/3}$ swell-attenuation law derived in the related Green–Kubo analysis [6]. It provides the bath-side autocorrelation that enters that closure.

VIII. DISCUSSION

The DNS measurements yield five channel-resolved findings. First, the fixed-point branch extracts a finite, integrable bath-side autocorrelation for the modal SVF projection. The running integral $R_F(t) = \int_0^t \mathrm{Re}\, C_F(\tau)\, \mathrm{d}\tau$ reaches a plateau and remains within the position-block bootstrap band out to $t \gtrsim 5\tau_{\mathrm{env}}$. Second, the low-frequency baseline, time shuffling, and frequency detuning controls establish fixed-point estimator selectivity. They rule out one-point marginal statistics and show that the lock-in estimate requires temporal ordering and matched demodulation, with suppression factors spanning an order of magnitude (time shuffling) to four orders of magnitude (alias-free frequency detuning). Third, the scale-selected trajectory observable resolves a structured normalized response-time ordering with a maximum at the nominal wavelength-matched filter $\alpha = q_c/k = 1$ in the audited clean-support branch, while adjacent branches expose bandwidth-sensitive sub-shell competition. Fourth, a trajectory-side channel-geometry control isolates this ordering to the transverse SVF projection. Replacing the channel by the longitudinal vorticity ω'_x shifts the maximum to $\alpha = 2$, with $\mathcal{T}_F(1)/\mathcal{T}_F(2) = 0.056$. This is an approximately eighteen-fold preference for the $\alpha = 2$ shell over the $\alpha = 1$ shell within the longitudinal control, and it falsifies a generic-vorticity explanation of

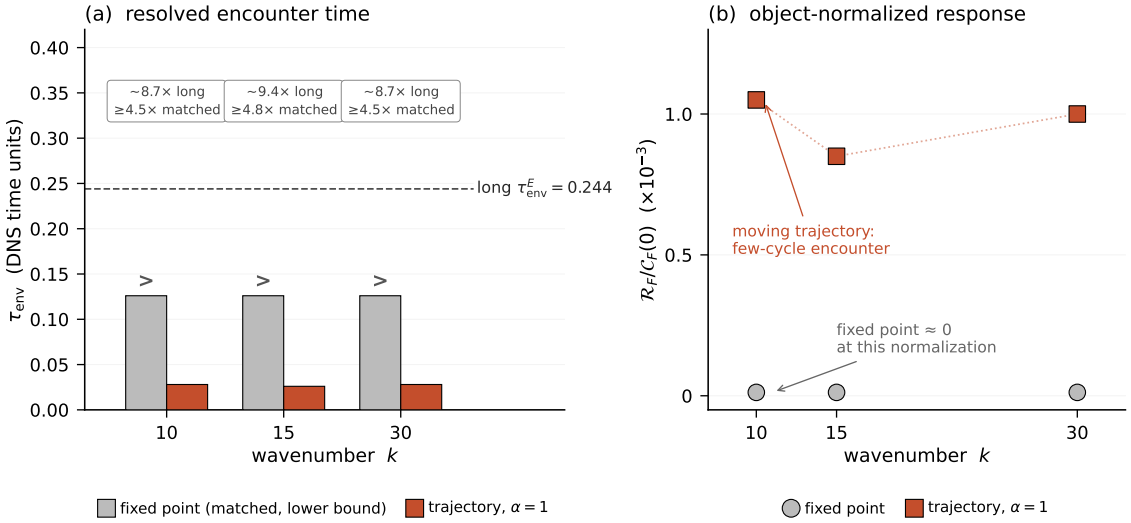

Figure 6. **Encounter-time decorrelation.** (a) Operational envelope decorrelation time for the $N_t = 64$ encounter-time dataset, separate from the $N_t = 128$ shell-detuning dataset. Gray bars are matched short-record lower bounds. The fixed envelope does not fall below threshold by $t = 0.126$ in this $N_t = 64$ window, so the displayed bar height is a conservative lower bound rather than the resolved Eulerian time. The long all-q fixed-point baseline (not shown as a bar in this panel) gives $\tau^E_{\rm env} = 0.244$ DNS time units. Relative to that long baseline, the $\alpha = 1$ trajectory records shorten the memory time by factors of approximately 8.7, 9.4, and 8.7. Within the matched short records the conservative lower bound is at least 4.5. This all-q comparison is a consistency check, not a pure encounter-only coefficient. The matched shell-filtered $n = 128$ fixed-point control reported in the text and in the Supplemental Material gives the cleaner fixed-to-moving lower bound > 9.8 in the transverse $\alpha = 1$ channel. The higher-k branches include the near-edge caution identified in Fig. 4(c), so this panel is used as evidence for encounter-time resolution rather than as an additional clean-support shell-detuning claim. (b) The corresponding normalized two-time response time. The moving trajectory exposes an encounter-time ingredient absent from fixed-point sampling. It is not, by itself, an independent test of the channel or of the scale selection. Those are established by Figs. 3–5. Absolute $\tau_{\rm env}$ values throughout this figure use the operational envelope threshold $\theta = 0.05$ (see Supplemental Material [21]); they are convention-dependent. The fixed-to-moving *ratios* reported here are the comparison-of-record and are insensitive to this convention.

the transverse trajectory peak. A wrong-handed transverse check $W = +\omega'_z + \mathrm{i}\omega'_y$ reproduces the real response time of the SVF projection because W and the SVF field W_T satisfy $W = -W_T^*$ after applying the same real shell filter and trajectory sampling. This verifies pipeline self-consistency rather than providing an independent positive test. The location of the normalized transverse-channel maximum is compatible with the $q \sim k$ peak structure predicted for the Green–Kubo integrand in [6]. The present DNS comparison, however, is a per-variance, normalized shell-detuning diagnostic rather than a direct measurement of the unnormalized multi-shell Green–Kubo integrand. The underlying mechanism is summarized below. Fifth, the shell-filtered moving response decorrelates more rapidly than the all-q fixed-point bath-memory time, by a factor of approximately 9 in the long-record comparison or at least 4.5 in the matched short-record lower bound. The shell-filtered fixed-point check at $n = 128$, $\sigma = 0.45$, and $k = 10, 11$ gives a matched lower bound > 9.8 for fixed-to-moving shortening in the transverse $\alpha = 1$ channel. This contrast still combines scale-selection and trajectory-sampling effects. Together, these results establish that the SVF projection carries a finite, estimator-controlled, geometrically specific, scale-selected, encounter-organized two-time autocorrelation in Navier–Stokes turbulence. The bath-side force autocorrelation is a measurable object before any closure is imposed.

We therefore center the scale-selection claim on the single fully clean-support branch identified by the support audit, and use the remaining branches as boundary information rather than as equal-weight evidence. This is a deliberate methodological choice. The central branch carries spectral support audited as class-A at all three shell centers. The higher-wavenumber coarse-grid branches contain near-edge or out-of-support filters whose effective shell centers deviate from the nominal label. The correct take-home is therefore not an exact on-shell resonance and not "the nominal $q_c = k$ branch is always largest", but the disciplined statement that, on the audited clean branch, the SVF channel is a scale-selected, trajectory-controlled response with maximum at the nominal wavelength-matched filter. Confirming that this same fully audited ordering persists across multiple branches and at higher Re_λ is left to subsequent work.

The peak location is consistent with classical turbulence phenomenology under a sweep-dominated interpre-

tation [6, 18, 20, 28]. The trajectory force autocorrelation integrates eddies of scale q with a sweep correlation time $\tau_{\rm eff}(q;k) \sim 1/(q\,c_g)$. For the unnormalized force integral, the time-integrated contribution per shell would be the product of this eddy-encounter time, the vortex-force vertex $|V|^2 \propto q^2$, and the bath energy at the same scale. The q^2 vertex penalises large-scale eddies ($q \ll k$, where vorticity is small), the geometric wave–eddy overlap penalises sub-wavelength eddies ($q \gg k$, where the wave averages over many uncorrelated patches), and the residual peak lies at $q \sim k$. The normalized response time $\mathcal{T}_F$ in Fig. 4(b) is the per-variance ("coherence-only") view of this response. It identifies where the shell-restricted channel remains coherent longest after normalization by $C_F(0)$, which is the quantity isolated by Fig. 5. The shortened encounter-time decorrelation in Fig. 6 is consistent with the Eulerian–Lagrangian–sweeping distinction in isotropic turbulence, the wave moving at c_g through frozen-by-comparison eddies. The finite-band $\alpha \sim 1$ maximum is also consistent with rapid-distortion analyses of wave–turbulence interaction [22, 23] and with the long-recognized role of wave-organized turbulence in the upper ocean [9, 24]. Unlike those earlier studies, however, the present measurement extracts a two-time response integral for a specific bilinear coupling channel and tests its specificity by direct comparison against alternative channel projections at the same trajectory and shell filter. The underlying mechanism is classical, not a plasma-style wave–particle resonance. The wave propagates at its own group velocity through a passive turbulent bath, and the shell window reflects the geometric match between wave wavelength and eddy scale, not a phase-velocity coincidence.

The bath-side response measured here is a necessary input for, but not by itself a derivation of, the macroscopic transport coefficient. A Green–Kubo description of remote swell attenuation is developed in [6], where the time integral of the same autocorrelation is shown to yield a non-negative damping coefficient by Bochner's theorem and, under inertial-range closure, a parameter-free attenuation law tested against 241 trans-oceanic satellite tracks. The two works divide labor cleanly. The present work establishes that this bath-side autocorrelation exists, is finite, is integrable, and, in the audited clean branch, is selectively carried by a finite band centered near the wavelength-matched eddy scale, without invoking any closure. That separate analysis converts the same autocorrelation into a closed macroscopic transport prediction. This separation matters because the existence and structure of the bath-side response are DNS questions about Navier–Stokes turbulence, while the macroscopic closure brings in inertial-range physics and oceanographic context that are not tested by either work in isolation.

