entropy

MDPI

Article

How Spectral Bias in Machine-Learning Forecast Emulators Propagates Through the Data-Assimilation Cycle: A Linear–Gaussian Spectral Theory

Nitin Lohan [1], Pankaj Goel [1], Sushil Kumar [2], R. K. Giri [3], Fahdah Falah Ben Hasher [4], Dimitris G. Kaskaoutis [5,6] and Mohamed Zhran [7,*]

[1] Department of Applied Sciences and Humanities, G. L. Bajaj Institute of Technology and Management, Greater Noida 201306, India; lohannitin@gmail.com (N.L.); pankajjnu1@gmail.com (P.G.)
[2] Department of Applied Mathematics, University School of Vocational Studies and Applied Sciences, Gautam Buddha University, Greater Noida 201312, India; sushil.kumar@gbu.ac.in
[3] India Meteorological Department (IMD), Ministry of Earth Sciences (MoES), Lodi Road, New Delhi 110003, India; rk.giriccs@gmail.com
[4] Department of Geography and Environmental Sustainability, College of Humanities and Social Sciences, Princess Nourah bint Abdulrahman University, P.O. BOX 84428, Riyadh 11671, Saudi Arabia; fafalhasher@pnu.edu.sa
[5] Climate and Atmosphere Research Center, The Cyprus Institute, 2121 Nicosia, Cyprus; dkask@noa.gr
[6] Institute for Environmental Research and Sustainable Development, National Observatory of Athens, Palaia Penteli, 15236 Athens, Greece
[7] Public Works Engineering Department, Faculty of Engineering, Mansoura University, Mansoura 35516, Egypt
* Correspondence: mohamedzhran@mans.edu.eg

Abstract

Machine-learning forecast emulators systematically under-represent small-scale variance, a spectral bias documented across architectures. Its effect inside a data-assimilation cycle has been conjectured rather than derived. Here, a closed-form theory is developed in the simplest setting that retains the following mechanism: a linear–Gaussian cycle on a periodic domain with statistically homogeneous errors, in which the covariance recursion separates into independent scalar recursions per wavenumber k. The emulator is modelled as the true propagator composed with a spectral damping $d(k) \in [0,1]$. Three results are proved. First, the filter is overconfident at damped scales by a closed-form gap that grows monotonically with the damping. Second, below an exact damping threshold, the Kalman gain falls beneath a prescribed tolerance, provided that the tolerance exceeds the gain floor set by the process noise, so that observations are assimilated with a weight below that tolerance; at weakly forced, locally growing modes, the actual error diverges. Third, without injected model-error variance, ensemble spread collapses geometrically at $d(k)^2 m(k)^2$ per cycle wherever $d(k)\, m(k) < 1$, against $m(k)^2$ for a perfect-model ensemble at decaying modes. Phase and scale-transfer errors and inflation are quantified as departures from this baseline, and localization is examined numerically. The closed forms of Theorems 1–4 are confirmed by computer algebra and reproduced in 50-digit arithmetic to better than 10^{-42}, and the full matrix recursions agree with the closed forms to better than 10^{-10}. Within the linear theory, the gain-miscalibration formulae bound the error of a fully emulator-driven cycle from below at decaying modes. A Lorenz-96 experiment shows that nonlinearity carries the degradation to undamped scales, where no such bound holds.

Academic Editor: Nikolai Leonenko

Received: 14 September 2026
Revised: 22 October 2026
Accepted: 16 November 2026
Published: 19 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

Keywords: data assimilation; Kalman filter; ensemble Kalman filter; machine-learning emulators; spectral bias; covariance inflation; filter divergence

1. Introduction

Data-driven emulators of atmospheric and oceanic dynamics now rival physics-based numerical weather prediction (NWP) models in headline forecast scores [1,2], and the integration of such emulators with data assimilation (DA) is an active field of research [3,4]. A robust and repeatedly documented deficiency of deterministic, regression-trained emulators is spectral bias: the systematic under-representation of small-scale (high-wavenumber) variance, visible as a progressive smoothing of forecast fields [5]. Smith et al. [6] demonstrated an inherent spectral bias in reservoir-computing emulators of surface quasi-geostrophic turbulence, aggravated by temporal subsampling of the training data and resembling numerical diffusion. Chattopadhyay and Hassanzadeh [7] connected the same bias to the long-term instability of autoregressive emulators. Molinaro et al. [8] showed that deterministic surrogates trained on mean square error (MSE) regress towards the conditional mean, and therefore, have poor spectral resolution on turbulent flows. Subich et al. [9] traced the same smoothing to the double penalty that an MSE loss places on fine-scale features predicted in the wrong place and recovered a sharp effective resolution in GraphCast only by modifying the loss to separate decorrelation from spectral-amplitude error. The phenomenon echoes the generic low-frequency learning bias of neural networks [10]. In operator terms, an ML emulator used as a forecast model multiplies the amplitude of Fourier mode k by a factor $d(k) \leq 1$ that decays towards high wavenumbers.

What has been missing is a theory of what this bias does once the emulator sits inside an assimilation cycle, where forecast-error statistics determine the weight given to observations. The literature is explicit that the question is open. Smith et al. [6] noted "… surmise that such errors will reduce ensemble spread within data assimilation algorithms, which could limit their usefulness within a forecasting system". The conjecture has two parts: a reduction in ensemble spread and a resulting loss of usefulness of the assimilation. The community perspective of Arcucci et al. [3] states (p. 3) that "perturbations in these models often fail to propagate forward in time in a physically consistent manner", posing "a major challenge for their integration into 4D-Var systems" (that is, four-dimensional variational assimilation). The same perspective reports (p. 5) that ML systems "may exhibit weaker error amplification from small-amplitude perturbations in their initial conditions than physical models". Empirical support for the latter has since appeared. Selz and Craig [11] found weaker error growth in an AI-based model than in a physics-based one. Kotsuki et al. [12], cycling an ensemble Kalman filter (EnKF) with an AI-based model, reported weaker error growth requiring larger covariance inflation factors. Slivinski et al. [13], cycling deterministic ML models in an EnKF with real surface-pressure observations, found that the models accumulate small-scale noise and fail to represent the cross-variable covariances that spread observational information. Tian et al. [14] documented unphysical tangent-linear and adjoint sensitivities in leading ML weather models. Standard treatments of DA [15,16] analyse suboptimality due to misspecified covariances, sampling error and model error, and the mathematical ML–DA literature [4] frames learned forecast models within filtering theory. Closest to the present work, Adrian et al. [17] showed empirically and theoretically that filtering estimates built on an ML surrogate can remain accurate over long assimilation windows despite the surrogate's own instability, while Sanz-Alonso and Waniorek [18] proved long-time accuracy bounds for ensemble Kalman filters driven by chaotic and machine-learned dynamics. Those results bound the

https://doi.org/10.3390/e28111276

aggregate accuracy of a surrogate-driven filter; the error, the gain and the spread are not resolved by scale, and spectral bias is not identified as the mechanism. What is missing, and what is supplied here, is a per-scale, closed-form account of how a spectrally biased forecast operator distorts a cycling filter.

This paper supplies that account in the most transparent setting that retains the mechanism: a linear–Gaussian cycle on a periodic domain with statistically homogeneous errors. In this setting, every covariance in the cycle is simultaneously diagonalized by the discrete Fourier transform, and the matrix Riccati recursion of the Kalman filter [19] splits into independent scalar recursions per wavenumber, the same idealization that underlies operational spectral background-error covariance models [16,20]. Within it, five theorems formalize the conjecture and its consequences. First, the biased emulator distorts the actual analysis-error spectrum away from the spectrum the filter believes (Theorems 1–3). Second, it suppresses the Kalman gain at biased scales beyond a sharp, closed-form threshold, producing spectral blind spots in which observations are assimilated with a weight below a prescribed tolerance that exceeds the gain floor $q(k)/[q(k)+r(k)]$, with outright divergence at weakly forced, locally growing scales (Theorem 4). Third, without injected model-error variance, it collapses ensemble spread at the exactly computable geometric rate $d(k)^2 m(k)^2$ wherever $d(k)\,m(k) < 1$, against $m(k)^2$ for a perfect-model ensemble at decaying scales (Theorem 5). With respect to the conjecture of Smith et al. [6], the first part is thereby confirmed within this setting in a sharpened form, with an exact collapse rate and a sharp survival criterion. Because a deterministic perfect-model ensemble also loses spread at decaying scales, the emulator-specific content is the factor $d(k)^2$ in the rate and the displacement of the survival boundary from $m(k) > 1$ to $d(k)\,m(k) > 1$. It is also qualified: the collapse is prevented by any injected model-error variance and, under multiplicative inflation λ, wherever $\lambda\, d(k)\, m(k) > 1$. The second part is made precise by the overconfidence gap and the blind-spot threshold. Four propositions and a corollary then quantify what changes when the idealizations are relaxed: phase errors, scale transfer and time dependence of the damping, covariance inflation, and the dependence of the blind-spot extent on the noise spectra and the observation density (see Sections 3.7–3.9).

Two features of the analysis deserve emphasis at the outset. First, two deployment configurations are distinguished. In the mechanism-isolating configuration (Level 1), the emulator determines the filter's covariances and gain, while the forecast mean is still provided by the physical model; in the emulator-in-the-loop configuration (Level 2), the emulator provides both mean and covariance. The actual error of Level 2 is proved to be bounded below by that of Level 1 at every decaying wavenumber. Within the linear theory, the transparent Level-1 formulae are therefore lower bounds for the realistic deployment; at growing wavenumbers Level 2 has no stationary spectrum, and only Level-1 statements are made. Second, the linear–Gaussian, homogeneous, fully observed setting is a deliberate first-order isolation of the mechanism, in the spirit of homogeneous–isotropic idealizations in turbulence theory; it makes every step provable and every formula checkable, but it is not a model of an operational system. All closed forms are verified symbolically. The scalar closed forms of Theorems 1–4 are reproduced by independent 50-digit computations to better than 10^{-42}, and the full matrix recursions agree with the closed forms to better than 10^{-10} (Section 3.11). A Lorenz-96 experiment (Section 3.6) illustrates the mechanism in a nonlinear system. There, the degradation is largest at the undamped, energy-containing scales, which the linear theory leaves unaffected; the linear per-scale formulae are not bounds for a nonlinear cycle, and none is claimed.

The paper is organized as follows. Section 2 states the assumptions, the diagonalization lemma and the two cycles. Section 3 presents the theorems and the numerical confirmation, followed by extensions beyond the idealized damping model, to inflated cycles, to the sensitivity of the blind-spot extent and to gridpoint and localized ensemble filters

(Sections 3.7–3.10). Section 4 draws out design rules, Section 5 states the limitations, and Section 6 summarizes the conclusions. Proofs follow the corresponding statements, and the verification protocol is summarized in Section 3.11.

2. Theoretical Framework

2.1. System, Emulator Model and Assumptions

On the periodic grid $\{0, \ldots, n-1\}$ consider the linear–Gaussian state-space system

$$x_{j+1}^t = \mathbf{M}\, x_j^t + w_j, \qquad y_j = x_j^t + \epsilon_j, \tag{1}$$

with $w_j \sim \mathcal{N}(0, \mathbf{Q})$ and $\epsilon_j \sim \mathcal{N}(0, \mathbf{R})$ white in time and mutually independent, where x_j^t denotes the true state at cycle j and y_j the observation vector. The ML emulator is modelled as

$$\widetilde{\mathbf{M}} = \mathbf{D}\,\mathbf{M}, \tag{2}$$

where $\mathbf{D}$ is a Fourier multiplier with real symbol $d(k) \in [0,1]$: the emulator applies the true dynamics and then damps each Fourier amplitude by the factor $d(k)$, with $d \approx 1$ at large scales and d small at small scales. Because D and M are both diagonal in the Fourier basis (Assumption A6), the two operators commute, $DM = MD$, and the order of damping and propagation in Equation (2) is immaterial. Two features of this model should be made explicit. First, $d(k)$ is an empirical description of a trained emulator, not a consequence of the idealized system (1). Within that system, the conditional mean of x_{j+1}^t given x_j^t is exactly $\mathbf{M}x_j^t$, and the conditional mean minimizes the expected squared one-step error over all functions of the full state. Therefore, an emulator that attains this minimum has $d \equiv 1$: system (1) alone produces no spectral bias. Damping arises when the emulator's effective input is coarser or noisier than the state it must predict, through unresolved scales, noisy training inputs, temporal subsampling of the training data or finite capacity. Mean-square-error training then returns a conditional mean over the unresolved degrees of freedom and suppresses the variance that is unpredictable from that input [8]. The measured spectra of trained emulators show exactly such scale-selective amplitude loss [5,6]. Second, $d(k)$ is therefore imposed as a diagnosed property of a given emulator, and the question addressed is what a cycling filter does with it; $d(k)$ is not derived from a training problem. The function $d(k)$ is referred to as the damping profile of the emulator. A real, time-independent multiplier is the baseline model. Complex factors (phase errors), errors that transfer variance between wavenumbers and time-dependent damping are treated in Section 3.7, together with a diagnostic for testing whether a given emulator is adequately described by the baseline.

All of $\mathbf{M}$, $\mathbf{D}$, $\mathbf{Q}$, $\mathbf{R}$ and the initial covariance are assumed to be circulant (statistically homogeneous) and thus simultaneously diagonalized by the unitary discrete Fourier transform (DFT) F. The propagator has symbol $M(k) = m(k)\, e^{i\theta(k)}$ with modulus $m(k) > 0$ and phase $\theta(k)$, and the noise spectra are denoted $q(k) > 0$ and $r(k) > 0$. The filter is a Kalman filter that is consistent with its own model: it propagates covariances with the emulator $\widetilde{\mathbf{M}}$ and uses the true q and r; it does not know that $d \neq 1$. Table 1 lists the Assumptions A1–A10 one by one, followed by the symbols used throughout; each proof indicates which assumptions it invokes.

Table 1. Assumptions (upper block, one per row) and symbol glossary (lower block). All per-wavenumber variances are real and non-negative; the argument k is suppressed where no confusion arises.

Item	**Statement or Meaning**
A1	Linear truth $x_{j+1}^t = \mathbf{M}x_j^t + w_j$ on a periodic grid of n points
A2	Emulator $\widetilde{M} = DM$, with D a Fourier multiplier of real symbol $d(k) \in [0,1]$
A3	Gaussian noises w_j and ϵ_j and Gaussian initial condition, white in time and mutually independent
A4	Statistical homogeneity: Q, R and the initial covariance are circulant
A5	Isotropy: spectra depend on $\lvert k\rvert$ only (automatic in one dimension)
A6	Commuting propagators: M and D are both diagonal in the Fourier basis; M has symbol $M(k) = m(k)\, e^{i\theta(k)}$, and only $m(k) = \lvert M(k)\rvert$ enters the variance recursions
A7	Direct, dense observations: $y_j = x_j^t + \epsilon_j$, i.e., $\mathbf{H} = \mathbf{I}$
A8	Stationarity: results concern the fixed points of the covariance recursions
A9	The filter propagates covariances with the emulator and uses the true q, r
A10	(Theorem 5(i)–(iii) and (v) only) mode-wise ensemble gains; relaxed in Theorem 5(iv) and assessed numerically in Section 3.10
Symbol	**Meaning**
$a = d^2m^2$, $a_0 = m^2$	Believed and true forecast-error amplification per cycle
f, p_{bel}^a, g	Believed stationary forecast variance, analysis variance and gain (Theorem 1)
Σ^a, $G = \Sigma^a - p_{\mathrm{bel}}^a$	Actual stationary analysis variance (Theorem 2) and overconfidence gap; $\Sigma(\gamma)$ is the same variance for an arbitrary fixed gain γ, so that $\Sigma^a = \Sigma(g)$
$\varphi = \varepsilon r/(1-\varepsilon)$, $d^*(\varepsilon)$	Believed forecast variance at which $g = \varepsilon$; blind-spot threshold (Theorem 4)
$\Sigma_{\mathrm{free}} = q/(1-m^2)$	Stationary error variance without assimilation ($m < 1$)
v, c	Climatological spectrum and error–state cross moment (Level 2, Theorem 3)
$\kappa = 2\lvert k\rvert/n$	Normalized wavenumber
$\hat{\gamma}$; γ^*, p^*	Optimal gain for a given forecast variance (Lemma 2); gain and analysis variance of the perfect-model filter ($d = 1$)
$u = 1-g$, P_n	Complement of the gain; untracked fraction of a truth perturbation after n cycles (Section 3.2)
$z = d\, e^{i\delta}$	Complex damping factor with phase error δ (Section 3.7)
λ, q_f	Multiplicative inflation factor (on standard deviations) and process-noise variance assumed by the filter (Section 3.8)
$s = q/r$, ρ	Noise ratio and observation density (Section 3.9)
N, ν; s^f, s^a	Ensemble size; real degrees of freedom of the per-mode sample variance, $\nu = 2(N-1)$ at complex Fourier modes and $\nu = N-1$ at $k = 0$ and $k = n/2$; forecast and analysis spread per mode (Theorem 5)

2.2. Per-Wavenumber Diagonalization

Lemma 1 (diagonalization; phases drop out). *Under A1–A7, every covariance matrix generated by the cycle remains circulant, and the matrix covariance recursion separates into* n *independent scalar recursions in the per-mode variances. Only the moduli enter: with* $a(k) := d(k)^2 m(k)^2$ *(believed amplification) and* $a_0(k) := m(k)^2$ *(true amplification), the phase* $\theta(k)$ *cancels identically.*

Proof. A matrix is circulant if and only if it is diagonalized by the unitary DFT F; circulants form a commutative algebra closed under products, sums, adjoints and inverses. By induction, every covariance produced by the recursion $\mathbf{P}^f = \mathbf{A}\mathbf{P}^a\mathbf{A}^* + \mathbf{Q}$, $\mathbf{K} = \mathbf{P}^f(\mathbf{P}^f + \mathbf{R})^{-1}$, $\mathbf{P}^a = (\mathbf{I} - \mathbf{K})\mathbf{P}^f$ with circulant $\mathbf{A}, \mathbf{Q}, \mathbf{R}, \mathbf{P}_0^a$ is circulant, and the eigenvalue identities are the diagonal entries of each matrix equation; the identity observation operator (A7) is what keeps the update within the circulant algebra. Only $|A(k)|^2$ enters, which equals d^2m^2 for $\mathbf{A} = \mathbf{DM}$ (A2: d real) and m^2 for $\mathbf{A} = \mathbf{M}$. □

The lemma is elementary but essential: it is what makes per-wavenumber reasoning exact here, and it is verified as a matrix statement in Section 3.5 by running the full $n \times n$ recursions in gridpoint space and comparing the Fourier diagonal of the result with the scalar theory.

2.3. The Believed Cycle

Per wavenumber (argument k suppressed), the filter's own covariance recursion, termed the believed cycle, reads

$$p_{j+1}^f = a\, p_j^a + q, \qquad g_j = \frac{p_j^f}{p_j^f + r}, \qquad p_j^a = \left(1 - g_j\right) p_j^f, \tag{3}$$

where p^f and p^a are the forecast- and analysis-error variances that the filter carries. Its stationary point defines the gain g that the filter actually applies to the innovations.

2.4. The Actual Cycle: Two Deployment Configurations

For the actual error, two configurations are considered, and it is essential to be explicit about which deployment each describes.

Level 1 (mechanism isolation). The error propagates with the true amplification, while the filter applies the biased stationary gain g of the believed cycle:

$$\Sigma_{j+1}^f = a_0\, \Sigma_j^a + q, \qquad \Sigma_j^a = (1 - g)^2\, \Sigma_j^f + g^2 r, \tag{4}$$

where the second relation is the Joseph form of the analysis update, valid for an arbitrary gain. Level 1 assumes that the forecast mean is still produced by the physical model; equivalently, it isolates the gain-miscalibration mechanism alone. Its deployment reading is a system whose uncertainty model is emulator-based, for example a background-error covariance derived from an ML ensemble of data assimilations [3,16], while the forecast mean comes from the physical model. It is not the full picture for an emulator that also produces the forecast mean.

Level 2 (emulator in the loop). If the forecast mean itself is $\widetilde{\mathbf{M}}x^a$, the realistic deployment in which the emulator provides mean and covariance, then per mode

$$e_{j+1}^f = M(k)\left[\, d\, e_j^a + (d - 1)\, x_j^t \,\right] - w_j, \tag{5}$$

so that the spectral bias acts as state-correlated model error. The variance recursion then couples Σ^a to the error–state cross moment $c = \mathbb{E}\left[\, e^a\, \overline{x^t} \,\right]$ and to the climatological spectrum $v = q/(1 - m^2)$, which requires $m < 1$ for stationarity. The derivation is short and

is given in full. Errors are defined as estimate minus truth. The emulator forecast $x^f_{j+1} = d\,M(k)\,x^a_j$ and the truth $x^t_{j+1} = M(k)\,x^t_j + w_j$ give Equation (5) on writing $x^a_j = x^t_j + e^a_j$. An equivalent form is

$$e^f_{j+1} = M(k)\,e^a_j - (1-d)\,M(k)\,x^a_j - w_j, \tag{6}$$

that is, the Level-1 error propagation plus a bias equal to the part of the analysed state that the emulator fails to propagate. The bias term is thus not an additional modelling assumption: it is fixed by Equation (2), and it is correlated with the state because the analysed state contains the truth. The analysis error obeys $e^a_j = (1-g)\,e^f_j + g\,\epsilon_j$. Let $v_j = \mathbb{E}\left|x^t_j\right|^2$ denote the climatological variance of the mode; $c_j = \mathbb{E}[e^a_j\,\overline{x^t_j}]$, the error–state cross moment; and $\Sigma^a_j = \mathbb{E}\left|e^a_j\right|^2$, the actual analysis variance. Because w_j and ϵ_j are independent of all earlier quantities (Assumption A3), the second moments of Equation (5) close

$$v_{j+1} = m^2 v_j + q, \qquad c_{j+1} = (1-g)\left[m^2\left(d\,c_j - (1-d)\,v_j\right) - q\right], \tag{7}$$

$$F_j = m^2(1-d)^2\,v_j - 2m^2 d(1-d)\,c_j + q, \tag{8}$$

$$\Sigma^a_{j+1} = (1-g)^2\left(a\,\Sigma^a_j + F_j\right) + g^2 r. \tag{9}$$

The system is triangular: v_j evolves autonomously, c_j is driven by v_j, and Σ^a_j is driven by both. The cross moment is real once its transient has decayed, because its imaginary part obeys a homogeneous contraction. The term F_j takes the place of the process-noise variance q in the Level-1 recursion (4); at $d = 1$ it reduces to q, and Level 1 is recovered. A stationary point exists if and only if $m < 1$, in which case $v = q/(1-m^2)$; the remaining contraction factors $(1-g)m^2 d$ and $(1-g)^2 a$ are then below one automatically. For $m \geq 1$, the climatological variance grows without bound, and no stationary Level-2 spectrum exists for any $d < 1$ (Remark 1 in Section 3.2).

Both configurations admit closed forms, and the two are ordered: $\Sigma^a_{\mathrm{L2}} \geq \Sigma^a_{\mathrm{L1}}$ at every wavenumber of the domain on which Level 2 is defined, namely the locally decaying modes $m(k) < 1$, with equality if and only if $d = 1$ (Theorem 3). On those modes, Level 1 is therefore a lower bound, a floor, on the actual error of the realistic emulator-in-the-loop cycle: whenever Level 1 predicts overconfidence or a blind spot there, the full deployment is affected at least as strongly. At locally growing modes ($m \geq 1$), where the climatological spectrum of the Level-2 recursion does not converge, only the Level-1 results apply and are stated as such; that is the regime of the stability wedge of Theorem 2 and of the divergence branch of Theorem 4. Both configurations are presented side by side, and Level 1 is never relied upon alone where the distinction matters.

3. Results

Throughout this section, $m > 0$, $q > 0$, $r > 0$ and $d \in [0,1]$, with $a = d^2 m^2$ and $a_0 = m^2$. Condensed proofs follow each statement; complete proofs with explicit assumption tags, together with the computer-algebra confirmation of every identity, are provided in the code repository (see the Data Availability Statement). The theorems are illustrated with a spectral toy model, defined in Section 3.5, whose propagator and damping profiles are shown in Figure 1. One notation is used for the actual stationary analysis variance throughout: Σ^a at the believed gain, and $\Sigma(\gamma)$ when an arbitrary fixed gain γ is applied, so that $\Sigma^a = \Sigma(g)$; the subscripts L1 and L2 are attached where the two configurations of Section 2.4 must be distinguished. The abbreviation $u = 1 - g$ is used in Sections 3.2 and 3.7.

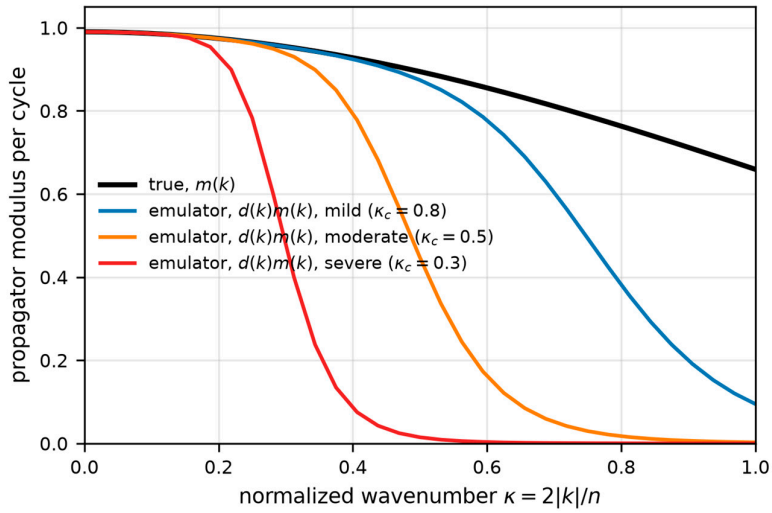


Figure 1. True propagator modulus $m(k)$ (black) and emulator modulus $d(k)m(k)$ for the three damping profiles of the spectral toy model (Section 3.5), as functions of the normalized wavenumber $\kappa = 2|k|/n$. The emulator is faithful at large scales and strongly damped at small scales. The sigmoid profiles are a modelling choice that spans mild to severe small-scale damping; none is fitted to a trained emulator.

3.1. The Believed Analysis Spectrum

Theorem 1 (believed stationary analysis spectrum). *The believed cycle (3) has a unique stationary point, which attracts every initial condition. The stationary believed forecast variance f is the unique positive root of $f^2 + f\,[\,r(1-a) - q\,] - qr = 0$, that is,*

$$f = \frac{1}{2}\Big(q + ar - r + \sqrt{[\,r(1-a) - q\,]^2 + 4qr}\,\Big), \tag{10}$$

$$p^a_{\mathrm{bel}} = \frac{f\,r}{f + r}, \qquad g = \frac{f}{f + r}, \tag{11}$$

and the linearized contraction factor at the fixed point is $a(1-g)^2 < 1$. Moreover, f, p^a_{bel} and g are strictly increasing in $a = d^2m^2$, and $f \geq q$, so that the gain satisfies the floor $g \geq q/(q+r)$ for every $d \in [0,1]$. At $d = 1$, Equations (10) and (11) reduce to the standard stationary scalar Kalman filter.

Proof. The map $\Phi_{\mathrm{bel}}(p) = r(ap + q)/(ap + q + r)$ is continuous, bounded by r, strictly increasing and strictly concave in p for $a > 0$ (constant for $a = 0$), with $\Phi_{\mathrm{bel}}(0) = qr/(q + r) > 0$. Strict concavity follows from $\Phi_{\mathrm{bel}}(p) = r - r^2/(ap + q + r)$, whose second derivative $-2a^2r^2/(ap + q + r)^3$ is negative. The function $h(p) = \Phi_{\mathrm{bel}}(p) - p$ is then strictly concave with $h(0) > 0$ and $h \to -\infty$ as $p \to \infty$, and therefore it has exactly one positive root; monotone bounded orbits give global attraction. The fixed-point identity $f = afr/(f + r) + q$ multiplies out to the stated quadratic, whose roots have product $-qr < 0$, so exactly one is positive. At the fixed point $ar/(f + r) = (f - q)/f < 1$ for $a > 0$, whence the contraction factor $a(1-g)^2 = [\,ar/(f + r)\,][\,r/(f + r)\,] < 1$. Implicit differentiation gives $\partial_a f\,[\,1 - a(1-g)^2\,] = p^a_{\mathrm{bel}} > 0$, so f, p^a_{bel} and g increase strictly in a. Finally, $f = ap^a_{\mathrm{bel}} + q \geq q$ yields the gain floor. □

The gain floor has a direct interpretation. Even a fully collapsed emulator ($d \to 0$) leaves the filter believing a forecast variance q, so observations are never formally discarded as long as the filter carries a process-noise term. Spectral blind spots (Theorem 4)

can therefore arise only where the process-noise spectrum $q(k)$ that the filter uses, which under Assumption A9 is the true one, is small relative to the observation-error variance $r(k)$: the emulator's damping suppresses the gain, and a small ratio $q(k)/r(k)$ is what allows it to be suppressed below a given tolerance. Section 3.9 quantifies this dependence. Both conditions are typically met at small scales, where geophysical process-noise spectra are weak relative to instrument error and emulator damping is strongest. The corresponding quantity in the ensemble setting of Theorem 5 is the injected variance q_{ens}, which is zero for any deterministic forecast step, emulator or physical model, unless it is added explicitly.

3.2. The Actual Analysis Spectrum and the Overconfidence Gap

Two structural identities drive the analysis. The first is an exact decomposition of the one-step analysis variance produced by an arbitrary gain.

Lemma 2 (one-step gain identity). *For $p, r > 0$ and any $\gamma \in \mathbb{R}$,*

$$(1-\gamma)^2 p + \gamma^2 r = \frac{pr}{p+r} + \left(\gamma - \frac{p}{p+r}\right)^2 (p+r). \tag{12}$$

Here and below, $\hat{\gamma} := p/(p+r)$ denotes the optimal (Kalman) gain for a forecast variance p, so that the last term of Equation (12) is $(\gamma - \hat{\gamma})^2(p+r)$. The hat on a gain is used in this sense only.

Proof. Expanding the square on the right-hand side of Equation (12), $pr/(p+r) + \hat{\gamma}^2(p+r) - 2\gamma p + \gamma^2(p+r) = p - 2\gamma p + \gamma^2 p + \gamma^2 r$ with $\hat{\gamma} = p/(p+r)$. □

The first term on the right-hand side is the optimal analysis variance and the second is the quadratic suboptimality penalty. The second identity, Equation (13) below, gives the derivative of the stationary actual variance with respect to the applied gain.

Theorem 2 (actual stationary spectrum and overconfidence gap). *Let g be the believed gain of Theorem 1.*

(i) Stability dichotomy. If $(1-g)^2 a_0 < 1$, the actual cycle (4) has the unique, globally attracting stationary point $\Sigma^a = [(1-g)^2 q + g^2 r]/[1-(1-g)^2 a_0]$; if $(1-g)^2 a_0 \geq 1$, then $\Sigma_j^a \to \infty$ for every initial condition.

(ii) Automatic stability at $d = 1$; a stability wedge for $d < 1$. At $d = 1$ stability always holds, because the actual contraction coincides with the believed contraction of Theorem 1. For $m \leq 1$, stability holds for every d. For $m > 1$, the cycle is stable at every $d \in [0,1]$ when the gain floor already exceeds $1 - 1/m$, that is, when $q > r(m-1)$; otherwise, there is a critical damping d_c, $0 \leq d_c < 1$, such that the cycle is stable if and only if $d > d_c$.

(iii) Gap non-negativity. The overconfidence gap $G := \Sigma^a - p_{bel}^a$ satisfies $G \geq 0$, with equality if and only if $d = 1$.

(iv) Monotonicity. On the stable range, Σ^a and G are strictly decreasing in d, i.e., the gap grows monotonically as the damping deepens, and $G \to \infty$ as $d \downarrow d_c$ when $d_c > 0$.

Writing $\Sigma(\gamma)$ for the stationary actual variance produced by an arbitrary fixed gain γ, differentiation of the stationary identity yields

$$\Sigma'(\gamma)\,[\,1-(1-\gamma)^2 a_0\,] = 2\,(a_0\Sigma + q + r)\big(\gamma - \hat{\gamma}(\Sigma)\big), \qquad \hat{\gamma}(\Sigma) = \frac{a_0\Sigma + q}{a_0\Sigma + q + r}, \tag{13}$$

which shows that Σ is strictly decreasing in the gain all the way up to the true optimum: an emulator-suppressed gain always lies on the harmful side. In Equation (13), $\hat{\gamma}(\Sigma)$ is the optimal gain $\hat{\gamma}$ of Lemma 2 evaluated at the forecast variance $p = a_0\Sigma + q$ that an analysis variance Σ generates; it is the same object, not a new symbol.