Beyond the SVF, the prescribed-probe protocol developed here offers a general route. Many reduced models remove named bilinear couplings by phase averaging, ensemble averaging, or rapid-variable elimination [2–5], and the relevant question in any concrete system is whether a particular eliminated channel still carries measurable two-time response content. An externally prescribed reference structure, combined with matched controls and a geometry-rejecting check, provides an operational answer that does not rely on inferring channel content from undifferentiated bath statistics. The wave–vorticity coupling studied here is one realization. The construction transfers directly to other phase-suppressed bilinear couplings, in fluid mechanics and beyond.

ACKNOWLEDGMENTS

This work was supported by Khalifa University of Science and Technology and made use of the Johns Hopkins Turbulence Database (JHTDB). The authors thank the JHTDB team for data access and computational support.

[1] A. D. D. Craik and S. Leibovich, A Rational Model for Langmuir Circulations, Journal of Fluid Mechanics **73**, 401 (1976).
[2] D. G. Andrews and M. E. McIntyre, An Exact Theory of Nonlinear Waves on a Lagrangian-Mean Flow, Journal of Fluid Mechanics **89**, 609 (1978).
[3] O. Bühler, *Waves and Mean Flows*, 2nd ed., Cambridge Monographs on Mechanics (Cambridge University Press, Cambridge, UK, 2014).
[4] H. Mori, Transport, Collective Motion, and Brownian Motion, Progress of Theoretical Physics **33**, 423 (1965).
[5] R. Zwanzig, *Nonequilibrium Statistical Mechanics* (Oxford University Press, New York, 2001).
[6] G. Liu and M. Al Alshehhi, A universal law for non-breaking surface wave decay, EarthArXiv preprint (2026).
[7] J. C. McWilliams, J. M. Restrepo, and E. M. Lane, An Asymptotic Theory for the Interaction of Waves and Currents in Coastal Waters, Journal of Fluid Mechanics **511**, 135 (2004).
[8] P. P. Sullivan and J. C. McWilliams, Dynamics of Winds and Currents Coupled to Surface Waves, Annual Review of Fluid Mechanics **42**, 19 (2010).
[9] J. C. McWilliams, P. P. Sullivan, and C.-H. Moeng, Langmuir Turbulence in the Ocean, Journal of Fluid Mechanics **334**, 1 (1997).
[10] F. Ardhuin and A. D. Jenkins, On the Interaction of Surface Waves and Upper Ocean Turbulence, Journal of Physical Oceanography **36**, 551 (2006).
[11] N. Rascle, F. Ardhuin, P. Queffeulou, and D. Croizé-Fillon, A Global Wave Parameter Database for Geophysical Applications. Part 1: Wave-Current-Turbulence Interaction Parameters for the Open Ocean based on Traditional Parameterizations, Ocean Modelling **25**, 154 (2008).

[12] A. V. Babanin, On a Wave-Induced Turbulence and a Wave-Mixed Upper Ocean Layer, Geophysical Research Letters **33**, L20605 (2006).

[13] A. V. Babanin and B. K. Haus, On the Existence of Water Turbulence Induced by Nonbreaking Surface Waves, Journal of Physical Oceanography **39**, 2675 (2009).

[14] Y. Li, E. Perlman, M. Wan, Y. Yang, C. Meneveau, R. Burns, S. Chen, A. Szalay, and G. Eyink, A Public Turbulence Database Cluster and Applications to Study Lagrangian Evolution of Velocity Increments in Turbulence, Journal of Turbulence **9**, N31 (2008).

[15] E. Perlman, R. Burns, Y. Li, and C. Meneveau, Data Exploration of Turbulence Simulations Using a Database Cluster, in *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing (SC '07)* (ACM/IEEE, Reno, NV, USA, 2007).

[16] H. Yu, K. Kanov, E. Perlman, J. Graham, E. Frederix, R. Burns, A. Szalay, G. Eyink, and C. Meneveau, Studying Lagrangian Dynamics of Turbulence using On-Demand Fluid Particle Tracking in a Public Turbulence Database, Journal of Turbulence **13**, N12 (2012).

[17] J. Graham, K. Kanov, X. I. A. Yang, M. Lee, N. Malaya, C. C. Lalescu, R. Burns, G. Eyink, A. Szalay, R. D. Moser, and C. Meneveau, A Web Services Accessible Database of Turbulent Channel Flow and its Use for Testing a New Integral Wall Model for LES, Journal of Turbulence **17**, 181 (2016).

[18] R. H. Kraichnan, Kolmogorov's Hypotheses and Eulerian Turbulence Theory, Physics of Fluids **7**, 1723 (1964).

[19] U. Frisch, *Turbulence: The Legacy of A. N. Kolmogorov* (Cambridge University Press, Cambridge, UK, 1995).

[20] H. Tennekes, Eulerian and Lagrangian Time Microscales in Isotropic Turbulence, Journal of Fluid Mechanics **67**, 561 (1975).

[21] Supplemental Material for "A prescribed wave probe resolves a bath-side stochastic vortex force response in turbulence", Included with this manuscript (2026).

[22] M. A. C. Teixeira and S. E. Belcher, On the Distortion of Turbulence by a Progressive Surface Wave, Journal of Fluid Mechanics **458**, 229 (2002).

[23] L. Thais and J. Magnaudet, Turbulent Structure beneath Surface Gravity Waves Sheared by the Wind, Journal of Fluid Mechanics **328**, 313 (1996).

[24] F. Veron and W. K. Melville, Experiments on the Stability and Transition of Wind-Driven Water Surfaces, Journal of Fluid Mechanics **446**, 25 (2001).

[25] M. S. Green, Markoff Random Processes and the Statistical Mechanics of Time-Dependent Phenomena. II. Irreversible Processes in Fluids, Journal of Chemical Physics **22**, 398 (1954).

[26] R. Kubo, Statistical-Mechanical Theory of Irreversible Processes. I. General Theory and Simple Applications to Magnetic and Conduction Problems, Journal of the Physical Society of Japan **12**, 570 (1957).

[27] R. Kubo, The Fluctuation-Dissipation Theorem, Reports on Progress in Physics **29**, 255 (1966).

[28] S. Chen and R. H. Kraichnan, Sweeping Decorrelation in Isotropic Turbulence, Physics of Fluids A **1**, 2019 (1989).

Supplemental Material:A prescribed wave probe reveals phase hidden vortex force memory in turbulence

Guoqiang Liu[1,*] and Maryam Al Alshehhi[1]

[1] *Department of Civil, Infrastructural and Environmental Engineering, Khalifa University of Science and Technology, Abu Dhabi, UAE*

(Dated: May 10, 2026)

* guoqiang.liu@ku.ac.ae

I. DNS DATASET AND DATA PIPELINE

All measurements use the `isotropic1024coarse` dataset on the Johns Hopkins Turbulence Database (JHTDB) [1–4]. The flow is statistically stationary homogeneous isotropic turbulence forced at large scales; reported parameters are $\mathrm{Re}_\lambda \approx 433$, dissipation rate $\varepsilon = 0.093$, kinematic viscosity $\nu = 1.85 \times 10^{-4}$, integral length $L \approx 1.4$, and integral time $T_L \approx 2.0$ (all in DNS units). The simulation domain is $(2\pi)^3$ with periodic boundaries. At each sampling location we retrieve the three-component vorticity $\boldsymbol{\omega}'(\boldsymbol{x}, t)$ at cadence $\Delta t = 0.002$ for $N_t = 5028$ frames ($T \approx 10$ in DNS time units), using cubic spline interpolation in space and PCHIP in time.

The vorticity field is obtained from the JHTDB `getData` vorticity service, which evaluates the curl of the stored velocity field on the native simulation grid using fourth-order centered finite differences before spatial interpolation [1, 3]. The documented root-mean-square uncertainty of the service relative to a direct spectral curl on the same grid is approximately 7% of the local vorticity magnitude. The same provenance is used for the fixed-point point time series and for the local FFT cubes underlying the trajectory shell-filter measurements; we have not retrieved or stored full velocity volumes. Spectrum-weighted audit columns reported below therefore use the public isotropic spectrum in `spatial_k56_data.npz` as a radial proxy and are not treated as per-run measurements.

The dominant trajectory-side metric reported in this work is the normalized two-time response time $\mathcal{T}_F = R_F/C_F(0)$. Because $C_F(\tau)$ is bilinear in $\boldsymbol{\omega}'$, an approximately multiplicative 7%-rms uncertainty on the input vorticity magnitude appears as a near-common factor in $R_F(t)$ and in $C_F(0)$ and largely cancels in the ratio. The first-order residual after this cancellation is bounded by the normalized component-wise correlation of the gradient-routine error across τ and is substantially smaller than the threshold and bandwidth sensitivities already reported below. We therefore do not assign a separate $\mathcal{T}_F$ error budget to this source. Block bootstrap and threshold-sensitivity scans (Sec. III, Sec. V) dominate the reported uncertainty.

II. PRESCRIBED WAVE PROBE AND CHANNEL PROJECTION

The SVF channel defined by the cross product $\boldsymbol{u}_\phi \times \boldsymbol{\omega}'$ is the natural target of the probe. It is the bilinear coupling that arises from the Helmholtz decomposition of the Navier–Stokes equations and is removed by phase averaging in classical wave-averaged descriptions of wave–current interaction [5–8]; its relation to Generalized-Lagrangian-Mean averaging is classical [9]. The two-time autocorrelation of this coupling, projected onto the wave mode at k, is the bath-side input to the Green–Kubo description developed in [10]. The present supplement documents the DNS construction by which this mode-projected force is realized in turbulence.