Proof. (i) The map $T(\Sigma) = (1-\gamma)^2(a_0\Sigma + q) + \gamma^2 r$ is affine with slope $\vartheta = (1-\gamma)^2 a_0$; for $\vartheta < 1$, the Banach fixed-point theorem applies, and for $\vartheta \geq 1$, the orbit is minorized by a divergent sequence. (ii) At $d = 1$, the actual and believed contractions coincide by Theorem 1; for $m > 1$, stability is equivalent to $g > 1 - 1/m$ with g strictly increasing in d. (iii) Write $p = a_0\Sigma(g) + q$ for the stationary actual forecast variance, so that the stationary identity reads $\Sigma(g) = (1-g)^2 p + g^2 r$. Lemma 2 with $\gamma = g$ and $\hat{\gamma} = p/(p+r)$ turns this identity into $\Sigma(g) = \Phi_{\text{true}}\big(\Sigma(g)\big) + (g-\hat{\gamma})^2(p+r) \geq \Phi_{\text{true}}\big(\Sigma(g)\big)$; the inequality holds because the term dropped is a square. Here, $\Phi_{\text{true}}(x) := r(a_0 x + q)/(a_0 x + q + r)$ is the analysis map of the correctly specified filter, that is, the believed map with a replaced by a_0. Its unique fixed point p^* is the optimal stationary analysis variance, with optimal gain $\gamma^* = (a_0 p^* + q)/(a_0 p^* + q + r)$. By Theorem 1 applied with a_0 in place of a, the function $\Phi_{\text{true}}(x) - x$ is strictly concave, positive at the origin and has a single positive zero, so that $\Phi_{\text{true}}(x) > x$ for $x < p^*$. If $\Sigma(g) < p^*$ held, the two inequalities would give $\Sigma(g) \geq \Phi_{\text{true}}(\Sigma(g)) > \Sigma(g)$, a contradiction; therefore $\Sigma(g) \geq p^*$. Also $p^a_{\text{bel}} \leq p^*$, strictly for $d < 1$, by the monotonicity of Theorem 1, because the believed analysis variance at $d = 1$ is p^*. Hence $G \geq 0$ with the stated equality case. (iv) Differentiating the stationary identity in γ yields Equation (13); for $\gamma = g < \gamma^*$, item (iii) gives $\hat{\gamma}\big(\Sigma(g)\big) \geq \hat{\gamma}(p^*) = \gamma^* > g$, so $\Sigma'(g) < 0$. Since the believed gain increases with d (Theorem 1), $\Sigma^a = \Sigma(g(d))$ decreases with d, and $dp^a_{\text{bel}}/dd > 0$ completes the monotonicity of G. The denominator of the closed form vanishes at d_c while the numerator stays bounded below by $(1-g)^2 q$. □

The physical content of item (ii) is a *stability wedge*. At $d = 1$, the filter's believed contraction $a(1-g)^2 < 1$ and the actual contraction $a_0(1-g)^2 < 1$ are the same inequality, so a correctly specified filter is automatically stable; spectral bias separates the two, and for growing modes, the actual cycle can diverge while the filter's own covariances remain finite and give no internal indication of failure.

Relation to classical suboptimal-filter theory. Both actual-error configurations are per-wavenumber instances of the classical covariance (sensitivity) analysis of suboptimal filters [21,22]. Level 1, Equation (4), is the case of a suboptimal gain applied with the correct propagator, in which the true error covariance obeys a linear recursion driven by the gain of the design model. Level 2, Equations (5)–(9), is the case of an incorrect dynamics model, in which the true error is driven by the state through the mismatch term of Equation (6), as in the augmented truth–estimate formulation of the classical analysis. The following elements are therefore classical results specialized to a diagonal setting: the Level-1 recursions (4) and the Level-2 recursion (5)–(9), the stationary solution in Theorem 2(i), the gain identity of Lemma 2, the inequality $\Sigma(\gamma) \geq p^*$ for any fixed gain, and the stationary scalar Riccati solution of Equations (10) and (11) at $d = 1$. The following are new. First, the model mismatch is reduced to the single scalar $d(k)$ per wavenumber (Lemma 1), and the overconfidence gap is shown to be non-negative and strictly monotone in the damping, with the stability wedge (Theorem 2(ii)–(iv)). Second, the blind-spot threshold is obtained exactly (Theorem 4). Third, the Level-2 configuration is solved in closed form and its ordering with respect to Level 1 is proved (Theorem 3). Fourth, the ensemble collapse law is derived with its finite-ensemble correction (Theorem 5).

Theorem 3 (emulator in the loop; Level 1 is a lower bound). *Assume* $m < 1$ *and let* g *be the believed gain. The Level-2 recursion has a unique, globally attracting stationary point with*

$$\Sigma^a_{\text{L2}} = \frac{(1-g)^2 F + g^2 r}{1-(1-g)^2 a}, \qquad F = m^2(1-d)^2\, v - 2m^2 d(1-d)\, c + q, \tag{14}$$

$$v = \frac{q}{1-m^2}, \qquad c = \frac{(1-g)\,[\,m^2(d-1)\,v - q\,]}{1-(1-g)\,m^2 d} \le 0, \tag{15}$$

and $\Sigma^a_{L2} \ge \Sigma^a_{L1}$ *on the whole admissible domain* $m < 1,\ d \in [0,1],\ q, r > 0$, *with equality if and only if* $d = 1$.

Proof. *Closed form.* The cascaded (v, c, Σ) recursion of Equations (7)–(9) is upper-triangular affine with spectral radius $\max(m^2, (1-g)m^2 d, (1-g)^2 a) < 1$; solving in order gives Equations (14)–(15). Ordering, step 1: reduction to a polynomial inequality. Parametrize the believed fixed point by its gain, $q = gr\,[\,1 - a(1-g)\,]/(1-g)$, which is bijective in q for fixed (a, r) because $dq/df = 1 - a(1-g)^2 > 0$. With $\mu = m^2,\ u = 1-g,\ D_1 = 1 - (1-g)^2\mu$ and $D_2 = 1 - (1-g)^2 d^2 \mu$, the difference then takes, in lowest terms, the form of Equation (16) below, in which $Q(g, d, \mu) = 1 + \mu u + d\mu u\,(1 + 2u - \mu u) - d^2 \mu u\,(g + 2\mu u) - d^3 \mu^2 u^2 (1+u)$. Every factor in front of Q is positive for $0 < g < 1,\ 0 < \mu < 1$ and $0 \le d < 1$, so that the ordering is equivalent to $Q > 0$. Step 2: positivity certificate. The three remaining parameters $g,\ d$ and μ each range over the unit interval. The substitution $g = 1/(1+G),\ d = 1/(1+T),\ \mu = 1/(1+U)$ with $G, T, U > 0$ maps the positive orthant bijectively onto the open cube $(0,1)^3$. A polynomial inequality on the cube thereby becomes one on the orthant, where a polynomial with non-negative coefficients is manifestly non-negative; this is the motivation for the substitution. Because Q has degrees 3, 3 and 2 in g, d and μ, the product $Q(1/(1+G), 1/(1+T), 1/(1+U))\,(1+G)^3 (1+T)^3 (1+U)^2$ is a polynomial in G, T and U. It has 47 monomials, and every coefficient is a positive integer between 1 and 31. The difference itself is

$$\Sigma^a_{L2} - \Sigma^a_{L1} = \frac{r\, g^2 \mu\, u\, (1-d)^2\, Q(g, d, \mu)}{(1-\mu)(1-d\mu u)\, D_1 D_2}. \tag{16}$$

Thus, $Q > 0$ on the open cube, so that $\Sigma^a_{L2} > \Sigma^a_{L1}$ for $d < 1$; the difference vanishes only for $d = 1$, and the factor $(1-d)^2$ shows that the two configurations separate quadratically in $1-d$. The boundary $d = 0$ needs no limiting argument, since there $Q = 1 + \mu u > 0$; it is also verified by hand below. The 47 coefficients are listed, and the identity is regenerated in exact arithmetic, by the accompanying code. □

Why the ordering holds: an impulse-response reading. The certificate establishes the ordering but does not explain it. An explanation follows from writing the stationary Level-2 error as a sum over the past noise realizations. An observation error enters the analysis error with weight g and is multiplied by $u\,d\,M(k)$ at every later cycle. A process-noise increment w_i perturbs the truth, and n cycles later a fraction $P_n(d)$ of the perturbation $M(k)^n w_i$ is still absent from the analysis, where

$$P_n(d) = 1 - g \sum_{l=1}^{n} u^{l-1} d^l = \alpha + \beta\,(ud)^n, \qquad \alpha = \frac{1-d}{1-ud}, \qquad \beta = \frac{g\,d}{1-ud}. \tag{17}$$

Summing the variances of these independent contributions gives

$$\Sigma^a_{L2} = \frac{g^2 r}{1-u^2 a} + q\,u^2 \sum_{n \ge 0} m^{2n}\, P_n(d)^2, \tag{18}$$

which is algebraically identical to Equations (14)–(15). For a perfect model $P_n(1) = u^n$: the untracked fraction decays geometrically, and Equation (18) reduces to the Level-1 closed form of Theorem 2(i). For a damped emulator $P_n(d) \ge u^n$ and $P_n(d) \to \alpha > 0$: each forecast removes a fraction $1-d$ of the signal that the filter has already acquired, so that a fraction α of every truth perturbation is never tracked. Subtracting the Level-1 form at the same gain gives

$$\Sigma^a_{L2} - \Sigma^a_{L1} = q\,u^2 \sum_{n\geq 0} m^{2n}\,[P_n(d)^2 - u^{2n}] - g^2 r\left[\frac{1}{1-u^2 a_0} - \frac{1}{1-u^2 a}\right]. \qquad (19)$$

The first term is non-negative for every gain. The second is non-positive: the emulator also damps the observation noise that the analysis has absorbed. The sign of the sum therefore depends on the gain, and for an imposed gain it can be negative: with $m = 0.95$, $d = 0.3$, $q = 0.01$, $r = 5$ and $\gamma = 0.5$, $\Sigma^a_{L2} = 1.29$ against $\Sigma^a_{L1} = 1.62$. The believed gain excludes this regime, because it is small exactly when r is large relative to q; this is what the certificate proves. At $d = 0$, the ordering can be verified by hand. There, $P_n \equiv 1$ and $g = q/(q + r)$, and the difference equals $u^2 m^2 g\,[(2 - g)\,v - g\,r]/(1 - u^2 m^2)$, which is positive because $g\,r = u\,q \leq q \leq v$.

Remark 1 (Level 2 has no closed form at growing modes). *Theorem 3 is restricted to $m < 1$ because no stationary Level-2 spectrum exists otherwise. For $m \geq 1$ the climatological variance v_j of the linear truth grows without bound, the bias term of Equation (6) grows with it, and $\Sigma^a_j \to \infty$ for every $d < 1$. This is not filter divergence in the sense of Theorem 2: the error grows at the rate of the truth itself, whatever the gain. The only invariant statement is relative. For $m > 1$, $\Sigma^a_j / v_j \to \zeta^2$ with $\zeta = u\,\alpha = (1 - g)(1 - d)/[1 - (1 - g)\,d]$, so that the analysis misses a fixed fraction ζ of the growing signal however long the cycle runs (confirmed to twelve digits for m between 1.02 and 1.5). Wherever $m \geq 1$, every statement of this paper about error is therefore a Level-1 statement, that is, a statement about gain miscalibration alone.*

Remark 2 (sensitivity to the structural form of the bias term). *Three variations were examined. (a) If the emulator error contains, in addition, a state-independent white component of variance q_{em} per cycle that is unknown to the filter, F in Equation (14) increases by q_{em} and Σ^a_{L2} by $u^2 q_{\mathrm{em}}/(1 - u^2 a)$; the ordering is preserved. (b) If the damping factor is complex, $z = d\,e^{i\delta}$, then $|P_n(z)| \geq P_n(d)$ and the Level-2 error can only increase (Proposition 1). (c) If the gain is not the believed gain, the ordering can fail, as the example above shows. The ordering is thus robust to the form of the emulator error but specific to a filter that computes its gain from the covariances of the emulator.*

Two further points complete the picture. First, at $d = 1$ the two configurations coincide identically, since then $F = q$ and $a = a_0$. Second, as to evidence independent of the certificate, a random scan of 2×10^6 parameter draws over the admissible domain found no violation beyond 10^{-15} relative round-off. Figures 2 and 3 display the believed and actual spectra and the gap for the toy model.

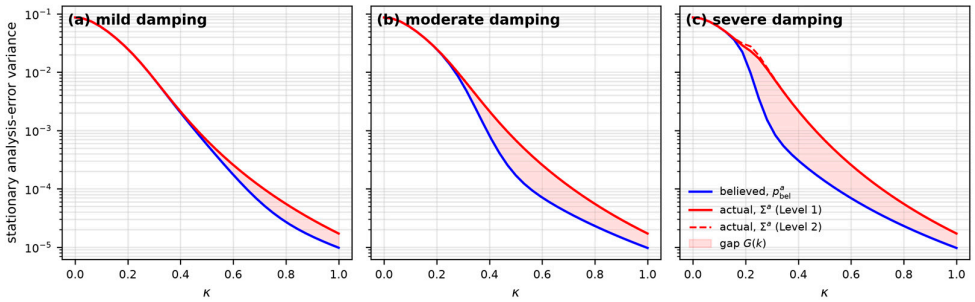


Figure 2. Believed (p^a_{bel}, blue) and actual (Σ^a, red; Level 2 dashed) stationary analysis-error spectra of the spectral toy model for the (**a**) mild, (**b**) moderate and (**c**) severe damping profiles of Figure 1. The shaded area is the overconfidence gap $G(k)$: the filter under-reports its own error exactly where the emulator is damped. In (a,b), the Level-2 curve coincides with the Level-1 curve to within the line width.

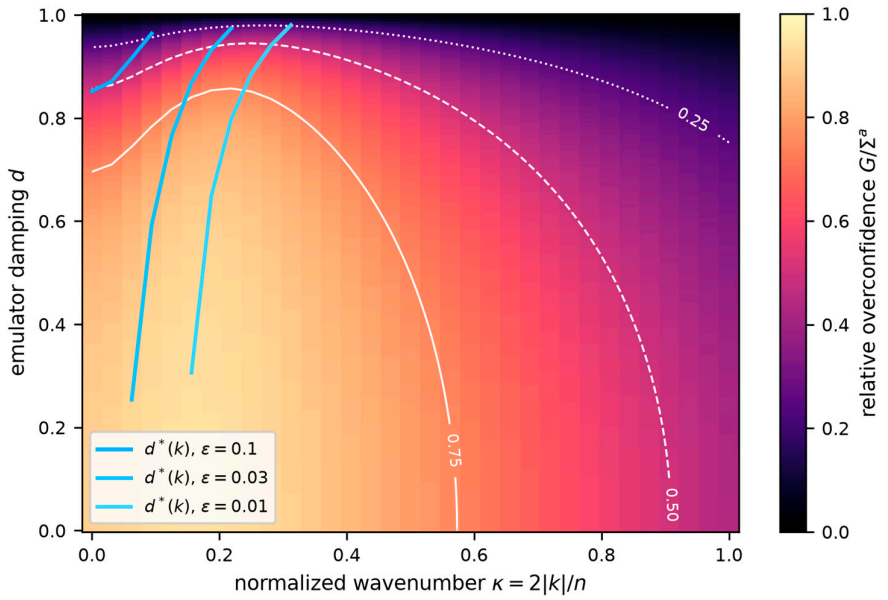


Figure 3. Relative overconfidence G/Σ^a, the fraction of the actual analysis-error variance that is invisible to the filter, over the plane of normalized wavenumber κ and uniform damping d for the stable toy configuration. Cyan curves are the blind-spot thresholds $d^*(k)$ of Theorem 4 for tolerances $\varepsilon = 0.1,\ 0.03$ and 0.01 (shown only where strictly inside $(0,1)$). White contours mark overconfidence fractions of 0.25 (dotted), 0.50 (dashed) and 0.75 (solid). The gap vanishes at $d = 1$ and grows monotonically as the damping deepens, approaching unity at the largest scales; at small scales it saturates below unity, because $m(k)$ is itself small there and the uncorrected error is correspondingly low.

3.3. Spectral Blind Spots

A wavenumber is termed a *spectral blind spot* at tolerance $\varepsilon \in (0,1)$ if the stationary gain satisfies $g < \varepsilon$: observations are assimilated every cycle, but with a weight below the prescribed tolerance.

Theorem 4 (Blind-spot threshold). *Fix $\varepsilon \in (0,1)$ and let $\varphi := \varepsilon r/(1-\varepsilon)$, the believed forecast variance at which $g = \varepsilon$. For $\varphi > q$, define*

$$a^*(\varepsilon) = \frac{(\varphi - q)(\varphi + r)}{\varphi\, r}, \qquad d^*(\varepsilon) = \frac{\sqrt{a^*(\varepsilon)}}{m}. \tag{20}$$

Then, exhaustively:

(i) If $\varphi \le q$, the gain floor keeps $g \ge \varepsilon$ at every d, and there is no blind spot at this wavenumber;

(ii) If $\varphi > q$ and $a^* > m^2$, then $g < \varepsilon$ for all *$d \in [0,1]$: the scale is under-observed even with a perfect model;*

(iii) Otherwise, $d^* \in (0,1]$ and $g(d) < \varepsilon$ exactly for $d < d^*$.

Furthermore, for $m < 1$ and every d, the stationary error at the believed gain improves on the no-assimilation error $\Sigma_{free} = q/(1-m^2)$ by the relative amount

$$0 \le \frac{\Sigma_{\text{free}} - \Sigma(g)}{\Sigma_{\text{free}}} = \frac{g\,[\,2 - g - (1-m^2)\,g\,r/q\,]}{1-(1-g)^2 m^2} \le \frac{g\,(2-g)}{1-(1-g)^2 m^2} \le \frac{2g}{1-m^2}. \tag{21}$$

The bound $g(2-g)/[1-(1-g)^2m^2]$ increases with g, so that within a blind spot the relative improvement is below $\varepsilon(2-\varepsilon)/[1-(1-\varepsilon)^2m^2] \le 2\varepsilon/(1-m^2)$. For $m > 1$, $\Sigma^a = \infty$ if and only if $g \le 1 - 1/m$, i.e., if and only if $d \le d_c = d^*(1-1/m)$ when $q \le r(m-1)$ (for $q >$

$r(m-1)$ *the mode is stable for every* d*): below the threshold, observations are assimilated at every cycle yet cannot stabilize the mode.*

Proof. The gain $g(d)$ is continuous and strictly increasing (Theorem 1), and $g \leftrightarrow f$ is a bijection with $g = \varepsilon$ if and only if $f = \varphi$. Solving the stationary quadratic for a at $f = \varphi$ gives $a^* = (\varphi - q)(\varphi + r)/(\varphi r)$, which requires $\varphi > q$; the discriminant at a^* is the perfect square $[\,(qr + \varphi^2)/\varphi\,]^2$, so $f(a^*) = \varphi$ identically. Cases (i)–(iii) follow from the gain floor and from $a \le m^2$. For Equation (21), subtracting the closed form of Theorem 2(i) from Σ_{free} gives $\Sigma_{\text{free}} - \Sigma(g) = g\,[\,q(2-g) - (1-m^2)\,g\,r\,]/\{(1-m^2)[1-(1-g)^2m^2]\}$, which is the stated identity. The upper bounds follow on dropping the non-positive term and from $1-(1-g)^2m^2 \ge 1-m^2$. The same subtraction holds for any fixed gain in place of g, so that the bracket is non-negative exactly for gains up to $\gamma_0 = 2q/[q + (1-m^2)r]$. Since $\Sigma(\gamma^*) = p^* \le \Sigma(0) = \Sigma_{\text{free}}$, the optimum satisfies $\gamma^* \le \gamma_0$, and $g \le \gamma^*$ (Theorem 1) gives the lower bound. The derivative of $g(2-g)/[1-(1-g)^2m^2]$ has the sign of $(1-g)(1-m^2) > 0$. The divergence branch is Theorem 2(i) at $\varepsilon = 1 - 1/m$, i.e., $\varphi = r(m-1)$. □

Equation (21) shows that the tolerance alone guarantees nearly ineffective assimilation in a blind spot only where ε is small compared with the decay rate $1-m^2$ of the mode. Whatever ε, the improvement is small wherever the gain is small compared with $1-m^2$, and it remains appreciable where the gain, although below ε, is comparable with $1-m^2$. In the typical mid-scale regime of Table 2, where every $d < 0.85$ is blind at $\varepsilon = 0.05$, the reduction is 28% at $d = 0$, 33% at $d = 0.5$ and 46% just below d^*. Figure 4 shows the gain suppression and the threshold for the toy model, and Table 2 lists representative regimes spanning the three cases. The divergence branch is confirmed exactly in Section 3.5: in the growing toy configuration with severe damping the theory flags eight Fourier modes as divergent, corresponding to four distinct wavenumber magnitudes together with the conjugate wavenumbers, and the Level-1 matrix recursion, followed for 300 cycles, grows without bound at exactly those eight.

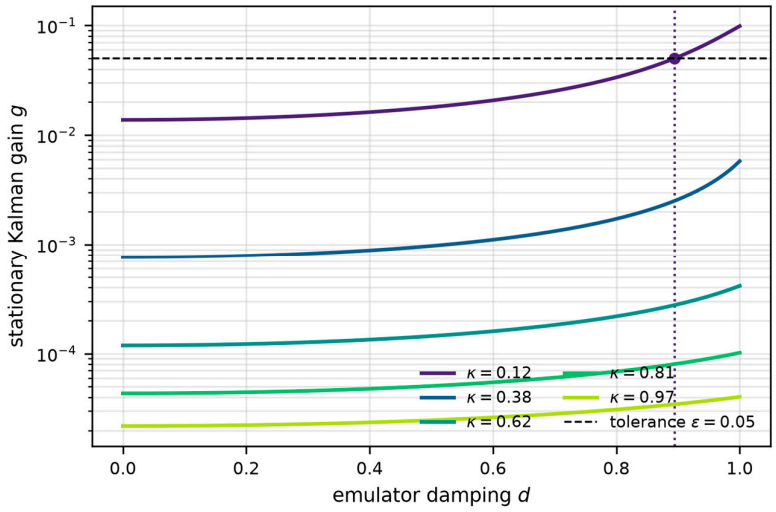


Figure 4. Stationary Kalman gain g as a function of the damping d at five wavenumbers of the stable toy configuration (log scale). The dashed line is the tolerance $\varepsilon = 0.05$ and the dotted vertical line marks the threshold d^* of the only wavenumber shown for which case (iii) of Theorem 4 applies; the remaining curves lie below ε for every d (case ii). All curves approach the floor $q/(q+r)$ as $d \to 0$.

Table 2. Blind-spot thresholds d^* at tolerance $\varepsilon = 0.05$ for representative regimes, illustrating the three cases of Theorem 4; $g(d = 1)$ is the gain of the perfect-model filter.

Regime	m	q	r	d^*	Interpretation	$g(d = 1)$
well-observed growing scale	1.05	10^{-2}	10^{-1}	–	no blind spot (case i)	0.302
well-observed neutral scale	1.00	10^{-2}	10^{-1}	–	no blind spot (case i)	0.270
typical mid-scale	0.95	10^{-2}	0.5	0.850	blind for $d < 0.850$ (case iii)	0.097
small scale, small process noise	0.80	10^{-4}	0.5	1	blind at all d (case ii)	5.5×10^{-4}
small scale, tiny process noise	0.70	10^{-6}	0.5	1	blind at all d (case ii)	3.9×10^{-6}
noisy observations, small scale	0.80	10^{-4}	5	1	blind at all d (case ii)	5.6×10^{-5}
strong process noise	0.90	1	0.5	–	no blind spot (case i)	0.721
unstable, tiny process noise	1.20	10^{-4}	0.5	0.853	blind for $d < 0.853$; here $d_c = 0.912 > d^*$, so the mode also diverges for $d \leq 0.912$	0.306

3.4. Ensemble Spread Collapse

Ensemble filters, in which the forecast covariance is estimated from a finite sample of propagated states [23], are considered next. Consider a deterministic square-root ensemble Kalman filter (EnSRF) [24] acting per wavenumber (mode-wise gains, Assumption A10), with N members whose perturbations are propagated by the emulator and with additive model-error variance $q_{\mathrm{ens}} \geq 0$ injected per cycle. Let $\hat{s}^f$ denote the ensemble sample variance of the forecast perturbations at the considered mode; $\hat{g} = \hat{s}^f/(\hat{s}^f + r)$, the corresponding sample gain; and $\psi(s) := a\,sr/(s + r) + q_{\mathrm{ens}}$, the noise-free spread map (one propagation followed by one update).

Theorem 5 (spread collapse and finite-ensemble bias).

(i) Exact update law (any $N \geq 2$). The sample spread obeys $\hat{s}^a = (1 - \hat{g})\,\hat{s}^f$ exactly, and $\hat{s}^f_{j+1} = a\,\hat{s}^a_j + \hat{q}_{ens,j}$, where $\hat{q}_{ens,j}$ is the increase in the sample variance caused by the injected perturbations, including the sample covariance between injected and propagated perturbations; it equals q_{ens} when the variance is injected exactly.

(ii) Collapse dichotomy ($q_{ens} = 0$). If $a = d^2m^2 \leq 1$, the spread decreases monotonically to zero with $s_{j+1} \leq a\,s_j$ and per-cycle ratio $s_{j+1}/s_j = ar/(s_j + r)$ increasing to exactly a: for $a < 1$ a geometric collapse at asymptotic rate d^2m^2, and for $a = 1$ an algebraic decay, $s_j = r\,s_0/(r + j\,s_0)$. If $a > 1$, the spread equilibrates at $s^+ = r(a - 1)$. Spread survives at mode k if and only if $d(k)m(k) > 1$ or $q_{ens} > 0$.

(iii) Finite-N bias of a freshly drawn ensemble. With $\hat{s} = sX/\nu$ and $X \sim \chi^2_\nu$, where ν is the number of real degrees of freedom of the per-mode sample variance, $\nu = 2(N - 1)$ at a complex Fourier mode $0 < |k| < n/2$ and $\nu = N - 1$ at the real modes $k = 0$ and $k = n/2$, the expected post-update spread satisfies

$$\mathbb{E}[\hat{s}^a] = h(s)\left[1 - \frac{2\hat{g}(1 - \hat{g})}{\nu}\right] + O\left(\nu^{-3/2}\right), \qquad h(s) = \frac{sr}{s + r}, \tag{22}$$

and, for $q_{ens} = 0$, $\mathbb{E}[s_{j+1}] \leq a\,\mathbb{E}[s_j]$. Sampling error therefore lowers the expected post-update spread of a freshly drawn ensemble below its infinite-ensemble value, most strongly at $\hat{g} = 1/2$, where the forecast spread equals the observation-error variance. With $q_{ens} = 0$, this deficit is

incurred once in the cycled deterministic filter, when the ensemble is drawn: by item (i), the realized spread then follows the map ψ exactly, and at modes with $a \le 1$, the per-cycle ratio tends to a for every N.

The expectation is also known exactly. With $x = \nu r/(2s)$ and E_p the generalized exponential integral [25],

$$\mathbb{E}[\hat{s}^a] = r\left[1 - x\, e^x E_{\nu/2}(x)\right], \qquad E_p(x) = \int_1^\infty e^{-xt}\, t^{-p}\, dt, \tag{23}$$

so that the remainder in Equation (22) is in fact $O(\nu^{-2})$, with leading term $4\hat{g}^2(1-\hat{g})(2-3\hat{g})\, h(s)/\nu^2$.

(iv) General bound without Assumption A10. Consider any deterministic square-root ensemble filter whose analysis sample covariance satisfies $\hat{P}^a = (I - \hat{K}H)\hat{P}^f$ with the sample gain $\hat{K} = \hat{P}^f H^T (H\hat{P}^f H^T + R)^{-1}$, for any $N \ge 2$, any linear observation operator H and any positive-definite R, without localization. Let the emulator be diagonal in the Fourier basis, let $q_{ens} = 0$, and let $\lambda \ge 1$ be a multiplicative inflation factor applied to the perturbations ($\lambda = 1$ without inflation). Then, the Fourier-diagonal sample spread obeys, for every wavenumber and every cycle,

$$\hat{s}^f_{j+1}(k) \le \lambda^2 a(k)\, \hat{s}^f_j(k). \tag{24}$$

Spread at wavenumber k therefore collapses at least geometrically wherever $\lambda\, d(k)\, m(k) < 1$, whatever the ensemble size and the observing network.

(v) Inflated mode-wise filter. Under A10 with multiplicative inflation λ and $q_{ens} = 0$, items (i) and (ii) hold with a replaced by $\lambda^2 a$: spread survives if and only if $\lambda\, d\, m > 1$, and then equilibrates at the forecast value $r(\lambda^2 a - 1)$.

Proof. (i) The scalar EnSRF scales the perturbations by $1 - \alpha\hat{g}$ with $\alpha = \left[1 + \sqrt{r/(\hat{s}^f + r)}\right]^{-1}$, and $(1-\alpha\hat{g})^2 = 1 - \hat{g}$ identically, so the sample variance obeys $\hat{s}^a = (1-\hat{g})\hat{s}^f$ exactly; propagation is linear. (ii) $\psi(s) = asr/(s+r)$ is increasing and concave with $\psi(s) < as$; for $a \le 1$, orbits decrease to zero with ratio $ar/(s_j + r)$ increasing to a; for $a > 1$, the interior fixed point $r(a-1)$ attracts $(0, \infty)$ by the standard monotone-map argument. (iii) Under Gaussianity, $\nu\hat{s}/s \sim \chi^2_\nu$ exactly, with the degrees of freedom of item (iii), since the real and imaginary parts of a complex mode are independent with variance $s/2$ each. A Taylor expansion with $\mathbb{E}[\hat{s} - s] = 0$, $\mathrm{Var}(\hat{s}) = 2s^2/\nu$, $|h'''| \le 6/r^2$ and $\mathbb{E}|\hat{s} - s|^3 = O(\nu^{-3/2})$ gives Equation (22), since $1/2\, h''(s) \cdot 2s^2/\nu = -h(s)\, 2\hat{g}(1-\hat{g})/\nu$; Jensen's inequality for the strictly concave h gives the sign, and for $q_{\mathrm{ens}} = 0$ the inequality $h(s) \le s$ gives $\mathbb{E}[s_{j+1}] \le a\, \mathbb{E}[s_j]$. For the exact expectation, $Y = X/2$ is Gamma distributed with shape $\nu/2$, $h(\hat{s}) = r - r\,(1 + Y/x)^{-1}$, and $\mathbb{E}[(1+Y/x)^{-1}] = x \int_0^\infty e^{-xt}\,(1+t)^{-\nu/2} dt = x\, e^x E_{\nu/2}(x)$; expanding in $1/\nu$ gives the stated remainder. (iv) The difference $\hat{P}^f - \hat{P}^a = \hat{P}^f H^T (H\hat{P}^f H^T + R)^{-1} H\hat{P}^f$ is positive semi-definite, so $\hat{P}^a \preceq \hat{P}^f$ in the Loewner order, and every diagonal element in any orthonormal basis, in particular the Fourier basis, satisfies $\hat{s}^a(k) \le \hat{s}^f(k)$. Deterministic square-root filters, including the serial EnSRF, share this analysis covariance [24,26]. Propagation by a Fourier-diagonal emulator followed by inflation multiplies the k-th diagonal element by exactly $\lambda^2 a(k)$. (v) This is item (ii) for the map $s \mapsto \lambda^2 a\, sr/(s+r)$. □

The scope of Theorem 5 is as follows. Items (i)–(iii) and (v) are exact laws for the mode-wise (spectral) filter of Assumption A10, and item (iv) is a rigorous bound for unlocalized gridpoint square-root filters. Neither is an operational law. Operational ensemble filters are localized, inflated, nonlinear and sparsely observed, and for those the theorem is a statement of mechanism, with spectral bias implying the spread collapse at a rate

of d^2m^2, and not a quantitative prediction of the spread. Section 3.10 maps where the exact law ceases to extrapolate. In brief, an unlocalized gridpoint serial EnSRF without the mode-wise idealization follows the infinite-ensemble map at moderately damped modes ($0.2 < a < 0.95$) to within 2% when the ensemble is about three times larger than the state dimension: the median ratio of spread to map after 20 cycles is 0.98 at $N = 200$ on the 64-point grid. At $N = 20$, that ratio is 0.49, so the spread has run about a factor of two below the infinite-ensemble prediction after 20 cycles; it continues to fall, to 0.45 after 40 cycles. This deficit is not the sampling bias of item (iii): the map is started from the realized initial spread, against which a mode-wise filter keeps a ratio of exactly one for any N (item (i)). It arises from spurious cross-wavenumber sample covariances of a finite ensemble, which Assumption A10 removes and item (iv) only bounds (Section 3.10). With covariance localization [27–29] of suitable width the agreement at moderately damped modes returns, but the per-mode bound of item (iv) no longer holds, because localization transfers sample variance between wavenumbers. Each deeply damped mode of the unlocalized filter continues to collapse at its own rate; an apparent floor in double-precision gridpoint arithmetic is round-off and not a sampling effect (Section 3.10).

At $d = 1$, the same map collapses at rate m^2 at every decaying mode: without injected variance, a deterministic ensemble loses spread at decaying scales for any model. The emulator-specific content of item (ii) is the factor d^2 in the rate and the displacement of the survival boundary from $m(k) > 1$ to $d(k)\,m(k) > 1$, which turns the modes with $1 < m(k) < 1/d(k)$ from surviving into collapsing. Combined with Theorem 4, item (ii) is the quantitative form of filter divergence by overconfidence [15,24,27,30]: the spread and hence the gain tend to zero, while the actual error tends to the free-run error for $m < 1$ or diverges for $m > 1$. Figure 5 shows the cycled EnSRF following the exact map to 10^{-13} and the sustained spread matching $r(a-1)$ to 3.2×10^{-15} (Table 3). The χ^2-quadrature check of the sampling coefficient at $\nu = 1279$ gives -0.4998 against the predicted value $-2\hat{g}(1-\hat{g}) = -0.5$ at $s = r$. Table 4 compares the exact expectation of Equation (23) with the first-order form of Equation (22) at complex modes, $\nu = 2(N-1)$: the remainder is below 0.5% of $h(s)$ at $N = 5$ and below 0.03% for $N \geq 20$, at every gain. A rigorous but loose bound on the absolute remainder follows from the Lagrange form of the Taylor expansion: the remainder does not exceed $12^{3/4}(1+4/\nu)^{3/4}s^3r^{-2}\nu^{-3/2}$ in modulus.

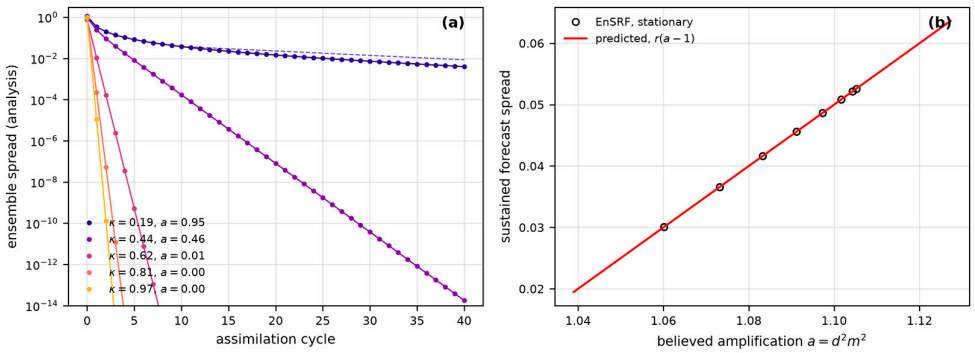


Figure 5. (**a**) Analysis spread of the mode-wise (spectral) EnSRF of Assumption A10 versus assimilation cycle (dots; $N = 20$; the first 40 of 60 cycles are shown) at five wavenumbers of the stable toy configuration, together with the exact spread map (solid) and the asymptotic geometric rate $a^j = (dm)^{2j}$ (dashed). (**b**) Where $a = d^2m^2 > 1$ (growing configuration), the stationary forecast spread of the ensemble after 2000 cycles (circles; $N = 20$) lies on the predicted line $r(a-1)$.