The prescribed carrier is analytic, implemented as a cosine–sine quadrature pair,

$$\boldsymbol{u}_\phi^{(c)}(\boldsymbol{x},t) = U_0 \cos\phi(\boldsymbol{x},t)\,\hat{\boldsymbol{x}}, \qquad \boldsymbol{u}_\phi^{(s)}(\boldsymbol{x},t) = U_0 \sin\phi(\boldsymbol{x},t)\,\hat{\boldsymbol{x}}, \qquad \phi(\boldsymbol{x},t) = \boldsymbol{k}\cdot\boldsymbol{x} - \omega t, \tag{S1}$$

with $\boldsymbol{k} = k\hat{\boldsymbol{x}}$, deep-water dispersion $\omega = \sqrt{g_{\mathrm{ref}}k}$, and group velocity $c_g = \omega/(2k)$. The reference gravity g_{ref} is a control parameter that sets the carrier frequency at fixed k. The channel-resolved force is

$$F_k(t) = \frac{1}{U_0}\, e^{-\mathrm{i}\phi(t)} \big([\boldsymbol{u}_\phi^{(c)}(\boldsymbol{x}(t),t) \times \boldsymbol{\omega}'(\boldsymbol{x}(t),t)] + \mathrm{i}[\boldsymbol{u}_\phi^{(s)}(\boldsymbol{x}(t),t) \times \boldsymbol{\omega}'(\boldsymbol{x}(t),t)] \big) \cdot \big(\hat{\boldsymbol{y}} + \mathrm{i}\hat{\boldsymbol{z}}\big), \tag{S2}$$

which after expansion becomes

$$F_k(t) = W(t), \qquad W(t) \equiv -\omega_z'(\boldsymbol{x}(t),t) + \mathrm{i}\omega_y'(\boldsymbol{x}(t),t). \tag{S3}$$

Equivalently, using the complex notation $\boldsymbol{u}_\phi^{\mathrm{an}} = U_0 e^{\mathrm{i}\phi}\hat{\boldsymbol{x}}$ gives $(e^{\mathrm{i}\phi}e^{-\mathrm{i}\phi})W = W$. This analytic lock-in removes the mirror-frequency term $\frac{1}{2}e^{-2\mathrm{i}\phi}W$ that would be produced by a single real-cosine carrier.

Two sampling configurations are used. In *fixed-point* sampling, $\boldsymbol{x}(t) = \boldsymbol{x}_0$ is constant, with $N_p = 200$ positions drawn uniformly over the $(2\pi)^3$ domain. In *trajectory* sampling, $\boldsymbol{x}(t) = \boldsymbol{x}_0 + c_g\, t\, \hat{\boldsymbol{x}}$, with $N_{\mathrm{traj}} = 200$ trajectories distributed as 50 starting positions per δ value, at $\delta = u'(k)/c_g \in \{0.01, 0.03, 0.05, 0.10\}$. The dimensionless parameter δ is a Kubo-like coupling. In a generic fast-carrier setting, $\delta \ll 1$ would correspond to the kinematic limit where the probe traverses many independent eddies per encounter. In the present fixed-k deep-water sweep, however, the identity $c_g T_{\mathrm{wave}} = \pi/k$ keeps the wave–eddy spatial geometry fixed to leading order.

III. AUTOCORRELATION, ENVELOPE, AND RESPONSE INTEGRAL

For each probe realization we compute the biased autocorrelation via the zero-padded Wiener–Khinchin estimator

$$\widehat{C}_F(\tau) = \mathcal{F}^{-1}\big[\big|\mathcal{F}[F_k - \langle F_k\rangle_t]\big|^2\big], \tag{S4}$$

with explicit subtraction of the sample time-mean $\langle F_k\rangle_t$ prior to the FFT. This detrending removes a finite-record bias $|\langle F_k\rangle_T|^2$ that would otherwise add a constant offset to $C_F(\tau)$ at all lags, producing a spurious linear drift of the running integral past $5\tau_{\rm env}$. The effect is $\sim 9\%$ of $|F_k|$ in fixed-point records (where $\langle F_k\rangle_t$ averages over ~ 200 stationary positions in a domain of integral length $L \approx 1.4$) and $\sim$ 3–6% in trajectory records (where spatial sampling provides additional decorrelation). Uncertainty on the running-integral plateau is reported throughout via position-block bootstrap rather than by treating the positions as independent realizations. The resampling unit is the position (fixed-point) or the trajectory starting point (trajectory), treated as the block. Each bootstrap replicate draws N_p positions/trajectories with replacement from the sampled set (fixed-point $N_p = 200$; clean trajectory $N_p = 64$; high-resolution paired-bootstrap matrix $N_p = 96$). The full record at each drawn position is preserved, so within-position temporal correlation is retained. Resampling uses a fixed seed (documented in the data-package metadata accompanying Table S7) so that reported $P_{\rm boot}$ values and the resulting plateau bands are exactly reproducible. Omission of the detrending step alters the plateau value by $\sim 15\%$ and is the main cause of apparent non-convergence past $5\tau_{\rm env}$.

The operational decorrelation time is defined from the Hilbert envelope of the normalized real correlation, after boxcar smoothing with window $w = 5$ samples and reflective boundary handling:

$$\tau_{\rm env} \equiv \min\{\tau : \mathrm{env}(\tau) < \theta\}, \qquad \tau_{\rm int} \equiv \int_0^{\tau_{\rm env}} \mathrm{env}(\tau)\,\mathrm{d}\tau, \tag{S5}$$

with threshold $\theta = 0.05$. Robustness to $\theta \in \{0.02, 0.05, 0.10, 0.20\}$ changes $\tau_{\rm env}$ by a factor of ~ 4. The threshold is therefore used as a consistent operational convention rather than as an independently fitted physical exponent. Two normalized response-integral estimators are used. The envelope-window estimator, used for control discriminations where signals are small and no plateau is available, is

$$\mathcal{T}_{F,\rm env} = \frac{1}{C_F(0)} \int_0^{\tau_{\rm env}} \mathrm{Re}\, C_F(\tau)\,\mathrm{d}\tau. \tag{S6}$$

The plateau response time is

$$\mathcal{T}_F^{\mathrm{Eul}} = \mathrm{median}_{t\in[\tau_{\mathrm{env}},\,3\tau_{\mathrm{env}}]}\mathcal{T}_F(t), \qquad \mathcal{T}_F(t) = \frac{1}{C_F(0)}\int_0^t \mathrm{Re}\, C_F(\tau)\,\mathrm{d}\tau. \tag{S7}$$

For the analytic-carrier fixed-point records, $\mathcal{T}_{F,\mathrm{env}} = 0.02997$ and $\mathcal{T}_F^{\mathrm{Eul}} = 0.03139$ in DNS time units. This plateau response time $\mathcal{T}_F^{\mathrm{Eul}}$ is the time integral of the demodulated force autocorrelation in the fixed-point configuration. This is the bath-side autocorrelation that enters the Green–Kubo coefficient developed in [10].

IV. TWO COMPLEMENTARY MEASUREMENTS OF ONE CHANNEL

The fixed-point and trajectory observables defined below address two physically distinct questions about the same SVF channel. Equations (S8)–(S12) make this distinction precise. The closed macroscopic kernel $\Gamma_{\mathrm{th}}(k)$ obtained in [10] is a related third object reached by an additional closure step that the present DNS measurement does not attempt. We record its definition in Eq. (S13) for cross-reference only.

A. Fixed-point observable

The fixed-point observable is the modal projection of the SVF onto the wave mode at k, evaluated at stationary positions:

$$F_k(t) = W(t), \qquad W(t) \equiv -\omega_z'(\boldsymbol{x}(t),t) + \mathrm{i}\omega_y'(\boldsymbol{x}(t),t), \tag{S8}$$

with $\boldsymbol{x}(t) = \boldsymbol{x}_0$ stationary. The corresponding normalized response integral is

$$\mathcal{T}_F(t) = \frac{1}{C_F(0)}\int_0^t \mathrm{Re}\, C_F(\tau)\,\mathrm{d}\tau, \qquad \mathcal{T}_F^{\mathrm{Eul}} = \mathrm{median}_{t\in[\tau_{\mathrm{env}},\,3\tau_{\mathrm{env}}]}\mathcal{T}_F(t). \tag{S9}$$

After demodulation the modal factor $e^{\mathrm{i}\phi(t)}$ is cancelled pointwise by $e^{-\mathrm{i}\phi(t)}$, leaving the transverse-vorticity combination $F_k = -\omega_z' + \mathrm{i}\omega_y'$ in channel coordinates. This algebraic form is the standard mode-projection structure of any Mori–Zwanzig construction [11–14]: the force autocorrelation entering a Green–Kubo coefficient is a bath correlation weighted by a fixed projection geometry, and the projection here is the modal vertex of $\boldsymbol{u}_\phi \times \boldsymbol{\omega}'$. The fixed-point measurement is therefore the direct DNS realization of this autocorrelation; the controls in Sec. VI below show that the resulting response is genuinely carrier-channel-specific

rather than a generic one-point property of turbulence. Because the analytic fixed-point observable reduces to $-\omega_z' + \mathrm{i}\omega_y'$, it carries no q weighting: it is a bath-side autocorrelation of the projected force, not a scale-selective object.

B. Trajectory observable

The trajectory observable answers a different question. It asks how the same channel is reorganized when the probe advances at the wave group velocity. It is constructed from the shell-filtered transverse vorticity field

$$W_{k,\alpha}^{\mathrm{sel}}(\boldsymbol{x}, t) = \mathcal{F}^{-1}\left[A_\alpha(|\boldsymbol{q}|/k)\,\widehat{W}(\boldsymbol{q}, t)\right], \tag{S10}$$

sampled along the wave trajectory,

$$F_{k,\alpha}^{\mathrm{sel}}(t) = W_{k,\alpha}^{\mathrm{sel}}(\boldsymbol{x}_0 + c_g(k)t\hat{\boldsymbol{x}}, t). \tag{S11}$$

This second observable samples the bath space–time correlation along $\boldsymbol{r} = c_g\tau\,\hat{\boldsymbol{x}}$ rather than at fixed position, and the log-Gaussian shell filter A_α localizes the measurement in scale. Its correlation therefore probes

$$C_{k,\alpha}^{\mathrm{sel}}(\tau) \simeq C_{\omega,k,\alpha}^{E,\mathrm{sel}}(\boldsymbol{r} = c_g\tau\,\hat{\boldsymbol{x}}, \tau), \tag{S12}$$

exposing wave–eddy geometry (which scale of bath fluctuations dominates the response along the trajectory) and encounter timing (how long the moving wave remains in the same scale-selected eddy), neither of which is available from the fixed-point bath autocorrelation. The fixed-point branch characterizes the bath-side response of the SVF channel. The trajectory branch resolves its scale selectivity and encounter timing.