Table 3. Ensemble spread-collapse rates: predicted, $a = d^2m^2$, versus measured in the cycled EnSRF ($N = 20$, stable toy configuration) after correction for the finite-spread factor $r/(as + r)$. An entry $<10^{-16}$ denotes agreement in every digit of double-precision arithmetic.

κ	Predicted a	Measured Asymptotic Rate	Relative Deviation
0.31	8.646194×10^{-1}	8.646194×10^{-1}	1×10^{-16}
0.44	4.646791×10^{-1}	4.646791×10^{-1}	3×10^{-16}
0.56	5.959801×10^{-2}	5.959801×10^{-2}	3×10^{-16}
0.81	2.326878×10^{-4}	2.326878×10^{-4}	$<10^{-16}$
0.97	1.146996×10^{-5}	1.146996×10^{-5}	$<10^{-16}$

Table 4. Finite-ensemble sampling factor $\mathbb{E}[\hat{s}^a]/h(s)$ from the exact expression (23), and remainder R left by the first-order form $1 - 2\hat{g}(1 - \hat{g})/\nu$ of Equation (22), for three values of the gain $\hat{g} = s/(s + r)$, at a complex Fourier mode, where an ensemble of N members has $\nu = 2(N - 1)$ degrees of freedom.

N	Exact, $\hat{g} = 0.1$	R	Exact, $\hat{g} = 0.5$	R	Exact, $\hat{g} = 0.9$	R
5	0.9784	9.0×10^{-4}	0.9415	4.0×10^{-3}	0.9740	-3.5×10^{-3}
10	0.9902	1.8×10^{-4}	0.9730	7.9×10^{-4}	0.9893	-7.1×10^{-4}
20	0.9953	4.2×10^{-5}	0.9870	1.8×10^{-4}	0.9951	-1.6×10^{-4}
40	0.9977	1.0×10^{-5}	0.9936	4.1×10^{-5}	0.9977	-3.7×10^{-5}
80	0.9989	2.4×10^{-6}	0.9968	1.0×10^{-5}	0.9989	-9.1×10^{-6}

3.5. Numerical Verification in a Spectral Toy Model

The toy model is a periodic advection–diffusion propagator on $n = 64$ grid points with symbol $M(k) = \exp(\lambda - \nu_d\kappa^2)\, e^{-ic\kappa}$, where $\kappa = 2|k|/n \in [0,1]$ is the normalized wavenumber, $\nu_d = 0.4067$, $c = 1.5$, and either weak large-scale damping ($\lambda = -0.01$, the stable configuration, $m \leq 0.99$) or growth ($\lambda = 0.05$, the growing configuration, $m > 1$ at large scales). The process-noise spectrum is red, $q(k) = q_0\,[\,1 + (\kappa/0.15)^2\,]^{-2}$ with $q_0 = 0.02$, the observation-error spectrum is white, $r = 0.5$, and the damping profiles are sigmoids $d(k) = [\,1 + (\kappa/\kappa_c)^8\,]^{-1}$ with $\kappa_c = 0.8$, 0.5 and 0.3 (mild, moderate and severe; Figure 1). The grid is deliberately small. By Lemma 1, the wavenumbers are exactly decoupled, so the grid size determines only how many independent scalar recursions are sampled, not the content of any one recursion; $n = 64$ provides 33 distinct wavenumber magnitudes over the full range of κ and keeps the dense matrix recursions inexpensive. Repeating the believed and Level-1 matrix recursions at $n = 256$ (stable configuration, moderate damping) gives agreement to 1.1×10^{-11} and 1.7×10^{-11}, against 9.2×10^{-12} and 1.5×10^{-11} at $n = 64$. Three independent implementations confirm the theory. The full 64×64 matrix recursions in gridpoint space test the diagonalization of Lemma 1 itself and not merely the scalar algebra. Iterated to stationarity, these recursions match the closed forms of Theorems 1 and 2 to within 4.2×10^{-11} and 5.6×10^{-11} at every wavenumber, for Theorem 1 in all six configurations and for Theorem 2 in the five configurations without divergent modes (at most 1.7×10^{-11} in the three strictly stable ones). The Level-2 form of Theorem 3, which is defined only for $m < 1$, is matched to within 1.3×10^{-12} in the three stable configurations. A stochastic Monte Carlo simulation of the Level-2 cycle with 2000 independent realizations over 400 stationary cycles matches Theorem 3 to within 1.3% (maximum over the three damping profiles), consistent with its sampling noise. The growing configuration with severe damping, in which the theory predicts a divergence band, cannot be iterated to stationarity. There the Level-1 matrix recursion was followed for 300 cycles and compared with the scalar recursion at every cycle. The two agree at all 64 wavenumbers to within 2×10^{-12} of the largest variance,

the one-cycle amplification equals $(1-g)^2 m^2$ to within 7×10^{-13}, and the variance grows without bound at exactly the eight predicted modes, with amplification factors between 1.0017 and 1.032. Figure 6 overlays all three against the closed forms.

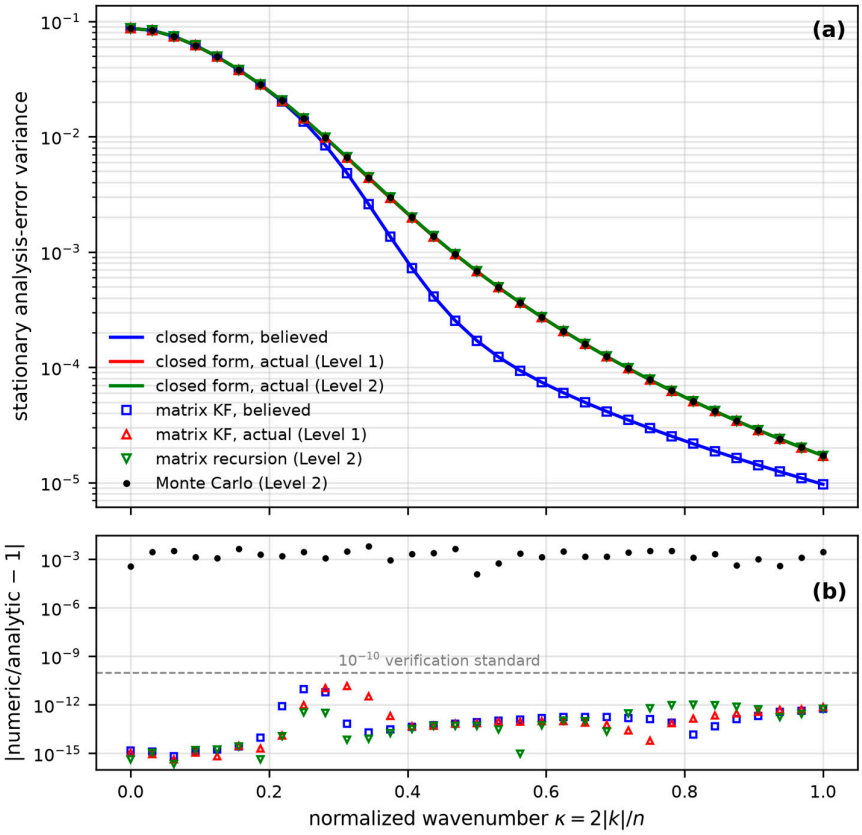


Figure 6. (**a**) Stationary analysis-error spectra of the spectral toy model for the stable configuration with moderate damping, the case for which all three numerical paths are defined: closed forms (lines), full matrix Kalman-filter recursions in gridpoint space (open markers) and stochastic Monte Carlo cycling (dots). (**b**) Relative deviations of the numerical results from the closed forms. The deterministic checks lie between 10^{-15} and 10^{-11}, below the 10^{-10} verification standard (dashed line); the Monte Carlo points scatter at the 10^{-3} sampling level.

3.6. Nonlinear Illustration: Lorenz-96

As motivation, and explicitly not as a test of the theorems, a serial EnSRF ($N = 20$, every variable observed at every step, $\sigma_{\mathrm{obs}} = 1$, multiplicative inflation 1.05) is cycled on the Lorenz-96 model [31] in the standard configuration of Lorenz and Emanuel [32], with $n = 40$ variables, forcing $F = 8$ and a fourth-order Runge–Kutta step of 0.05 time units. The forecast step is not a trained emulator. It applies the true integrator followed by a prescribed partial spectral damping $d(k) = d_{\min} + (1 - d_{\min})\,[\,1 + (\kappa/0.5)^8\,]^{-1}$ with $d_{\min} = 0.7$, i.e., a 30% amplitude loss at small scales. It is thus a controlled mimic of the amplitude loss of a smoothed emulator, without the phase errors, scale transfer and state dependence of a real one. The observation-error standard deviation of one unit is 27% of the climatological standard deviation of the model (3.64), so that the observations are informative without being dominant. The inflation factor 1.05 is the smallest of the tested values for which the true-model filter is stable at all three tested observation-error levels (Table 5). Three regimes are compared over 4000 cycles after a 500-cycle spin-up (Figure 7). The baseline uses the true model (root mean square error, RMSE, 0.209). The collapse

regime uses the damped forecast step without model-error injection (RMSE 3.65). The mitigated regime adds noise of standard deviation $\sigma_{\mathrm{add}} = 0.2$ to the damped forecast step, the $q_{\mathrm{ens}} > 0$ remedy of Theorem 5 (RMSE 2.09). In the collapse regime, the cycle behaves as Theorems 4 and 5 predict qualitatively. Diagnostics quoted here for the damped scales are medians over the wavenumbers with $d(k) < 0.75$, that is, those appreciably affected by the damping. There, the spread-to-error ratio falls from 1.33 to 0.0003, the gain vanishes, and the analysis-error variance inflates by a factor of about 220 relative to the baseline; in the mitigated regime, the inflation factor is about 130 and the spread-to-error ratio 0.013.

Sensitivity to inflation and observation error. The experiment was repeated for multiplicative inflation factors 1.00, 1.02, 1.05, 1.10 and 1.20 and observation-error standard deviations 0.5, 1 and 2 (Table 5). Without inflation, the true-model filter itself diverges (RMSE about 4), and at inflation 1.02, it fails for $\sigma_{\mathrm{obs}} = 2$; these four settings provide no baseline. In each of the remaining 11 settings, the qualitative result is unchanged. The collapse regime raises the RMSE by a factor between 3.2 and 47 relative to the baseline, the spread-to-error ratio at the damped scales stays below 0.02 in the collapse regime, and model-error injection lowers the RMSE. Larger inflation mitigates the collapse (RMSE 3.65 at inflation 1.05 against 2.18 at 1.20 for $\sigma_{\mathrm{obs}} = 1$) but does not remove it. This is consistent with the survival criterion $\lambda\, d\, m > 1$ of Theorem 5(v), which at the fully damped scales ($d = 0.7$) requires $\lambda\, m > 1.43$.

Table 5. Lorenz-96 experiment: sensitivity of the analysis RMSE and of the spread-to-error ratio at the damped scales (median over the wavenumbers with $d(k) < 0.75$) to the multiplicative inflation factor and to the observation-error standard deviation σ_{obs}. The setting of Figure 7 is inflation 1.05 and $\sigma_{\mathrm{obs}} = 1$; a dagger marks a setting in which the true-model baseline itself fails.

Inflation	σ_{obs}	RMSE, Baseline	RMSE, Collapse	RMSE, Mitigated	Spread/Error: Collapse	Spread/Error: Mitigated
1.02	0.5	0.086	4.06	1.68	1.2×10^{-5}	0.011
1.02	1	0.187	4.08	2.16	4.8×10^{-5}	0.012
1.02	2	2.07†	4.16	2.59	1.9×10^{-4}	0.014
1.05	0.5	0.100	3.72	1.65	7.6×10^{-5}	0.012
1.05	1	0.209	3.65	2.09	3.2×10^{-4}	0.013
1.05	2	0.471	3.70	2.47	1.3×10^{-3}	0.017
1.10	0.5	0.122	2.74	1.60	4.0×10^{-4}	0.013
1.10	1	0.250	2.72	2.00	1.6×10^{-3}	0.015
1.10	2	0.547	2.71	2.30	6.5×10^{-3}	0.022
1.20	0.5	0.156	2.16	1.53	1.2×10^{-3}	0.015
1.20	1	0.318	2.18	1.87	4.8×10^{-3}	0.020
1.20	2	0.685	2.21	2.10	1.9×10^{-2}	0.034

The nonlinear result does not validate the per-scale formulae; it shows the limit of the linear picture. In the linear theory, the blind spot is confined to the damped wavenumbers; in Lorenz-96, it is not: nonlinear coupling exports the damage across the spectrum, and the error peaks at the energy-containing scales. Wavenumber-localized blind spots are therefore a linear-theory statement and should not be read as a nonlinear prediction; what survives nonlinearity is the mechanism (bias, spread collapse, gain suppression, uncorrected error) together with the remedy. The linear theory predicts no degradation at all at undamped wavenumbers, where $G = 0$ whenever $d = 1$ (Theorem 2), whereas in the nonlinear cycle precisely those energy-containing scales degrade most (Figure 7b). The linear analysis is therefore neither an upper nor a per-scale lower bound for a nonlinear cycle. It locates where the mechanism originates and quantifies the

consequence within the linear setting, whereas the total degradation of the nonlinear cycle is larger and differently distributed. The only bound proved here that relates the two deployment configurations, the Level-1 lower bound of Theorem 3, holds within the linear theory and is not asserted for a nonlinear system. No per-scale linear prediction is evaluated for Lorenz-96, and none should be read into the comparison.

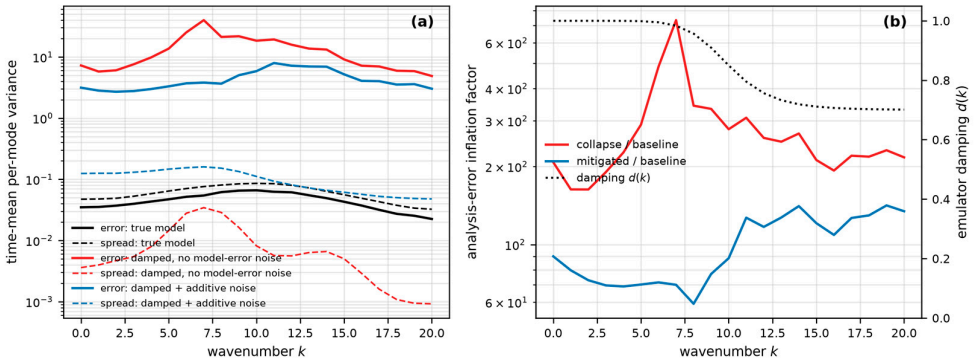


Figure 7. Lorenz-96 EnSRF with a spectrally damped forecast step ($N = 20$, $\sigma_{\mathrm{obs}} = 1$, multiplicative inflation 1.05; statistics over 4000 cycles). (**a**) Time-mean analysis-error (solid) and ensemble-spread (dashed) spectra for the true model (black), the damped forecast step without model-error injection (red) and the damped forecast step with additive model-error noise (blue). (**b**) Analysis-error inflation relative to the baseline for the two damped regimes, together with the damping profile $d(k)$ (dotted, right axis). Time- and domain-mean analysis RMSE of the three regimes: 0.209 (baseline, true model), 3.65 (collapse, damped forecast step) and 2.09 (mitigated, damped forecast step with additive noise).

3.7. Beyond the Real-Damping Model: Phase Errors, Scale Transfer and Time Dependence

The emulator model (2), a real multiplier that is fixed in time, is the most restrictive assumption of the theory. This subsection quantifies what changes when it is relaxed in three directions and gives a diagnostic for testing whether the baseline is adequate for a given emulator. The results are illustrated in Figure 8.

Proposition 1 (complex damping factor). *Let the emulator symbol be* $z(k)\,M(k)$ *with* $z = d\,e^{i\delta}$, $d \in [0,1]$ *and phase error* $\delta \in (-\pi, \pi]$.

(i) Theorems 1, 2, 4 and 5 hold verbatim with $d = |z|$*: the believed spectrum, the gain, the blind-spot threshold, the Level-1 error and the ensemble spread do not depend on* δ.

(ii) For $m < 1$*, the Level-2 stationary error is given by Equation (14) with*

$$F = m^2|z-1|^2\,v + 2m^2\,\mathrm{Re}\left[z\,\overline{(z-1)\,c}\right] + q, \qquad c = \frac{(1-g)[m^2(z-1)\,v - q]}{1-(1-g)\,m^2 z}, \tag{25}$$

where $|z-1|^2 = (1-d)^2 + 4d\sin^2(\delta/2)$.

(iii) For every gain $\gamma \in [0,1]$, $\Sigma_{L2}(\gamma; d, \delta) \geq \Sigma_{L2}(\gamma; d, 0)$. *For* $\gamma \in (0,1)$*, an interval that contains the believed gain, equality holds if and only if* $\delta = 0$ *or* $d = 0$*; at* $\gamma = 0$*, the error equals the climatological variance* v*, and at* $\gamma = 1$ *it equals* r*, whatever the phase. In particular, the ordering of Theorem 3 persists for every phase error.*

Proof. (i) By Lemma 1, the believed recursion, the Level-1 recursion and the spread map involve the emulator only through $|z\,M(k)|^2 = d^2m^2$. (ii) Equation (5) holds with d replaced by z, and the second moments close as in Equations (7)–(9). (iii) With g the applied gain and $u = 1 - g$, Equation (18) holds with $P_n(d)^2$ replaced by $|P_n(z)|^2$, where $P_n(z) = 1 - g\sum_{l=1}^{n} u^{l-1} z^l$, while the observation-noise term depends on $|z|$ only. The

triangle inequality gives $|P_n(z)| \geq 1 - g \sum_{l=1}^{n} u^{l-1} d^l = P_n(d)$, strictly for $n = 1$ unless $\delta = 0$, $d = 0$ or $g = 0$; the term $n = 1$ enters Equation (18) with the weight $q\,u^2 m^2$, which is positive for $g < 1$. □

The consequence is an asymmetry. A phase error changes neither the covariances that the filter carries nor the gain nor the Level-1 error, so the internal uncertainty estimate of the filter gives no indication of it. Yet it increases the emulator-in-the-loop error, because the bias term of Equation (6) involves $|z-1|^2$, to which the phase contributes $4d\sin^2(\delta/2)$. For $m = 0.95$, $q = 0.01$ and $r = 0.5$, a pure phase error ($d = 1$) of 0.05π raises Σ_{L2}^{a} by a factor of 1.47, one of 0.1π by 1.89, and the factor saturates at 2.26; at $d = 0.9$, the corresponding factors are 1.15, 1.34 and 1.59 (Figure 8a). The effect weakens as the damping deepens, since a mode that is already removed has no phase left to corrupt. The real-damping model is therefore adequate at Level 2 wherever $4d\sin^2(\delta/2) \ll (1-d)^2$, a condition that is most demanding where the emulator is most accurate in amplitude.

Proposition 2 (first-order invariance under scale transfer). *Let the emulator be* $\widetilde{M} = (D + \eta B)M$, *where* B *is a bounded operator whose Fourier representation has zero diagonal, so that* ηB *carries the part of the emulator error that transfers variance between wavenumbers. Assume that the stationary cycle at* $\eta = 0$ *is stable and, for the Level-2 covariance, that* $m(k) < 1$ *at every wavenumber. Then, the Fourier-diagonal parts of the stationary believed, Level-1 and Level-2 covariances are differentiable in* η *at* $\eta = 0$ *with zero derivative: the per-wavenumber closed forms of Theorems 1–3 hold up to* $O(\eta^2)$.

Proof. At $\eta = 0$, every operator in the three recursions is diagonal in the Fourier basis. Differentiating the stationary equations at $\eta = 0$ gives, for each covariance, a linear equation $\dot{P} - \mathcal{L}(\dot{P}) = S$. For the believed covariance, the derivative of the gain, $\dot{K} = (I - K)\,\dot{P}^f (P^f + R)^{-1}$, is linear in $\dot{P}^f$ and belongs to the linear map $\mathcal{L}$; together with the propagation, $\mathcal{L}$ multiplies the (k, k') Fourier element by $(1 - g_k)(1 - g_{k'})\,A(k)\,\overline{A(k')}$, whose modulus is the geometric mean of the contraction factors $a(1-g)^2$ of the two wavenumbers and lies below one. The source S consists of the terms $\dot{A}P^a A^* + AP^a\dot{A}^*$ with $\dot{A} = BM$, each of which contains exactly one factor B between diagonal operators, so that it has zero Fourier diagonal, and so it has the unique solution $\dot{P}$. For the Level-1 and Level-2 covariances, the derivative of the gain is a source term, which has zero Fourier diagonal by the believed result, and the same argument applies with the contraction factors of the actual cycles. Differentiability follows from the implicit function theorem. □

The size of the second-order term depends on how different the variances of the coupled scales are. For the red toy spectra, in which the propagated analysis variance at the largest scale exceeds the smallest believed forecast variance by a factor of 8.8×10^3, full matrix recursions with $\|B\| = 1$ on a 32-point grid give maximal relative departures of the believed, Level-1 and Level-2 spectra of $648\,\eta^2$, $88\,\eta^2$ and $3.5 \times 10^3\,\eta^2$ for broadband coupling. When only wavenumbers at most two apart are coupled, the departures are $26\,\eta^2$, $4.8\,\eta^2$ and $57\,\eta^2$ (Figure 8b; fitted exponents between 1.999 and 2.001). A 1% accuracy of the per-wavenumber believed spectrum thus requires $\eta \lesssim 4 \times 10^{-3}$ for broadband and $\eta \lesssim 2 \times 10^{-2}$ for banded transfer. Transfer from energetic large scales into weak small scales is consequently the most damaging departure from the damping model, and the per-wavenumber formulae should not be applied where the diagnostic below reports an appreciable incoherent fraction.

Proposition 3 (envelope for time-dependent damping). *Let the damping vary in time,* $d_j \in [d_{min}, d_{max}] \subseteq [0,1]$, *independently of the state. Then, the believed analysis variance satisfies*

$$p(d_{\min}) \le \liminf_{j\to\infty} p_j^a \le \limsup_{j\to\infty} p_j^a \le p(d_{\max}), \tag{26}$$

where $p(d)$ *denotes the stationary value of Theorem 1 at constant damping* d*, and the gain* $g_j = p_j^a/r$ *obeys the same bounds. In particular, a wavenumber that is a blind spot at* d_{max} *remains one, and a wavenumber that is not a blind spot at* d_{min} *never becomes one.*

Proof. The believed map $\Phi_{\mathrm{bel}}(p) = r(ap+q)/(ap+q+r)$ is non-decreasing in both p and a. By induction, the orbit driven by d_j lies between the orbits driven by the constants $d_{\min}$ and $d_{\max}$ from the same initial condition, and these converge to $p(d_{\min})$ and $p(d_{\max})$ by Theorem 1. The gain identity is the update of Equation (3). □

The proposition covers flow-dependent damping to the extent that the damping realized along a trajectory can be bracketed. It does not cover a genuinely nonlinear emulator, in which different ensemble members are damped differently.

A diagnostic for the adequacy of the damping model. Whether a given emulator is well described by a real damping profile can be tested from one-step forecasts alone. Let $\tilde{x}_k$ and x_k^{ref} denote the Fourier coefficients of the emulator forecast and of a reference forecast (the physical model, or verifying analyses) started from the same states, and let $\langle\cdot\rangle$ denote an average over a sample of initial states. The cross-spectral transfer function and the squared coherence,

$$\hat{z}(k) = \frac{\langle \tilde{x}_k\, \overline{x_k^{\mathrm{ref}}} \rangle}{\langle |x_k^{\mathrm{ref}}|^2 \rangle}, \qquad \mathrm{coh}^2(k) = \frac{\left|\langle \tilde{x}_k\, \overline{x_k^{\mathrm{ref}}} \rangle\right|^2}{\langle |\tilde{x}_k|^2 \rangle \langle |x_k^{\mathrm{ref}}|^2 \rangle}, \tag{27}$$

give the amplitude ratio $\hat{d}(k) = |\hat{z}(k)|$, the phase error $\hat{\delta}(k) = \arg \hat{z}(k)$ and the incoherent fraction $1 - \mathrm{coh}^2(k)$. The last quantity is the share of the emulator variance at a wavenumber that is not linearly related to the reference at the same wavenumber, whether through scale transfer, nonlinearity or noise. The real-damping model is adequate at wavenumber k if (a) $4\hat{d}\sin^2(\hat{\delta}/2) \ll (1-\hat{d})^2$ and (b) $1 - \mathrm{coh}^2 \ll 1$. If (a) fails, Proposition 1 applies with the measured $\hat{z}$; if (b) fails, no per-wavenumber formula applies. On synthetic forecasts, the estimator recovers a prescribed complex multiplier to 6×10^{-15} and reports an incoherent fraction below 2×10^{-14} in magnitude. With a banded transfer of amplitude $\eta = 0.05$ added, the incoherent fraction rises to between 6×10^{-5} and 0.96, increasing towards the small scales, while the amplitude estimate stays within 7×10^{-4} of the prescribed profile (Figure 8c). The square root of the ratio of forecast to reference power spectra, the usual way in which spectral bias is reported, is less suitable: in the same test, it errs by up to 4×10^{-2}, because it cannot separate damping from incoherent variance.

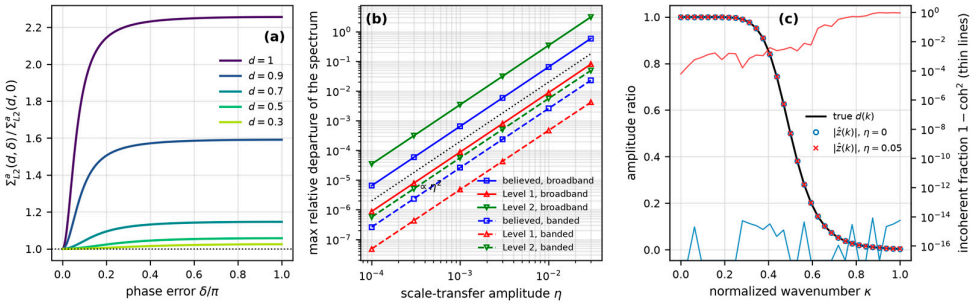


Figure 8. Departures from the real-damping model. (**a**) Ratio of the Level-2 error with a phase error δ to that without, for $m = 0.95$, $q = 0.01$, $r = 0.5$ and five amplitude factors d; the believed spectrum, the gain, the Level-1 error and the ensemble spread are independent of δ (dotted

line). (**b**) Maximal relative departure of the Fourier-diagonal stationary spectra from the per-wavenumber closed forms for an emulator $(D + \eta B)M$ with broadband (solid) and banded (dashed) scale transfer, against η; the dotted line has slope 2. (**c**) Cross-spectral diagnostic on synthetic one-step forecasts: estimated amplitude ratio $|\hat{z}(k)|$ without and with scale transfer (η = 0, blue circles; η = 0.05, red crosses, markers left axis) against the prescribed profile (line), and incoherent fraction $1 - \mathrm{coh}^2$ (thin lines in the same colours, right axis).

3.8. The Inflated Cycle: Corrected Fixed Point and Exact Budgets

Covariance inflation is the standard remedy for an under-dispersive filter [24,30]. Inflating the forecast variance, or carrying a process-noise variance q_f that differs from the true q, steps outside Assumption A9, under which the filter uses the true process noise. The believed spectrum of Theorem 1 then no longer applies as stated. The corrected fixed point, the resulting actual error and the inflation that is exactly required are as follows.

Proposition 4 (inflated cycle). *Let the filter propagate* $p^f_{j+1} = \lambda^2(a\, p^a_j + q_f)$*, with a multiplicative inflation factor* $\lambda \geq 1$ *(on standard deviations) and an assumed process-noise variance* $q_f > 0$*, the update being unchanged.*

(i) Corrected fixed point. The stationary believed forecast variance f_λ *is the unique positive root of*

$$f^2 + f\left[r\,(1 - \lambda^2 a) - \lambda^2 q_f\right] - \lambda^2 q_f\, r = 0, \tag{28}$$

that is, Theorem 1 holds with (a, q) *replaced by* $\left(\lambda^2 a, \lambda^2 q_f\right)$*, and the applied gain is* $g_\lambda = f_\lambda/(f_\lambda + r)$.

(ii) Actual error. The Level-1 error is $\Sigma(g_\lambda)$*, given by Theorem 2(i) with the true* q *and* a_0 *and the gain* g_λ*. The Level-2 error is given by Equations (14) and (15) with* g *replaced by* g_λ*. Neither equals the inflated believed variance.*

(iii) Exact budgets. The applied gain equals the optimal gain γ^* *of the perfect-model filter if and only if* $f_\lambda = a_0 p^* + q$*. This is achieved by additive inflation alone* $(\lambda = 1,\ q_f = q^*_f)$ *or by multiplicative inflation alone* $(q_f = q,\ \lambda = \lambda_*)$*, with*

$$q^*_f = q + (1 - d^2)\, m^2 p^*, \qquad \lambda^2_* = \frac{m^2 p^* + q}{d^2 m^2 p^* + q}, \qquad 1 \leq \lambda^2_* \leq \frac{1}{d^2}. \tag{29}$$

(iv) Sufficiency at Level 1 only. At either budget $\Sigma^a_{L1} = p^*$*, the minimum over all gains; any other inflation gives a larger Level-1 error. For* $m < 1$ *and* $d < 1$*,* $\Sigma^a_{L2}(\gamma) > p^*$ *for every gain* γ*: no inflation removes the emulator-in-the-loop error.*

Proof. (i) The stationary identity $f = \lambda^2(a\, fr/(f + r) + q_f)$ multiplies out to Equation (28), and Theorem 1 applies. (ii) The actual recursions (4) and (7)–(9) are driven by the true noise and by whatever gain is applied. (iii) The gain is a bijective function of the believed forecast variance. At $f_\lambda = a_0 p^* + q$, the believed analysis variance is p^*, and inserting this value into the stationary identity gives both budgets; the bounds on λ^2_* follow from $0 \leq d \leq 1$ and $q > 0$. (iv) By Lemma 2 and Equation (13), $\Sigma(\gamma)$ attains its minimum p^* at γ^* only. At Level 2, the analysis is a fixed linear combination of the past observations, with weight $\gamma\,[(1 - \gamma)\, d\, M(k)]^i$ on the observation that is i cycles old. The optimal filter is the unique minimum-variance estimator and has the weights $\gamma^*[(1 - \gamma^*)\, M(k)]^i$. The two sets coincide only if $\gamma = \gamma^*$ and $d = 1$. □

Three consequences follow. First, the additive budget has a direct reading: the filter must be told, as process noise, the true process noise plus exactly the forecast variance $(1 - d^2)\, m^2 p^*$ that the emulator fails to propagate. Second, the multiplicative budget never exceeds $1/d^2$ and tends to one where the process noise dominates the forecast

variance. Since it depends on scale through $d(k)$, $m(k)$ and $q(k)/r(k)$, a single global factor cannot satisfy it at all scales. Third, the gain-floor condition $q_f \geq \varepsilon r/(1-\varepsilon)$ discussed in Section 4 is a different requirement from the additive budget: it keeps the gain above a tolerance and does not aim at the optimal gain. It is weaker than the budget only where $\varepsilon r/(1-\varepsilon) \leq q + (1-d^2)\, m^2 p^*$. Where the perfect-model gain itself lies below the tolerance (case (ii) of Theorem 4), the floor demands more process noise than the budget and drives the gain beyond γ^*, which by Equation (13) raises the Level-1 error above p^*. In the small-scale regime of Table 6 with $d = 0.5$ and $\varepsilon = 0.05$, the floor requires 113 times the additive budget, and the resulting Level-1 error is 15 times p^*. Table 6 evaluates the budgets for the regimes of Table 2. At the typical mid-scale with $d = 0.5$, the required additive inflation is 4.3 times the true process noise or, equivalently, $\lambda_* = 1.60$. It reduces the Level-1 error from 1.42 to exactly 1 times the optimum but reduces the Level-2 error only from 1.93 to 1.59.

Comparison with reported inflation factors. The only study cited here that reports inflation magnitudes is that of Kotsuki et al. [12], in which the adaptive multiplicative inflation exceeded 5 in places with the AI-based model. The same study contrasts this with values of at most about 2 that it quotes for an earlier global system with a dynamical model and real observations; the two systems differ in model, resolution and observations. A like-for-like test of Equation (29) is therefore not possible: the damping profile and the noise spectra of either system are not published, and the adaptive inflation also compensates sampling error, which Proposition 4 does not contain. Only the sign of the effect can be compared. It is consistent with $\lambda_* > 1$ wherever $d < 1$, but any under-dispersive model error produces the same sign, so the comparison does not discriminate the spectral mechanism. If the ratio 2.5 of the two reported values were attributable to the emulator alone, the bound $\lambda_*^2 \leq 1/d^2$ would require $d \leq 0.63$ at the scales that control the adaptive estimate; this reading presumes that the reported factors multiply variances and is not a validation.

Table 6. Exact inflation budgets of Proposition 4 and the resulting errors, relative to the optimal variance p^* of the perfect-model filter, for regimes of Table 2 (parameters m, q, r in parentheses). g is the believed gain without inflation and γ^* the gain restored by either budget; the last column is the Level-2 error after the gain has been restored. A dash marks growing modes, at which Level 2 has no stationary value.