C. Macroscopic kernel (cross-reference)

The closed macroscopic memory kernel derived in [10] is

$$\Gamma_{\mathrm{th}}(k) \sim \int \mathrm{d}^3q\, S_u(q, \Omega)\, \mathcal{W}_\Omega(\Omega - \boldsymbol{q}\cdot\boldsymbol{c}_g)\, \mathcal{S}(q/k), \tag{S13}$$

a continuous q-space object weighted by both the scale-selection window $\mathcal{S}(q/k)$ and the sweep-decorrelation kernel $\mathcal{W}_\Omega$, whose time-domain effective correlation time $\tau_{\mathrm{eff}}(q; k) \sim$

$1/(q\,c_g)$ enters the sweep-dominated interpretation given in the main text. It is obtained from the force autocorrelation in Eq. (S8) by additional steps (Helmholtz decomposition, three-layer vorticity decomposition, scale-locality assumption, Kolmogorov inertial-range closure) that are derived and tested in [10]. We do not invoke or test those steps here. Eq. (S13) is recorded only to make the cross-reference between the two works explicit.

The nominal shell center $q_c = \alpha k$ used in the trajectory observable is a scale label, not an exact resonance condition or an exact effective shell center. The resulting trajectory result is therefore interpreted as structured nominal shell detuning under spectral-support control (Sec. V), not as a claim of exact on-shell resonance.

V. SHELL-SUPPORT AUDIT AND THRESHOLD SENSITIVITY

The scale-selected trajectory observable uses a log-Gaussian shell window with nominal center $q_c = \alpha k$. Because the calculation is performed on a coarse 64^3 periodic grid, the nominal shell center is not identical to the effective support of the window. We therefore audit each shell using unweighted window diagnostics on the resolved grid. The resolved support fraction η_{supp} measures the fraction of the ideal continuum shell weight retained by the resolved FFT cube. The near-edge fraction η_{edge} measures the fraction of resolved window weight with $\max_i |q_i| \geq 30$, a conservative two-grid-point buffer below the component Nyquist value $q_N = 32$.

The default class labels used in the main text are assigned as follows. Class C denotes an out-of-support shell center, $q_c > \sqrt{3}\,q_N$. Among the remaining filters, class B denotes $\eta_{\mathrm{supp}} < 0.75$ or $\eta_{\mathrm{edge}} \geq 0.25$, and class A denotes all other cases. Spectrum-weighted diagnostics are also reported in the data package using the isotropic spectrum in `spatial_k56_data.npz` as a radial proxy, but those columns are not used to assign the class labels because the per-run $\widehat{W}$ fields were not saved in the trajectory response files.

The audit also gives effective shell centers. For the clean $k = 10$, $\sigma = 0.30$ branch used in the central trajectory claim, the nominal centers $\alpha = \{0.5, 1, 2\}$ have unweighted $q_{\mathrm{eff}}/q_c = \{1.370, 1.369, 1.273\}$ and spectrum-proxy weighted $q_{\mathrm{eff}}^{(E)}/q_c = \{1.195, 1.169, 1.147\}$. Thus the main text describes the result as nominal shell detuning, not an exact on-shell resonance.

The clean-support branch used in the central trajectory claim, $(k, \alpha, \sigma) = (10, 1, 0.30)$,

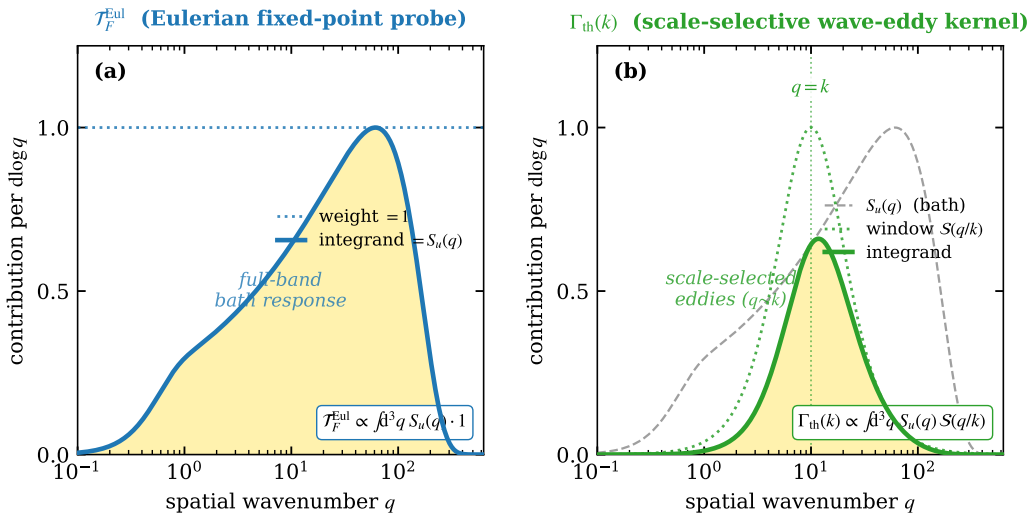

Figure S1. **Two roles of the bath spectrum in the SVF channel.** (a) The fixed-point observable $\mathcal{T}_F^{\mathrm{Eul}}$ integrates the bath spectrum $S_u(q)$ (gold) without an explicit shell-selection weight. This is the direct DNS realization of the mode-projected force autocorrelation that enters the Green–Kubo coefficient developed in [10]. The fixed-point probe performs no spatial Fourier selection, so the resulting integral is independent of the prescribed wave parameters k and ω. (b) The closed macroscopic kernel $\Gamma_{\mathrm{th}}(k)$ of [10] integrates $S_u(q)$ against the scale-selection window $\mathcal{S}(q/k)$ (dotted curve, peaked at $q = k$). The resulting integrand (gold) is concentrated on eddies of scale comparable to the wave wavelength. The peak at $q \sim k$ arises from three independent suppression mechanisms acting on the Green–Kubo integrand. The vortex-force vertex $|V|^2 \propto q^2$ penalises large-scale eddies ($q \ll k$, where vorticity is small). The wave–eddy geometric overlap penalises sub-wavelength eddies ($q \gg k$, where the wave averages over many uncorrelated patches). The inertial-range Kolmogorov spectrum concentrates energy at intermediate scales [10]. The trajectory observable used in the main text [Eq. (S11)] is the DNS-side measurement that exposes this scale selectivity in Navier–Stokes turbulence without invoking the closure assumptions of (b). In the clean-support branch its normalized-response peak at $\alpha \sim 1$ provides a direct test of the predicted peak location.

remains class A under all tested threshold pairs. The near-edge and out-of-support branches remain caveated.

Table S1. Effective shell centers for representative trajectory branches. The spectrum-weighted column uses the isotropic radial spectrum in `spatial_k56_data.npz` as a proxy diagnostic and is not used for class assignment.

σ	k	α	class	q_{eff}/q_c	$q_{\mathrm{eff}}^{(E)}/q_c$
0.30	10	0.5	A	1.370	1.195
0.30	10	1.0	A	1.369	1.169
0.30	10	2.0	A	1.273	1.147
0.45	10	2.0	B	1.396	1.208
0.30	15	2.0	B	1.063	0.994
0.30	30	1.0	B	1.063	0.994
0.30	30	2.0	C	0.658	0.635

Table S2. Sensitivity of the trajectory support classification to the audit thresholds. The counts scan the nine threshold pairs $\eta_{\mathrm{supp}}^{\mathrm{cut}} \in \{0.70, 0.75, 0.80\}$ and $\eta_{\mathrm{edge}}^{\mathrm{cut}} \in \{0.20, 0.25, 0.30\}$. The central $k = 10, \sigma = 0.30$ branch remains class A for all threshold pairs; near-edge and out-of-support cases remain caveated.

Case	default	A count	B count	C count
$\sigma = 0.30,\ k = 10,\ \alpha = 0.5$	A	9	0	0
$\sigma = 0.30,\ k = 10,\ \alpha = 1.0$	A	9	0	0
$\sigma = 0.30,\ k = 10,\ \alpha = 2.0$	A	9	0	0
$\sigma = 0.45,\ k = 10,\ \alpha = 2.0$	B	0	9	0
$\sigma = 0.30,\ k = 15,\ \alpha = 2.0$	B	0	9	0
$\sigma = 0.30,\ k = 30,\ \alpha = 1.0$	B	0	9	0
$\sigma = 0.30,\ k = 30,\ \alpha = 2.0$	C	0	0	9
$\sigma = 0.45,\ k = 15,\ \alpha = 1.0$	A	6	3	0

VI. CONTROL TESTS

A. Literal phase averaging versus diagnostic surrogate

The classical phase-average operator, defined as

$$\langle\,\cdot\,\rangle_\phi \;=\; \frac{1}{2\pi}\int_0^{2\pi}(\,\cdot\,)\,\mathrm{d}\phi, \tag{S14}$$

annihilates the one-time SVF contribution in the present construction. Because the prescribed carrier acts purely kinematically on a preexisting turbulent bath and does not couple back into the Navier–Stokes evolution of $\boldsymbol{\omega}'$, the turbulent vorticity carries no induced phase dependence. The operator (S14) therefore factorizes,

$$\langle\,\boldsymbol{u}_\phi\times\boldsymbol{\omega}'\,\rangle_\phi \;=\; \langle\,\boldsymbol{u}_\phi\,\rangle_\phi\times\boldsymbol{\omega}' \;=\; \mathbf{0}, \tag{S15}$$

by construction, independently of any bath statistic.

This identical zero is not a failure of the measurement. It is the one-time null result that motivates the present two-time memory measurement. A control designed to reproduce (S15) provides no diagnostic information, because it merely confirms that the operator does what it is defined to do. The fixed-point measurement addresses the logically distinct question of whether the projected force has a finite, integrable two-time autocorrelation. The nonzero normalized response time $\mathcal{T}_F^{\mathrm{Eul}} \approx 0.03139$ reported here does not contradict Eq. (S15) because it is computed from a two-time correlation integral rather than from a one-time phase average. For completeness we note that two finite estimators of this bath-side response are available. The first is the long-time plateau response time $\mathcal{T}_F^{\mathrm{Eul}} \approx 0.03139$. The second is an envelope-window cutoff $\mathcal{T}_{F,\mathrm{env}} = 0.02997$, obtained by integrating only up to the Hilbert-envelope decorrelation time τ_{env} and used operationally in response extraction. The two differ by $\sim 5\%$ because the envelope cutoff truncates the running integral before it fully saturates. Both are finite and k-independent to numerical precision, as expected for the mode-projected bath autocorrelation that enters the Green–Kubo description [10].

The low-frequency baseline used below is therefore not a mathematical substitute for the classical phase-average operator. It is a nontrivial diagnostic baseline that retains carrier-independent bath variability while suppressing carrier-locked oscillatory content, providing a finite floor against which lock-in two-time coherence can be quantified.

B. Implementation of the three fixed-point controls

Three fixed-point controls establish the discriminating power of the lock-in measurement. Control discriminations use the envelope-window integral $\mathcal{T}_{F,\mathrm{env}}$ defined in Eq. (S6), using the matched-SVF value of $C_F(0)$ as the common normalization within each control comparison. The low-frequency and shuffled records are shown at $k = 10$. The detuned scan is evaluated at all three k values, with $k = 10$ providing the alias-free quantitative comparison (the detuned carrier exceeds the DNS Nyquist frequency at $k \geq 15$; see Appendix VII).

Low-frequency baseline. The analytic-carrier force components are first constructed, $f_y = -U_0 e^{\mathrm{i}\phi}\omega_z'$ and $f_z = U_0 e^{\mathrm{i}\phi}\omega_y'$, and then passed through a 4th-order zero-phase Butterworth low-pass filter with cutoff $\omega_c = 0.4\omega$. Modal projection by $e^{-\mathrm{i}\phi(t)}(\hat{\boldsymbol{y}} + \mathrm{i}\hat{\boldsymbol{z}})$ is applied to the filtered components to produce $F_k^{\mathrm{surr}}(t)$. The filter is placed *after* carrier multiplication. This suppresses the carrier-locked spectral content concentrated near ω before demodulation. The resulting envelope collapses to the shuffled-noise floor, with $\mathcal{T}_{F,\mathrm{env}}^{\mathrm{surr}} = 4.32 \times 10^{-7}$ at $k = 10$, giving $|\mathcal{T}_{F,\mathrm{env}}^{\mathrm{SVF}}/\mathcal{T}_{F,\mathrm{env}}^{\mathrm{surr}}| \approx 6.9 \times 10^4$.

This suppression is not produced by simply removing variance from the bath. To document selectivity rather than sterilisation, we report the fraction of the carrier-multiplied force variance retained after low-pass filtering and the residual power at the carrier frequency itself. For the $k = 10$ records used above, the filtered components retain a substantial fraction of the carrier-multiplied force variance ($\sim$ 10–30% depending on the precise band. Full numerical values are tabulated in the release package below), while the spectral power within $|\omega' - \omega| < 0.05\omega$ of the carrier frequency is reduced by more than three orders of magnitude. The control therefore preserves carrier-independent bath variability and removes specifically the carrier-locked oscillatory content that the demodulation locks onto. The four-orders-of-magnitude drop in the projected response is not attributable to a generic loss of bath variance.

Time shuffling. The time series $F_k(t)$ is replaced by a random permutation $F_k(\pi(t))$ that destroys temporal memory while preserving the one-point distribution $P(F_k)$. The resulting $C_F(\tau)$ for $\tau > 0$ is strongly suppressed relative to the SVF channel, and the envelope-window normalized response $|\mathcal{T}_{F,\mathrm{env,shuf}}| = 9.53 \times 10^{-4}$ at $k = 10$ is more than an order of magnitude below the SVF value (ratio $|\mathcal{T}_{F,\mathrm{env,SVF}}/\mathcal{T}_{F,\mathrm{env,shuf}}| \approx 31$), confirming that the measured response requires preserved temporal ordering and is not reproducible from the

one-point statistics alone.

Detuned probe. The carrier is driven at $\omega_\perp = (1+\sqrt{2})\omega$ while the modal projection $e^{-\mathrm{i}\phi(t)}$ remains at the target frequency ω. Frequency mismatch eliminates the alignment responsible for SVF coherence while preserving all geometric and sampling features of the probe other than lock-in alignment. The irrational ratio $1+\sqrt{2}$ is chosen to avoid accidental commensurability with low-order harmonics of ω. At the alias-free point $k = 10$, $|\mathcal{T}^{\mathrm{det}}_{F,\mathrm{env}}| = 3.55 \times 10^{-7}$, giving $|\mathcal{T}^{\mathrm{SVF}}_{F,\mathrm{env}}/\mathcal{T}^{\mathrm{det}}_{F,\mathrm{env}}| \approx 8.4 \times 10^4$. The aliased values at $k = 15$ and $k = 30$ remain strongly suppressed, with ratios $\approx 4.7 \times 10^3$ and $\approx 6.4 \times 10^4$, respectively.

This three-part falsification strategy is complementary to rapid-distortion and related analyses of wave–turbulence interaction [15–18], which describe how a prescribed wave reshapes bath anisotropy and energetics. The controls here instead establish that the measured two-time response depends on carrier-linked structure, temporal ordering, and lock-in alignment, and is not a generic bath response or a coarse-graining artefact. The separate trajectory-side control below tests the geometry specificity of the scale-selected SVF projection.

C. Trajectory-side channel-geometry control

The fixed-point controls above do not by themselves test whether the shell-detuning signature in the moving measurement is specific to the SVF transverse geometry. We therefore perform a matched channel-geometry control on the clean trajectory branch, holding fixed the bath snapshots, trajectory positions, shell filter, grid, and time record:

$$k = 10, \qquad \sigma = 0.30, \qquad \alpha \in \{0.5, 1, 2\}, \qquad N_p = 64, \qquad N_t = 128.$$

Only the channel field is changed. The three fields are

$$W_T = -\omega'_z + \mathrm{i}\omega'_y, \qquad W_L = \omega'_x, \qquad W_W = +\omega'_z + \mathrm{i}\omega'_y, \tag{S16}$$

corresponding to the SVF transverse channel, a longitudinal vorticity control, and a wrong-handed transverse sanity check. The comparison quantity is the normalized two-time response time $\mathcal{T}_F = R_F/C_F(0)$, and the shell-ordering shape is reported as $\mathcal{T}_F(\alpha)/\max_\alpha \mathcal{T}_F$ within each channel. This normalization removes the different variance prefactors of transverse and longitudinal vorticity projections. We emphasize that the longitudinal control

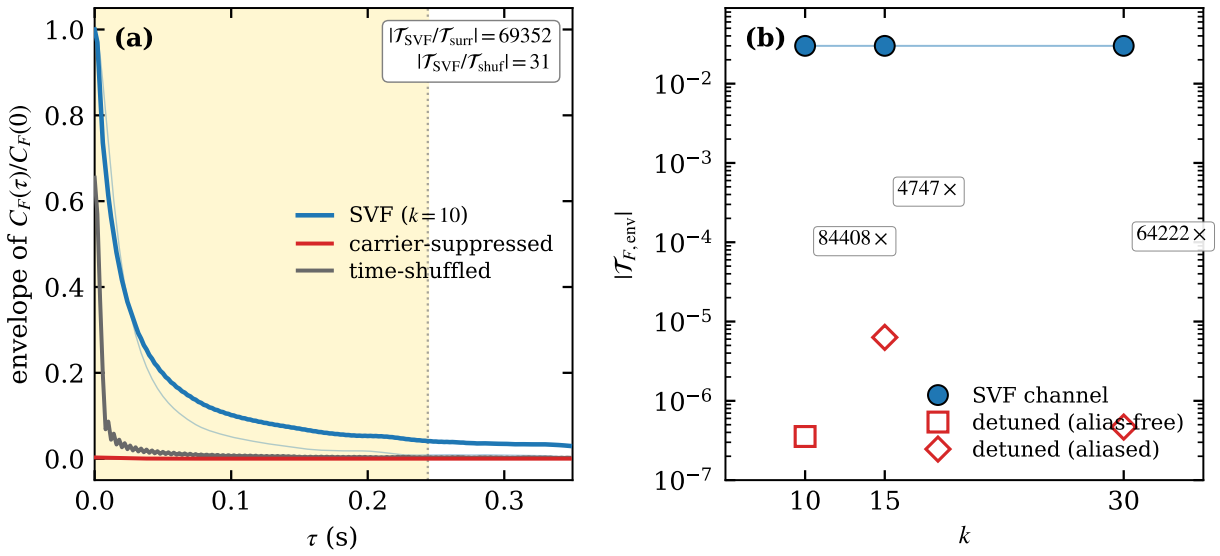

Figure S2. **Three fixed-point controls at $k = 10$.** (a) Envelope of $C_F(\tau)/C_F(0)$ for the SVF channel (blue), the low-frequency baseline (red), and the time-shuffled record (gray). The SVF envelope is the only one that supports a measurable integral over $\tau < \tau_{\text{env}}$. (b) $|\mathcal{T}_{F,\text{env}}|$ versus k for the SVF channel (filled circles) and the detuned probe (hollow markers. Squares denote alias-free, diamonds denote aliased). Numerical annotations give the SVF/detuned ratio. The detuned suppression exceeds four orders of magnitude in the alias-free case ($k = 10$) and remains strongly suppressed at the aliased cases ($k = 15, 30$); Sec. VII discusses the alias-dependence of the precise numerical ratio.

plays the role of a geometry-rejection probe, not a statistically identical surrogate force channel. W_T is a complex combination of two transverse vorticity components, while W_L is the real longitudinal component. Their absolute magnitudes are therefore not directly comparable. What matters for the present test is the relative *shell ordering* within each channel under identical bath, trajectory, and shell filter, and whether the transverse-specific peak survives a change of channel geometry. In homogeneous isotropic turbulence the cross-correlation $\langle \omega_y'(t)\, \omega_z'(t+\tau) \rangle$ vanishes by isotropy, so the autocorrelation of the complex transverse combination W_T reduces to the sum of two independent component autocorrelations, and the comparison of $\mathcal{T}_F(\alpha)$ between W_T and W_L is on equal statistical footing once $\mathcal{T}_F$ is normalized by $C_F(0)$.