Regime	d	g	γ^*	q_f^*/q	λ_*	Σ_{L1}^a/p^*	Σ_{L2}^a/p^*	$\Sigma_{L2}^a(\gamma^*)/p^*$
typical mid-scale (0.95, 10^{-2}, 0.5)	0.9	0.060	0.097	1.83	1.09	1.08	1.30	1.13
typical mid-scale (0.95, 10^{-2}, 0.5)	0.7	0.034	0.097	3.24	1.31	1.29	1.78	1.42
typical mid-scale (0.95, 10^{-2}, 0.5)	0.5	0.025	0.097	4.30	1.60	1.42	1.93	1.59
strong process noise (0.90, 1, 0.5)	0.5	0.681	0.721	1.22	1.10	1.01	1.39	1.27
small scale, small process noise (0.80, 10^{-4}, 0.5)	0.5	0.00024	0.00055	2.33	1.39	1.00	1.00	1.00
well-observed neutral scale (1.00, 10^{-2}, 0.1)	0.5	0.114	0.270	3.03	1.49	1.58	–	–
well-observed growing scale (1.05, 10^{-2}, 0.1)	0.5	0.117	0.302	3.50	1.54	2.17	–	–
unstable, tiny process noise (1.20, 10^{-4}, 0.5)	0.9	0.144	0.306	419	1.11	∞	–	–

3.9. Sensitivity of the Blind-Spot Extent to Noise Spectra and Observation Density

Theorem 4 classifies each wavenumber separately. How much of the spectrum lies in a blind spot depends on the spectra of the process noise and of the observation error and on the density of the observations. This dependence is made explicit here.

Corollary 1 (reduction to the noise ratio). *The stationary gain depends on* q *and* r *only through the ratio* $s = q/r$. *With* $\tau = \varepsilon/(1-\varepsilon)$, *a wavenumber is never a blind spot if* $s \geq \tau$; *it is a blind spot for every* $d \in [0,1]$, *including the perfect model, if* $s < \tau\,[1-(1-\varepsilon)\,m^2]$; *and otherwise, it is a blind spot exactly for* $d < d^*$, *with*

$$d^{*2}m^2 = \left(1-\frac{s}{\tau}\right)(1+\tau). \tag{30}$$

Proof. Dividing the stationary quadratic of Theorem 1 by r^2 gives $(f/r)^2 + (f/r)\,[(1-a)-s]-s=0$, so that f/r and the gain depend on (a,s) only. The three cases are those of Theorem 4 written with $\varphi/r = \tau$. □

Observation density. If only a fraction ρ of the grid points is observed, on a regular sub-grid with $1/\rho$ an integer that divides n, and the observation error is white, $R = r\,I$, then $H^T R^{-1} H = (\rho/r)(I+C)$, where C couples each wavenumber to its $1/\rho - 1$ aliases and has zero Fourier diagonal. The diagonal part is the fully observed problem with r replaced by r/ρ, that is, with s replaced by $\rho\,s$ in Corollary 1. For a coloured observation-error spectrum, this substitution is an approximation that is not examined here. The coupling enters the Fourier-diagonal variances only through products of the gains of aliased wavenumbers, which are small when at most one member of each alias family is well observed. Full matrix Riccati solutions for the toy model confirm this reduction. With every second, fourth and eighth point observed, the blind fraction (0.83, 0.89 and 0.95, against 0.80 for dense observations) coincides with that of the diagonal theory with r/ρ, and the Fourier-diagonal analysis variance agrees to within 3×10^{-5}, 3×10^{-4} and 1×10^{-2}.

Spectral slopes. Let the noise ratio follow $s(\kappa) = s_0\,[1+(\kappa/0.15)^2]^{-\Delta\beta/2}$, so that $\Delta\beta = \beta_q - \beta_r$ is the difference between the spectral slopes of the process noise and of the observation error and s_0 is the ratio at the largest scale; the toy model of Section 3.5 has $s_0 = 0.04$ and $\Delta\beta = 4$. In Figure 9, the observation error is white and the slope difference is carried by the process noise, so that the density enters through the diagonal reduction $s \to \rho\,s$ described above. Figure 9 shows the fraction of wavenumbers in a blind spot at $\varepsilon = 0.05$ over the plane $(\Delta\beta, \rho\,s_0)$ for the propagator of the stable toy configuration. Three regimes appear. Where $\rho\,s(\kappa) \geq \tau$ at every wavenumber, no wavenumber is blind for any emulator. Where $\rho\,s_0$ lies far below τ, almost every wavenumber is blind even for the perfect model: the scale is under-observed, and the emulator is immaterial. The emulator matters in the band between, where the perfect-model filter retains a working gain that the damping removes. There the mild, moderate and severe profiles add up to 0.32, 0.55 and 0.73 to the blind fraction.

The toy model itself lies in the second regime. At $s_0 = 0.04 < \tau = 0.053$ and $\Delta\beta = 4$, its blind fraction is 0.80 for the perfect model and 0.80, 0.80 and 0.82 for the three damping profiles. In that configuration, the blind spots of Figure 4 are therefore caused by the noise ratio and not by the emulator; what the emulator changes there is the believed variance and the ensemble spread, not the classification. With $s_0 = 1$ and $\Delta\beta = 2$, by contrast, the blind fraction rises from 0.10 for the perfect model to 0.26, 0.36 and 0.36 for the mild, moderate and severe profiles. Reducing the observation density lowers $\rho\,s$ at every wavenumber and moves the system downwards in Figure 9a,b. For the toy model, the blind fraction grows as quoted above and reaches one below $\rho \approx 0.09$. For $s_0 = 1$ and $\Delta\beta = 1$,

no wavenumber is blind at full density, whereas at $\rho = 1/8$ the moderate profile blinds 53% of the wavenumbers against 21% for the perfect model (Figure 9c).

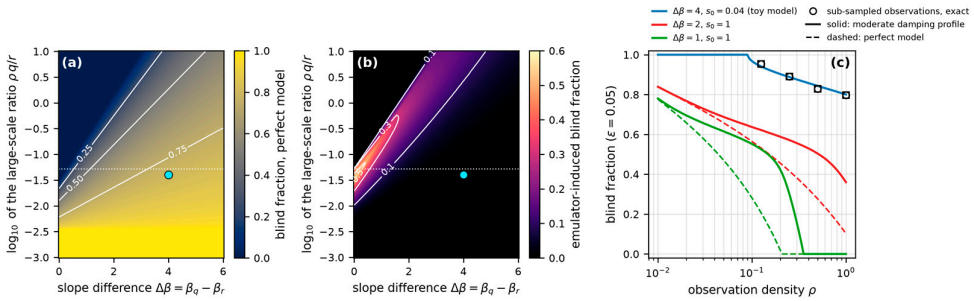


Figure 9. Fraction of wavenumbers in a blind spot at tolerance $\varepsilon = 0.05$ for the propagator of the stable toy configuration. (**a**) Perfect model, over the plane of the slope difference $\Delta\beta$ between process-noise and observation-error spectra and the large-scale noise ratio $\rho\, q/r$ (logarithmic). (**b**) Additional fraction caused by the moderate damping profile. The dotted line is $\rho\, s_0 = \tau$ and the dot marks the toy model. (**c**) Blind fraction against the observation density ρ for three spectra (solid: moderate profile; dashed: perfect model); markers are full matrix Riccati solutions with regularly sub-sampled observations.

3.10. Ensemble Filters Beyond Assumption A10: Gridpoint, Localized and Sparsely Observed Cases

Theorem 5(i)–(iii) is exact for the mode-wise filter. Three departures are examined with a serial EnSRF in gridpoint space, on the stable toy configuration with the moderate damping profile and without model-error injection: finite ensembles without localization, covariance localization, and sparse observations (Figure 10 and Table 7).

Unlocalized filter. The bound of Theorem 5(iv) holds in every experiment: for ensembles of 5 to 200 members, the largest value of $\hat{s}^f_{j+1}(k)/[a(k)\,\hat{s}^f_j(k)]$ over all wavenumbers and cycles lies between 0.992 and 0.99999. The spread follows its own rate down to the smallest representable numbers. In an implementation that carries the ensemble in Fourier coordinates, so that every mode keeps its own relative precision, the spread of the deeply damped modes ($a < 0.2$) after 20 cycles has the median 3.6×10^{-63}, below the bound $a^{20}\hat{s}_0$, whose median is 5.0×10^{-62}. In ordinary gridpoint double-precision arithmetic, the same modes stall at 6.7×10^{-34}, that is, at round-off relative to the largest-scale variance. This plateau is numerical and is not a sampling floor (Figure 10a); the two implementations agree to 8×10^{-13} at moderately damped modes. The exact infinite-ensemble map is approached from below as the ensemble grows (Figure 10b): the median ratio of spread to map at moderately damped modes after 20 cycles is 0.20, 0.49, 0.83, 0.94, 0.98 and 0.99 for $N = 10$, 20, 40, 80, 200 and 400 at $n = 64$. The controlling parameter is the ratio N/n: at $N/n = 0.08$ the ratio is 0.09 both for $(n, N) = (256, 20)$ and for $(512, 40)$, and at $N/n = 0.04$ it is 0.04. An under-sampled unlocalized filter thus collapses faster than the mode-wise law, which then overestimates the surviving spread. The cause is not the sampling bias of Theorem 5(iii) but the spurious cross-wavenumber sample covariances of a finite ensemble. For $H = I$ and $R = rI$, the analysis sample covariance is $h(\hat{P}^f)$ with $h(x) = rx/(x + r)$. Since h is strictly concave, Jensen's inequality over the spectral decomposition of $\hat{P}^f$ gives $\hat{s}^a(k) \le h(\hat{s}^f(k))$ at every wavenumber, with equality only where the sample covariance has no cross-wavenumber terms at that wavenumber. The mode-wise filter attains equality, whereas the gridpoint filter at $n = 64$ and $N = 20$ has the median ratio $\hat{s}^a(k)/h(\hat{s}^f(k)) = 0.98$ per update at the moderately damped modes.

Localized filter. With Gaspari–Cohn localization [20] of half-width c grid points, the picture changes in two ways. First, the spread at moderately damped modes returns to the mode-wise map. For $c = 8$, the ratio lies between 0.97 and 1.04 for all six combinations of $n \in \{64,256,512\}$ and $N \in \{20,40\}$, independently of the dimension; it lies between 1.10 and 1.14 for $c = 4$ and falls to between 0.54 and 0.86 for $c = 32$. Second, the per-mode bound of Theorem 5(iv) fails: the 95th percentile of $\hat{s}_{j+1}^f/(a\,\hat{s}_j^f)$ lies between 1.04 and 1.27 at $c = 8$ and reaches 1.62 at $c = 4$. The reason is that tapering the sample covariance in grid space convolves it in wavenumber space, so that sample variance is transferred between wavenumbers at every update. The deeply damped modes ($a < 0.05$) are the clearest case. Instead of decaying at the own rate, the spread of these modes decays, for $c \leq 16$, by a factor per cycle of 0.34 to 0.41 times the amplification of the least damped mode, and after ten cycles it exceeds the own-rate value by nine to fifteen orders of magnitude. Localization therefore restores the mode-wise law where the spread is appreciable and replaces it by a slaved, slower decay where the spread is negligible. Neither effect prevents the collapse.

Sparse observations. With every s-th grid point observed ($N = 40$, no localization), the bound of Theorem 5(iv) continues to hold, as it must, since it does not depend on the observation operator. The spread is, however, larger than the dense-observation map predicts, by factors of 1.45, 2.37 and 3.65 for $s = 2$, 4 and 8 after 20 cycles (0.84 for $s = 1$): fewer observations remove less spread per cycle. The collapse criterion $d\,m < 1$ is unaffected, because without any observations the spread still decays at exactly a per cycle.

Where the law ceases to extrapolate. In summary, four statements delimit the quantitative reach of Theorem 5. (a) The exact law of items (i)–(iii) is accurate for gridpoint filters when the ensemble resolves the state, with N about three times n without localization, or when the localization half-width is close to the value that balances sampling error, about eight grid points here. (b) It overestimates the surviving spread of undersampled unlocalized filters, by an order of magnitude for $N/n \leq 0.1$. (c) It underestimates the surviving spread under sparse observation and under tight localization. (d) Under any localization, it does not describe the deeply damped modes. The mode-wise experiments of Figure 5 and Table 3 and the 64-point, densely observed gridpoint test are thus the favourable end of the range. The theorem should be used as a statement of mechanism and of bounds, with Table 7 as a guide to its quantitative reach.

Table 7. Gridpoint serial EnSRF against the mode-wise law of Theorem 5 (stable toy configuration, moderate damping, dense observations, no model-error injection; means over three realizations, with standard deviations of the upper block at most 0.06). Upper block: median over the moderately damped modes ($0.2 < a < 0.95$) of the ratio of the analysis spread to the infinite-ensemble mode-wise map after 20 cycles. Lower block: 95th percentile of $\hat{s}_{j+1}^f/(a\,\hat{s}_j^f)$ over the same modes and all cycles; a value above one violates the per-mode bound of Theorem 5(iv), which is proved for the unlocalized filter only. c is the Gaspari–Cohn half-width in grid points. The unlocalized entries differ from the eight-realization values of Figure 10b within the sampling scatter.

(n, N)	$(64, 20)$	$(64, 40)$	$(256, 20)$	$(256, 40)$	$(512, 20)$	$(512, 40)$
Spread/mode-wise map, cycle 20						
no localization	0.46	0.85	0.090	0.21	0.039	0.087
$c = 4$	1.10	1.12	1.12	1.14	1.14	1.11
$c = 8$	0.97	1.00	1.02	1.04	1.02	1.01
$c = 16$	0.82	0.93	0.86	0.95	0.85	0.94
$c = 32$	0.65	0.86	0.57	0.78	0.54	0.78
95th percentile of the per-step ratio						

no localization	0.99	1.00	0.97	0.97	0.97	0.97
$c = 4$	1.21	1.12	1.53	1.35	1.62	1.42
$c = 8$	1.10	1.04	1.21	1.12	1.27	1.15
$c = 16$	1.03	1.02	1.05	1.02	1.06	1.03
$c = 32$	1.01	1.01	1.00	1.00	1.00	1.00

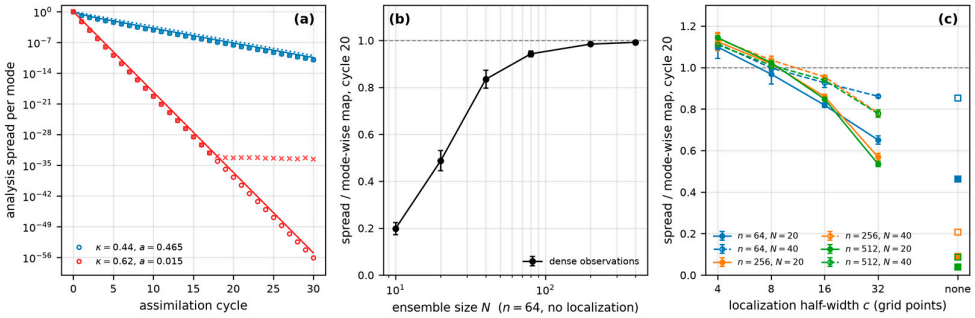


Figure 10. Gridpoint serial EnSRF beyond Assumption A10 (stable toy configuration, moderate damping, $q_{\mathrm{ens}} = 0$). (**a**) Analysis spread against cycle at a moderately and a deeply damped wavenumber for $N = 20$ and $n = 64$ without localization: Fourier-coordinate arithmetic (circles), gridpoint double-precision arithmetic (crosses), mode-wise map (solid) and the bound $a^j \hat{s}_0$ of Theorem 5(iv) (dotted). (**b**) Median ratio of spread to mode-wise map at moderately damped modes after 20 cycles against ensemble size (mean and standard deviation over eight realizations). (**c**) The same ratio with Gaspari–Cohn localization of half-width c for three grid sizes and two ensemble sizes; squares at the right give the unlocalized values.

3.11. Verification Summary and Protocol

Every theorem is reproduced by an independent numerical path, and the standard demanded, agreement to 10^{-10} in the linear–Gaussian cases, is met by every comparison of a closed form with an exact numerical solution of the same recursion. The margin is many orders of magnitude for the scalar closed forms and at least a factor of about 2 for the full matrix recursions. Comparisons with stochastic simulations agree to within the sampling level, and the gridpoint ensemble check of Theorem 5(v), which departs from the mode-wise setting, agrees to 5.7×10^{-7} (Table 8). The scalar closed forms were certified in 50-digit arithmetic against fixed points obtained *only* from the raw one-cycle maps (bisection for Theorem 1; affine solves assembled by evaluating the raw maps on basis vectors for Theorems 2 and 3) over a grid of $10 \times 21 \times 5 \times 5$ values of (m, d, q, r). The maximal relative deviations are 3.2×10^{-46} (Theorem 1, 5250 cases), 6.7×10^{-43} (Theorem 2, 4158 stable cases, with all 1092 predicted divergences matched), 1.3×10^{-47} (Theorem 3, 2625 cases) and 7.0×10^{-49} (Theorem 4, 84 thresholds). The double-precision implementation lies within 1.9, 1.8 and 61 units of $\epsilon_{\mathrm{mach}}/(1 - \vartheta)$ of these 50-digit reference values, where ϑ is the contraction factor of the corresponding recursion, i.e., at its conditioning limit. Symbolic verification in exact arithmetic covers every identity used in the proofs, including the positivity certificate of Theorem 3. Further results of Sections 3.2, 3.4 and 3.7–3.10 (Remark 2, Propositions 1–4, Corollary 1 and Theorem 5(iii)–(v)) were verified in the same manner and are listed in the lower part of Table 8. The protocol is as follows.