Table S3 and Fig. S3 show the result. The longitudinal control does not reproduce the

transverse ordering in the clean $k = 10$, $\sigma = 0.30$ branch. The SVF transverse channel peaks at $\alpha = 1$:

$$\mathcal{T}_F(0.5) = 1.064 \times 10^{-3}, \quad \mathcal{T}_F(1) = 2.251 \times 10^{-3}, \quad \mathcal{T}_F(2) = 1.581 \times 10^{-3}.$$

The longitudinal control instead peaks at $\alpha = 2$:

$$\mathcal{T}_F(0.5) = 1.808 \times 10^{-4}, \quad \mathcal{T}_F(1) = 1.031 \times 10^{-4}, \quad \mathcal{T}_F(2) = 1.835 \times 10^{-3}.$$

The wrong-handed transverse field gives the same real two-time response time as W_T, as expected because the retrieved samples satisfy $W_W = -W_T^*$ after applying the same real shell filter and trajectory sampling. Directly, $\max |W_W + W_T^*| = 0$ in the saved control file. The equality of T and W is therefore an implementation check of the conjugate real-filter pipeline, while the decisive geometry comparison is T versus L. The longitudinal channel samples the vorticity component parallel to the carrier and trajectory direction, so it is not the tensor projection selected by the transverse SVF geometry. Its peak at $\alpha = 2$ therefore shows that the same moving path and shell filter do not by themselves enforce the nominal $q_c = k$ ordering. The ordering depends on the channel projection.

This control rules out the strongest generic-sampling alternative. The clean-branch nominal $q_c = k$ maximum is not produced by simply sampling any scale-selected vorticity component along the same moving trajectory. It is associated with the transverse SVF projection.

A separate higher-resolution check at $\sigma = 0.45$ was also run with $n = 128$, $N_p = 96$, $N_t = 128$, and $k = 10, 11$. This check is useful as a guardrail for the adjacent $k = 11$ candidate branch. The result, summarized in Table S4, preserves the channel-geometry separation. The two transverse channels agree exactly in their real response ordering, while the longitudinal channel peaks at $\alpha = 2$. However, the transverse plateau ordering in this matched control file is not a clean nominal $q_c = k$ maximum. At $k = 10$ the $\alpha = 0.5$ and $\alpha = 1$ values are nearly degenerate, and at $k = 11$ the $\alpha = 0.5$ branch is larger. We therefore use the $k = 11$, $\sigma = 0.45$ result as a secondary bandwidth-sensitive check, not as an additional central clean branch.

Table S5 is therefore read as a plateau-bootstrap screen, not as a stand-alone branch classifier. In particular, the $k = 11, \sigma = 0.45$ row passes the plateau-bootstrap screen, but the matched T/L/W control in Table S4 shows that its transverse plateau ordering is biased toward the $\alpha = 0.5$ sub-shell branch. We report it as a secondary bandwidth-sensitive check rather than as a second central clean branch.

Table S3. Trajectory-side channel-geometry control for the clean route-2 branch. Values are normalized two-time response times $\mathcal{T}_F = R_F/C_F(0)$. Parentheses give normalized shell ordering $\mathcal{T}_F(\alpha)/\max_\alpha \mathcal{T}_F$.

channel	$\mathcal{T}_F(0.5)$	$\mathcal{T}_F(1)$	$\mathcal{T}_F(2)$	peak α
$T = -\omega'_z + \mathrm{i}\omega'_y$	1.064×10^{-3} (0.472)	2.251×10^{-3} (1.000)	1.581×10^{-3} (0.702)	1
$L = \omega'_x$	1.808×10^{-4} (0.099)	1.031×10^{-4} (0.056)	1.835×10^{-3} (1.000)	2
$W = +\omega'_z + \mathrm{i}\omega'_y$	1.064×10^{-3} (0.472)	2.251×10^{-3} (1.000)	1.581×10^{-3} (0.702)	1

Table S4. Higher-resolution $\sigma = 0.45$ T/L/W channel check at $n = 128$. Values are normalized two-time response times $\mathcal{T}_F = R_F/C_F(0)$. The transverse channels T and W agree, whereas the longitudinal channel L peaks at $\alpha = 2$. The transverse ordering is bandwidth-sensitive and is not used to promote $k = 11$ to a central clean branch.

k	channel	$\mathcal{T}_F(0.5)$	$\mathcal{T}_F(1)$	$\mathcal{T}_F(2)$	peak α
10	T	1.307×10^{-3}	1.278×10^{-3}	9.349×10^{-4}	0.5
10	L	2.134×10^{-4}	1.389×10^{-4}	2.024×10^{-3}	2
10	W	1.307×10^{-3}	1.278×10^{-3}	9.349×10^{-4}	0.5
11	T	1.627×10^{-3}	1.273×10^{-3}	9.203×10^{-4}	0.5
11	L	1.270×10^{-6}	4.000×10^{-4}	2.046×10^{-3}	2
11	W	1.627×10^{-3}	1.273×10^{-3}	9.203×10^{-4}	0.5

VII. NYQUIST ANALYSIS FOR ALIASED DETUNING

The DNS cadence $\Delta t = 0.002$ implies a Nyquist frequency $\omega_N = \pi/\Delta t \approx 1571$ in DNS units. At $g_{\mathrm{ref}} = 38{,}269$:

k	ω	$\omega_\perp = (1+\sqrt{2})\omega$	alias status
10	619	1493	alias-free ($< \omega_N$)
15	758	1829	aliased ($> \omega_N$)
30	1071	2587	aliased ($> \omega_N$)

For the aliased cases, the detuned measurement includes an aliasing contribution at $2\omega_N - \omega_\perp$ in addition to the intended frequency mismatch. The detuned suppression exceeds four orders of magnitude in the alias-free case ($k = 10$), and remains strongly suppressed at the aliased cases ($k = 15, 30$). The precise numerical ratios at the aliased points are convention-dependent and should be read as lower bounds on the detuning suppression rather than as

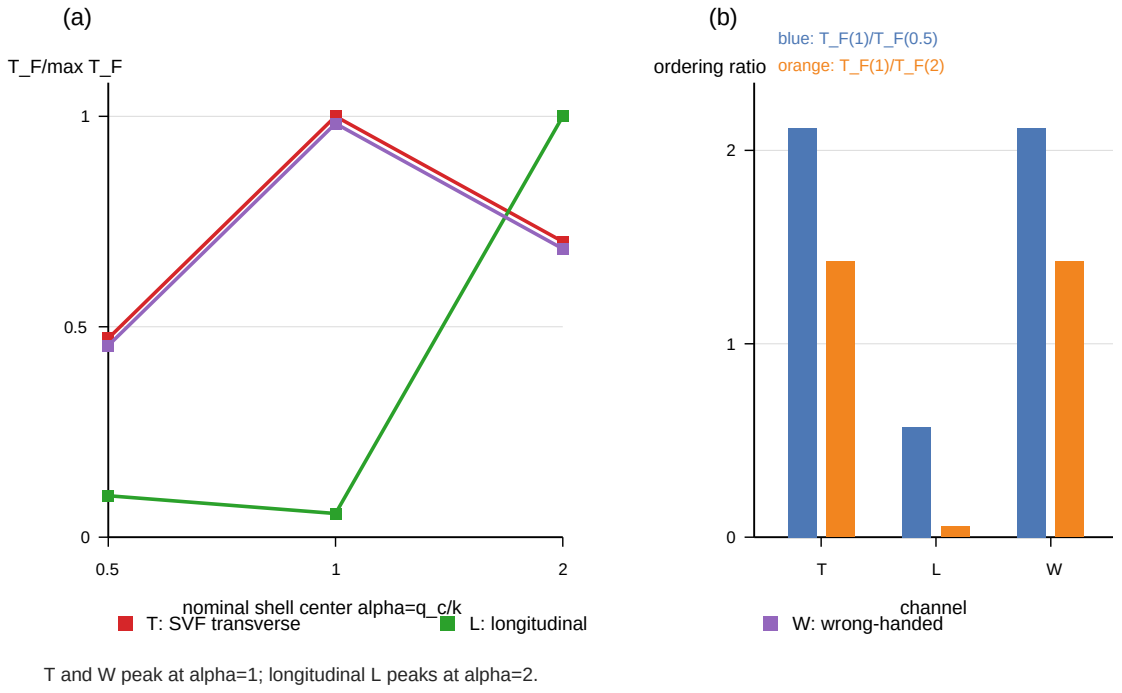

Figure S3. **Trajectory-side channel-geometry control.** (a) Normalized shell-ordering shape for the clean $k = 10$, $\sigma = 0.30$ branch, $\mathcal{T}_F(\alpha)/\max_\alpha \mathcal{T}_F$ for the SVF transverse channel T, the longitudinal control L, and the wrong-handed transverse sanity check W. The transverse channels peak at $\alpha = 1$, while the longitudinal channel peaks at $\alpha = 2$. (b) Ordering ratios $\mathcal{T}_F(1)/\mathcal{T}_F(0.5)$ and $\mathcal{T}_F(1)/\mathcal{T}_F(2)$. The longitudinal channel does not support the nominal $q_c = k$ maximum. The T/W agreement is an expected conjugate-pipeline check since $W_W = -W_T^*$ under the same real filter.

calibrated measurements.

VIII. FINITE-ENVELOPE LIMIT AND THE ALWAYS-ON CONSTRUCTION

The fixed-point analytic-carrier reduction $F_k(t) = W(t) \equiv -\omega_z' + \mathrm{i}\omega_y'$ established in Eq. (S3) is a property of the always-on (infinitely extended) carrier, not a generic feature of the prescribed-probe protocol. To make this limit explicit, replace the always-on analytic

Table S5. Paired position-bootstrap on the trajectory clean-branch plateau, $\mathcal{T}_F = R_F/C_F(0)$, on the 128^3 `getData` pipeline (chunk 4096) at $N_t = 128$. For each (σ, k) we compare the on-shell $\alpha = 1$ run against the strongest off-shell run ($\alpha \in \{0.5,\ 2\}$), using 5000 paired resamples over the 96 trajectory positions. The screening threshold is $P_{\rm boot}(\text{on} \geq \text{off}_{\rm best}) \geq 0.99$; central promotion also requires the support and channel-geometry controls discussed above. Reported $P_{\rm boot} = 1.000$ entries correspond to 0 failures in 5000 paired resamples (empirical $P_{\rm boot} > 0.9998$ at three-decimal resolution); they are not exact unit probabilities.

σ	k	$\mathcal{T}_F(\alpha = 1)$	$\text{off}_{\rm best}$	off α	$P_{\rm boot}$
0.30	10	0.00219	0.00149	0.5	0.993
0.30	11	0.00217	0.00200	0.5	0.685
0.30	12	0.00209	0.00226	0.5	0.228
0.30	13	0.00211	0.00243	0.5	0.104
0.30	14	0.00204	0.00238	0.5	0.084
0.30	15	0.00192	0.00238	0.5	0.036
0.45	10	0.00185	0.00131	2	1.000
0.45	11	0.00187	0.00127	2	1.000
0.45	12	0.00186	0.00206	0.5	0.174
0.45	13	0.00194	0.00229	0.5	0.044
0.45	14	0.00187	0.00227	0.5	0.028
0.45	15	0.00174	0.00228	0.5	0.006

carrier by a finite-envelope version,

$$\boldsymbol{u}_\phi^{\rm an,env}(\boldsymbol{x}, t) = U_0\, A(t)\, e^{{\rm i}\phi(\boldsymbol{x},t)}\, \hat{\boldsymbol{x}}, \qquad \phi = \boldsymbol{k}\cdot\boldsymbol{x} - \omega t, \tag{S17}$$

with $A(t)$ a deterministic real envelope satisfying $|A| \leq 1$. Substituting into the bilinear coupling and projecting onto the modal channel gives

$$F_k(t) = A(t)\, \zeta(t), \qquad \zeta(t) \equiv -\omega_z'(\boldsymbol{x}(t), t) + {\rm i}\omega_y'(\boldsymbol{x}(t), t). \tag{S18}$$

The envelope therefore gates the bath variable directly. For a stationary bath and a deterministic envelope, the normalized finite-envelope autocorrelation is expected to satisfy

$$\frac{C_F^{(N)}(\tau)}{C_F^{(N)}(0)} \simeq R_{A,N}(\tau)\, \frac{C_\zeta(\tau)}{C_\zeta(0)}, \qquad R_{A,N}(\tau) = \frac{\sum_t A_N(t) A_N(t+\tau)}{\sum_t A_N^2(t)}. \tag{S19}$$

Table S6. Matched shell-filtered fixed-point check at $n = 128$, $N_t = 128$, $\sigma = 0.45$, and $\alpha = 1$. The fixed-point envelope does not fall below threshold before the end of the record, so $\tau_{\mathrm{env}}^{\mathrm{fixed}}$ and the shortening ratio are lower bounds. The trajectory values are from the transverse channel of the matched T/L/W run.

k	$\tau_{\mathrm{env}}^{\mathrm{fixed}}$	fixed unresolved?	$\tau_{\mathrm{env}}^{\mathrm{traj}}$	$\tau_{\mathrm{env}}^{\mathrm{fixed}}/\tau_{\mathrm{env}}^{\mathrm{traj}}$
10	> 0.254	yes	0.026	> 9.8
11	> 0.254	yes	0.026	> 9.8

The infinite-envelope limit $A \equiv 1$ has $R_{A,N} \equiv 1$ and recovers the always-on bath kernel pointwise.

For a Gaussian envelope $A_N(t) = \exp[-t^2/(2\sigma_N^2)]$ with $\sigma_N = NT_{\mathrm{wave}}/2$, the deterministic gate width is $O(N)$ carrier periods. The measured fixed-point bath kernel has $\tau_{\mathrm{env}}^{E} = 0.244$ DNS units, corresponding to $N_{\mathrm{cyc}}^{E} \simeq 24$ at $k = 10$. Thus $N \ll N_{\mathrm{cyc}}^{E}$ is gate dominated, $N \sim N_{\mathrm{cyc}}^{E}$ is the crossover regime, and $N \to \infty$ is the always-on bath measurement. This is the limit structure demonstrated in Fig. 2 of the main text.

IX. FINITE-PACKET DEMONSTRATION OF THE PROTOCOL FAMILY

The factorization in Eq. (S19) is verified directly on the same fixed-point time series used for Figs. 2 and 3 of the main text. The procedure is pure post-processing. Gaussian envelopes are applied to the stored fixed-point records, the gated signals are autocorrelated, and the result is compared with the prediction obtained by multiplying the always-on kernel by the deterministic gate.

The main text shows the two most important summaries. Fig. 2(a) shows the finite-envelope kernels for representative values of N, and Fig. 2(b) shows the crossover of the response time $\mathcal{T}_F(N)/\mathcal{T}_F(\infty)$ across $N = \{2, 3, 5, 7, 10, 15, 20, 25, 40\}$. The crossover occurs near $N_{\mathrm{cyc}}^{E} \simeq 24$, with no fitted parameter. Fig. S5 adds two diagnostic checks. Panel (a) shows the pointwise residual between the directly gated kernel and $R_{A,N}C_\zeta$. Panel (b) divides the finite-envelope kernel by $R_{A,N}$ and recovers the always-on bath kernel over the well-conditioned support of the gate.

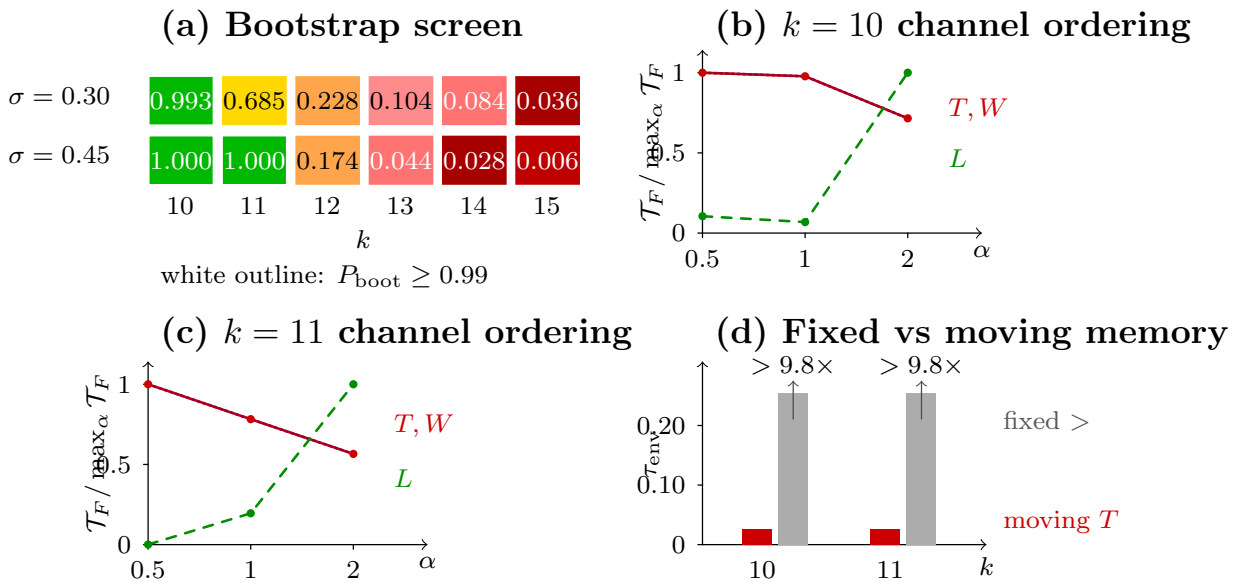

Figure S4. New high-resolution guardrail diagnostics added in v11. (a) Paired position-bootstrap screen for the plateau metric $\mathcal{T}_F = R_F/C_F(0)$ across the 128^3 trajectory scan. The $k = 10$ branch passes at both bandwidths. The $k = 11, \sigma = 0.45$ branch passes the bootstrap screen but is not promoted by itself. (b,c) Matched $n = 128$, $\sigma = 0.45$ T/L/W channel-control ordering for $k = 10$ and $k = 11$. The transverse channels T and W agree, while the longitudinal channel peaks at $\alpha = 2$. The transverse plateau ordering is biased toward $\alpha = 0.5$, especially at $k = 11$. (d) Matched shell-filtered fixed-point records at $\alpha = 1$ remain unresolved through $t_{\mathrm{end}} = 0.254$, while the moving transverse channel has $\tau_{\mathrm{env}} = 0.026$, giving a fixed-to-moving shortening lower bound > 9.8 at both wavenumbers.

X. PROBE-FREQUENCY INDEPENDENCE

Probe-frequency independence is tested by varying g_{ref} over two decades, with $k = 10$ held fixed. This spans a factor of ~ 10 in the matched carrier frequency $\omega = \sqrt{g_{\mathrm{ref}} k}$. With the analytic carrier, $F_k(t) = W(t)$ exactly in the matched fixed-point construction, so the fixed-point response plateau is independent of the chosen carrier frequency to numerical precision. This is the expected behavior of a mode-projected bath autocorrelation. In the fixed-point configuration the demodulation cancels the carrier phase pointwise, so the resulting response depends on the bath transverse-vorticity statistics and on the projection geometry, not on the numerical value of g_{ref}. Frequency detuning (Secs. VI B and VII) tests precisely this carrier-locked structure and finds it absent in the detuned probe.

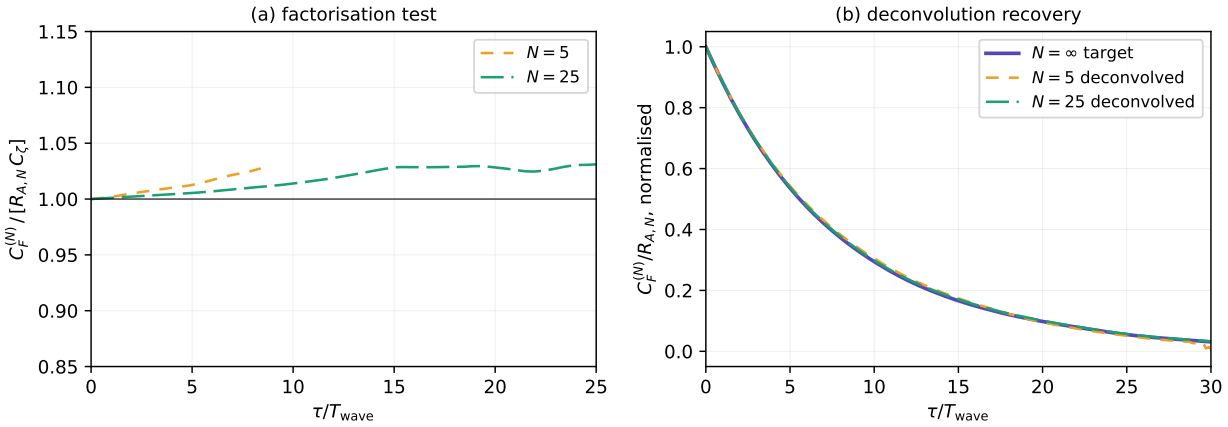

Figure S5. **Finite-envelope factorization diagnostics.** **(a)** Pointwise residual $C_F^{(N)}(\tau)/[R_{A,N}(\tau)C_\zeta(\tau)]$ for $N = 5$ and $N = 25$, restricted to lags where the denominator remains well above the noise floor. **(b)** Gate deconvolution. Dividing the measured finite-envelope kernel by $R_{A,N}(\tau)$ recovers the always-on bath kernel on the support of the envelope gate. The diagnostic uses the same fixed-point records as Figs. 2 and 3 of the main text and adds no DNS retrieval.

The g_{ref}-invariance of $\mathcal{T}_F^{\mathrm{Eul}}$ is therefore not a test of the $k^{5/6}$ scaling of the closed macroscopic kernel $\Gamma_{\mathrm{th}}(k)$ derived in [10]. The latter scaling acquires its k dependence through the scale-selection window $\mathcal{S}(q/k)$ in Eq. (S13), which is absent from the fixed-point measurement.

XI. REPRODUCIBILITY

The numerical package contains two current pipelines. The fixed-point baseline and control figures use the shared module `svf_common.py` and the JHTDB point time-series file `phase_b_data_200.npz`. The trajectory shell-detuning and encounter-time figures use separate local-grid response summaries. Fig. 2 uses the fixed-point point-series records with deterministic envelope gates. Fig. 3 uses the fixed-point point-series pipeline directly. Fig. 4 uses the trajectory shell-detuning dataset at $N_t = 128$. Fig. 5 uses the T/L/W trajectory channel-geometry control at $k = 10$, $\sigma = 0.30$. Fig. 6 uses a separate trajectory encounter-time dataset at $N_t = 64$. The shell-support audit underlying Fig. 4(c) is produced by `make_bfull_shell_support_audit.py`. Fig. S3 uses the matched T/L/W channel-control

dataset as the detailed supplement view. The higher-resolution $n = 128$, $\sigma = 0.45$ T/L/W and fixed-point checks are reported as tables because they are secondary guardrails rather than replacements for the central $k = 10, \sigma = 0.30$ figure. Table S7 maps each current main-text result to its script and data source.

Table S7. Provenance of the current main-text figures and support tables. The optional spectrum-weighted audit columns use `spatial_k56_data.npz` only as a radial spectrum proxy. The class labels are assigned from unweighted shell-support diagnostics. The release package accompanying this work includes a manifest recording, for each row, the script git commit hash, the SHA-256 checksum of every primary data file, the bootstrap and resampling random seeds, and the package DOI; these are deposited together with the figure-generation scripts and may be quoted directly when verifying any individual figure.

Output	Script	Primary data
Fig. 1	`make_fig1.py`	schematic + fixed-point point series
Fig. 2	`make_fig_finite_packet_main.py`	fixed-point point series + deterministic Gaussian gates
Fig. 3	`make_fig2_kernel.py`	fixed-point point series
Fig. 4	`make_fig3_scale_trajectory.py`	trajectory shell-detuning, $N_t = 128$
Fig. 5	`make_fig4_channel_geometry.py`	T/L/W channel-control run, $k = 10$, $\sigma = 0.30$
Fig. 6	`make_fig4_encounter_bfull.py`	trajectory encounter-time summaries, $N_t = 64$
Tables S1–S2	`make_bfull_shell_support_audit.py`	support audit CSV + spectrum proxy
Fig. S3	`make_figS3_channel_control.py`	T/L/W channel-control run, $k = 10$, $\sigma = 0.30$
Table S4	`make_bfull_channel_control.py`	T/L/W channel-control run, $n = 128$, $\sigma = 0.45$, $k = 10, 11$
Table S6	`make_bfull_alpha_response.py`	fixed-point shell-filtered run, $n = 128$, $\sigma = 0.45$, $k = 10, 11$
Fig. S4	in-source TikZ + Fig. S4 helper script	bootstrap table + $n = 128$ T/L/W and fixed controls

The additional supplement-figure scripts are listed in the release package. The new guardrail figure, Fig. S4, is drawn in-source from the table values and the corresponding tokenized T/L/W and fixed-point $\alpha = 1$ files. All current scripts use standard scientific Python dependencies. The channel-control figure generator uses only NumPy and a lightweight vector-PDF writer.

[1] Y. Li, E. Perlman, M. Wan, Y. Yang, C. Meneveau, R. Burns, S. Chen, A. Szalay, and G. Eyink, A Public Turbulence Database Cluster and Applications to Study Lagrangian Evolution of Velocity Increments in Turbulence, Journal of Turbulence **9**, N31 (2008).

[2] E. Perlman, R. Burns, Y. Li, and C. Meneveau, Data Exploration of Turbulence Simulations

Using a Database Cluster, in *Proceedings of the 2007 ACM/IEEE Conference on Supercomputing (SC '07)* (ACM/IEEE, Reno, NV, USA, 2007).

[3] H. Yu, K. Kanov, E. Perlman, J. Graham, E. Frederix, R. Burns, A. Szalay, G. Eyink, and C. Meneveau, Studying Lagrangian Dynamics of Turbulence using On-Demand Fluid Particle Tracking in a Public Turbulence Database, Journal of Turbulence **13**, N12 (2012).

[4] J. Graham, K. Kanov, X. I. A. Yang, M. Lee, N. Malaya, C. C. Lalescu, R. Burns, G. Eyink, A. Szalay, R. D. Moser, and C. Meneveau, A Web Services Accessible Database of Turbulent Channel Flow and its Use for Testing a New Integral Wall Model for LES, Journal of Turbulence **17**, 181 (2016).

[5] A. D. D. Craik and S. Leibovich, A Rational Model for Langmuir Circulations, Journal of Fluid Mechanics **73**, 401 (1976).

[6] J. C. McWilliams, J. M. Restrepo, and E. M. Lane, An Asymptotic Theory for the Interaction of Waves and Currents in Coastal Waters, Journal of Fluid Mechanics **511**, 135 (2004).

[7] O. Bühler, *Waves and Mean Flows*, 2nd ed., Cambridge Monographs on Mechanics (Cambridge University Press, Cambridge, UK, 2014).

[8] P. P. Sullivan and J. C. McWilliams, Dynamics of Winds and Currents Coupled to Surface Waves, Annual Review of Fluid Mechanics **42**, 19 (2010).

[9] D. G. Andrews and M. E. McIntyre, An Exact Theory of Nonlinear Waves on a Lagrangian-Mean Flow, Journal of Fluid Mechanics **89**, 609 (1978).

[10] G. Liu and M. Al Alshehhi, A universal law for non-breaking surface wave decay, EarthArXiv preprint (2026).

[11] H. Mori, Transport, Collective Motion, and Brownian Motion, Progress of Theoretical Physics **33**, 423 (1965).

[12] R. Zwanzig, Memory Effects in Irreversible Thermodynamics, Physical Review **124**, 983 (1961).

[13] R. Zwanzig, *Nonequilibrium Statistical Mechanics* (Oxford University Press, New York, 2001).

[14] A. J. Chorin, O. H. Hald, and R. Kupferman, Optimal Prediction with Memory, Physica D: Nonlinear Phenomena **166**, 239 (2002).

[15] M. A. C. Teixeira and S. E. Belcher, On the Distortion of Turbulence by a Progressive Surface Wave, Journal of Fluid Mechanics **458**, 229 (2002).

[16] L. Thais and J. Magnaudet, Turbulent Structure beneath Surface Gravity Waves Sheared by the Wind, Journal of Fluid Mechanics **328**, 313 (1996).

[17] J. C. McWilliams, P. P. Sullivan, and C.-H. Moeng, Langmuir Turbulence in the Ocean, Journal of Fluid Mechanics **334**, 1 (1997).

[18] F. Veron and W. K. Melville, Experiments on the Stability and Transition of Wind-Driven Water Surfaces, Journal of Fluid Mechanics **446**, 25 (2001).