In the *symbolic layer*, eighteen computer-algebra checks in exact rational arithmetic cover every identity used in the proofs, including the perfect-square discriminant of Theorem 4 and the positivity certificate of Theorem 3. Thirteen further checks cover the impulse-response form of the Level-2 error, the certificate of Theorem 3 in lowest terms, the inflated fixed point and the budgets of Proposition 4, the growing-mode limit of Remark 1 and the expansion coefficients of Theorem 5. The numerical layer consists of the

following: (i) 50-digit certification of all scalar closed forms against raw-map fixed points on the grid $m \in \{0.3,0.7,0.9,0.95,0.99,1.0,1.02,1.05,1.2,1.5\}$, $d \in \{0,0.05,\ldots,1\}$, $q \in \{10^{-6}, 10^{-4}, 10^{-2}, 1, 10^{2}\}$, $r \in \{10^{-3}, 10^{-1}, 1, 10, 10^{3}\}$; (ii) full-matrix (64×64) gridpoint-space recursions for the believed, Level-1 and Level-2 cycles against the per-mode closed forms; (iii) stochastic Monte Carlo cycling; (iv) cycled spectral and gridpoint EnSRF, including the exact update identity, collapse rates, sustained spread and the χ^2 quadrature of the sampling bias; (v) Lorenz-96. Every comparison met its prescribed tolerance, and no tolerance was relaxed to obtain agreement. Two numerical choices made during verification are documented, with the rationale, in the repository. The first is a cancellation-stable evaluation of the quadratic root of Theorem 1 (the textbook form loses about four digits when $r(1-a) > q$). The second is a strictly stable toy configuration ($\lambda = -0.01$ rather than a marginal $m = 1$ mode, whose random-walking climatology contaminated neighbouring modes with round-off in the long matrix recursions).

Table 8. Master verification summary. The lower part lists the checks of Remark 2, Propositions 1–4, Corollary 1 and items (iv) and (v) of Theorem 5, and further checks of Lemma 1, Theorem 3 and Theorem 5(iii).

Theorem or Claim	Independent Check	Result
Theorem 1, closed form	50-digit raw-map bisection, 5250 cases	3.2×10^{-46}
Theorem 2, closed form, $G \geq 0$, divergences	50-digit raw-map affine solve, 4158 cases; 1092 divergences matched	6.7×10^{-43}
Theorem 3, closed form and ordering	50-digit raw-map 3×3 solve, 2625 cases; positivity certificate	1.3×10^{-47}
Theorem 4, threshold formula	believed gain at d^* equals ε, 84 cases	7.0×10^{-49}
Lemma 1 and Theorems 1–3 in matrix form	full 64×64 gridpoint-space recursions	$\leq 5.6 \times 10^{-11}$
Theorem 3, stochastic interpretation	Monte Carlo cycling, 2000×400	1.3% (sampling level)
Theorem 5, map, rate, sustained spread	cycled spectral EnSRF, $N = 20$	$\leq 1.1 \times 10^{-13}$
Theorem 5, sampling coefficient	χ^2 quadrature, $\nu = 1279$	-0.4998 vs. -0.5
Theorem 5, gridpoint filter (A10 relaxed)	serial EnSRF, $N = 200$, 40 cycles, eight realizations; ratio at cycle 20	0.984 ± 0.003 of the map
Divergence band ($m > 1$, $d < d_c$)	Level-1 matrix recursion over 300 cycles: 8 predicted vs. 8 observed divergent modes	exact
Lorenz-96 illustration	qualitative	consistent
Lemma 1 on a larger grid	full 256×256 gridpoint-space recursions	$\leq 1.7 \times 10^{-11}$
Theorem 3, impulse-response form and certificate structure	computer algebra: identity with the closed form; reduced certificate with 47 positive integer coefficients	exact
Theorem 3, ordering	random scan of 2×10^6 draws, real damping, believed gain	minimum -2×10^{-16}; no violation
Remark 2, state-independent emulator noise	Monte Carlo cycling; 2×10^5-draw scan of the ordering	0.24%; no violation
Proposition 1, complex damping	closed form against iteration; Monte Carlo; 3×10^6-draw scan of both orderings	9.7×10^{-16}; 0.35%; no violation

https://doi.org/10.3390/e28111276

Proposition 2, scale transfer	full matrix recursions, $n = 32$, six amplitudes η	exponents 1.999–2.001
Proposition 3, envelope	300 random time-dependent damping sequences	no violation
Proposition 4, inflated fixed point and budgets	50-digit bisection; 4000 random cases	3.0×10^{-47}; $\leq 4.3 \times 10^{-14}$
Corollary 1, band formulae	20,000 random cases; matrix Riccati with sub-sampled observations	no mismatch
Theorem 5(iii), exact expectation	50-digit quadrature, $\nu = 4$ to 318	1.9×10^{-48}
Theorem 5(iv), general bound	unlocalized gridpoint EnSRF, $N = 5$ to 200; sparse observations	largest ratio 0.99999
Theorem 5(v), inflated survival	gridpoint EnSRF, $N = 400$, $\lambda = 1.05$ to 1.5	$\leq 5.7 \times 10^{-7}$

4. Discussion

From conjecture to design rules. The conjecture of Smith et al. [6], that the spectral errors of an emulator reduce ensemble spread within data assimilation and thereby limit the usefulness of the assimilation, is confirmed in the linear setting and sharpened. Absent model-error injection, spread reduction at damped scales is a theorem, with the exact per-cycle rate $d(k)^2 m(k)^2$ and the sharp survival criterion $d(k)m(k) > 1$. A perfect-model ensemble without injected variance also loses spread at decaying scales, at the rate $m(k)^2$, so that the emulator-specific content is the factor $d(k)^2$ and the displacement of the survival boundary. The conjecture is also qualified: the reduction is conditional, since any injected variance prevents it and multiplicative inflation λ relaxes the criterion to $\lambda\, d_k m_k > 1$. The practically actionable objects are the closed forms. Given an emulator's measured amplitude ratio $d(k)$, obtained preferably from the cross-spectral diagnostic of Section 3.7 rather than from a ratio of power spectra, a system designer can compute, scale by scale, the believed and actual error spectra of Equations (10)–(11) and Theorem 2, the overconfidence fraction G/Σ^a, the blind-spot threshold of Equation (20), and inflation budgets. Two budgets must be distinguished, and the first is a necessary condition only. This *gain-floor budget* follows from Theorem 1: the believed forecast variance obeys $f \geq q$, so a filter carrying a process-noise variance q_f at a fully damped scale retains a gain of at least $q_f/(q_f + r)$, and holding this above a tolerance ε requires $q_f \geq \varepsilon r/(1 - \varepsilon)$. The condition is necessary for a working gain, not sufficient for a good analysis: it restores weight on the observations without restoring the missing forecast variance, so the residual error is bounded by Theorem 2 rather than removed. Where the perfect-model gain itself lies below the tolerance, the floor exceeds the exact budget and over-weights the observations (Section 3.8). The *exact budget* of Proposition 4 restores the optimal gain; it is sufficient for an optimal analysis at Level 1 and provably insufficient at Level 2. In either case, inflating q_f beyond the true q steps outside Assumption A9. The actual error must then be evaluated from Proposition 4(ii), with the true q in the forcing and the inflated gain in the update, and not from the believed spectrum of Theorem 1. The empirical finding of Kotsuki et al. [12] that an AI-based model requires larger inflation factors when cycled in an EnKF is consistent in direction with these budgets, although it does not discriminate the spectral mechanism (Section 3.8).

A recipe for practitioners. The results translate into the following procedure for a given emulator and observing system. Step 1: from one-step emulator and reference forecasts started from common states, estimate the transfer function $\hat{z}(k)$ and the coherence by Equation (27). Step 2: check the two adequacy conditions of Section 3.7; where the incoherent fraction is appreciable, stop, since no per-wavenumber formula applies. Step 3:

Estimate the per-cycle amplification $m(k)$ of the reference model and the noise ratio $\rho\, s(k)$, with $s(k) = q(k)/r(k)$. The process-noise spectrum required here is the true one, not the tuned value carried by the existing system, which already contains compensatory inflation; it must be estimated, for example, from innovation statistics resolved by scale, and its uncertainty propagates into Steps 4–6. The factor ρ applies to a regular observing network with white observation error (Section 3.9). Step 4: compute the gain from Equations (10)–(11) and classify every wavenumber by Corollary 1 as never blind, blind even for a perfect model, or blind below $d^*(k)$. Step 5: compute the believed and actual spectra and the overconfidence fraction G/Σ^a, using Theorem 2 if only the covariances are emulator-based and Theorem 3 or Proposition 1 if the forecast mean is emulator-based as well; at growing scales, only Theorem 2 is available. Step 6: compute the exact inflation budget of Equation (29) per scale and compare it with the inflation actually applied; a uniform factor below $1/(d_k m_k)$ leaves the spread at scale k collapsing. Step 7: verify in the cycling system with innovation statistics resolved by scale, since the covariances of the filter give no warning. The recipe inherits every limitation listed in Section 5 and yields a per-scale, first-order guide, not a prediction for a nonlinear, localized system.

Two deployment configurations. Level 1 (gain miscalibration only) and Level 2 (emulator in the loop) bracket practice, and the distinction should be kept in view when any formula in this paper is applied. Level 1 describes systems whose covariance model, rather than forecast mean, is emulator-based, for example ML-generated ensembles feeding a hybrid background-error covariance [3,16]; Level 2 describes fully ML-driven cycles in which the emulator supplies both mean and covariance. The ordering $\Sigma^a_{\mathrm{L2}} \geq \Sigma^a_{\mathrm{L1}}$ of Theorem 3 means that, at every decaying mode, each Level-1 formula is a floor for the full emulator loop: whenever it already signals overconfidence or a blind spot, the full deployment is affected at least as strongly. The comparison is made at the gain that each configuration actually produces, both being the believed gain of Theorem 1 for the same (a, q, r); for an arbitrary gain imposed on both configurations the ordering need not hold.

Why the filter cannot diagnose itself. A striking feature of Theorem 2 is that the believed covariances remain well behaved at all times: the filter has no internal signal of failure. The stability wedge of item (ii) makes this concrete for growing modes, where the believed contraction always holds while the actual contraction fails below d_c. Diagnosing spectral blind spots therefore requires observation-space consistency statistics resolved per scale, such as innovation spectra, rather than global error scores; this is a concrete recommendation for the evaluation of ML–DA systems, in line with the diagnostic use of ensemble DA proposed by Kotsuki et al. [12]. Phase errors sharpen the point: by Proposition 1, a phase error changes no quantity that the filter computes, and at Level 2, it is visible only in the innovations.

Generative emulators and distributional training. The mechanism operates through the deterministic amplitude factor $d(k)$. Emulators trained to match distributions rather than conditional means, such as diffusion-based generators with correct spectra [8,33] or ensembles trained on proper scoring rules [3], address it, though not by restoring $d(k) = 1$ in the propagator. A stochastic emulator that reproduces the correct climatological spectrum supplies the small-scale variance that the deterministic map removes: in the language of Theorem 5, it acts as a non-zero q_{ens}, which is precisely the quantity that prevents spread collapse, and in that of Theorem 1, it raises the process-noise term that sets the gain floor. The present framework does not represent such emulators directly, since $\mathbf{D}$ is deterministic, and extending it to stochastic propagators is a natural next step. The same holds for end-to-end data-driven systems that learn the assimilation step itself [34] and for diffusion-based assimilation, whose target posterior has recently been shown to differ from that of a cycling ensemble filter [35]. The qualitative target it supplies is nonetheless clear: what the assimilation cycle requires of an emulator is not small-scale skill

but small-scale variance calibration, since even an unpredictable scale with the correct variance keeps the gain, the spread and the error budget mutually consistent.

Overconfidence and accuracy are different quantities. The gap G measures how far the filter under-reports its error, not how far the error exceeds the best attainable value p^*. The two can differ widely. In the toy model with the moderate profile, the overconfidence fraction G/Σ^a reaches 0.75, while the actual error exceeds p^* by at most 1.3% at Level 1 and 2.5% at Level 2; for the severe profile, the figures are 0.88, 20% and 36%. The reason is that the largest overconfidence occurs at scales where even the optimal gain is small, so that losing it costs little accuracy. The accuracy penalty is large where the optimal gain is appreciable and the damping substantial, as at the typical mid-scale of Table 6 (42% at Level 1 and 93% at Level 2 for $d = 0.5$), and it is unbounded at growing scales below the critical damping. In a linear, per-scale setting, the first casualty of spectral bias is therefore the reliability of the uncertainty estimate, and the loss of accuracy follows where the observations matter. In a nonlinear system, the two are not separable, because under-dispersion at one scale degrades the analysis at others, as the Lorenz-96 experiment shows.

5. Limitations

The theory is deliberately idealized, and its boundaries are as follows. (i) Linearity and Gaussianity. The closed forms are exact only for linear dynamics and Gaussian errors. Lorenz-96 shows that the mechanism persists under nonlinearity but also that nonlinearity exports damage across scales; the linear formulae are not bounds for a nonlinear cycle. (ii) Homogeneity and dense direct observation. Circulant statistics and the identity observation operator (Assumption A7) are what diagonalize the cycle; here, dense direct observation plays the role that homogeneous–isotropic statistics play in turbulence theory, namely the idealization that makes the per-wavenumber closed forms possible. Gappy or indirect observation operators couple wavenumbers, and localization strengthens the coupling. The structure of the required extension is already visible. For observations on a regular sub-grid of density ρ with white observation error, $H^T R^{-1} H = (\rho/r)(I + C)$, where C links each wavenumber to its aliases. The diagonal part is the present theory with r replaced by r/ρ. In a formal expansion in C, the first-order effect on the Fourier-diagonal variances vanishes, as in Proposition 2, and the leading correction involves products of the gains of aliased wavenumbers. In the toy model, the diagonal theory with r/ρ reproduces the full matrix solution to 3×10^{-5}, 3×10^{-4} and 1×10^{-2} for $\rho = 1/2,\ 1/4$ and $1/8$ (Section 3.9). A one-step estimate of the correction, the gain at a wavenumber times the sum of the gains at its aliases, has the right order of magnitude but underestimates these departures by up to an order of magnitude. For a general network, C is the Fourier transform of the observation-precision field evaluated at wavenumber differences, and it is banded when that field varies smoothly in space; this is the case in which a banded (quasi-diagonal) perturbation expansion applies. Such an expansion is the natural next step and is not attempted here. Until it is available, the per-wavenumber formulae should be applied to densely observed fields, or with r/ρ, as a scale-by-scale guide only. (iii) The damping model. The baseline is a real, time-independent **D** commuting with **M**. Section 3.7 shows that phase errors leave Theorems 1, 2, 4 and 5 unchanged and increase the Level-2 error, that scale-transfer errors enter at second order with a coefficient that is large for red spectra, and that time-dependent damping is bracketed by the constant-damping results. State-dependent damping, in which ensemble members are damped differently, lies outside the theory. (iv) Stationarity. The results concern the asymptotic cycle; transients are covered only by the global-attraction statements. Level 2 has no stationary spectrum at growing modes (Remark 1). (v) Assumption A10. The exact ensemble laws of Theorem 5(i)–(iii) hold for mode-wise (spectral) filters, and the bound of Theorem 5(iv) holds for unlocalized square-root filters. Section 3.10 shows that an

unlocalized gridpoint serial EnSRF follows the exact law to within 2% at moderately damped modes when the ensemble comfortably resolves the state ($N = 200$ against dimension 64), whereas at $N = 20$ its spread after 20 cycles is about half the infinite-ensemble prediction and falls further as the dimension grows. Localization restores the agreement at moderately damped modes but invalidates the per-mode bound. Theorem 5 quantifies the mechanism, not the behaviour of an operational, localized, inflated, nonlinear ensemble. (vi) The toy configuration. In the toy model, the extent of the blind spots is set by the noise ratio and not by the emulator (Section 3.9), and the accuracy penalty is small where the overconfidence is largest (Section 4). The figures illustrate the formulae and are not estimates of operational impact. (vii) Knowledge of the noise statistics. Under Assumption A9, the filter carries the true q and r. Section 3.8 treats a process-noise variance that differs from q, but a misspecified r, or a q estimated with error, changes the believed gain and every threshold derived from it. (viii) Univariate, one-dimensional fields. The theory concerns a single scalar field on a one-dimensional periodic grid. Cross-variable covariances, whose misrepresentation by ML models has been reported in cycling experiments [13], and multivariate balance are not represented. (ix) No trained emulator. The damping profiles are prescribed and none is fitted to a trained emulator; the Lorenz-96 forecast step is the true model composed with a prescribed damping. (x) The diagnostic. The transfer function of Equation (27) is estimated against a reference forecast whose own error is not represented in the theory. Items (i)–(x) are to be regarded as homogeneous–isotropic turbulence theory is regarded: a first-order isolation of a mechanism, not a simulation of practice.

6. Conclusions

In a linear–Gaussian, homogeneous assimilation cycle, spectral bias in a forecast emulator has three provable consequences. It makes the filter overconfident at damped scales by a closed-form, monotone gap (Theorems 1–3). Beyond an exact damping threshold, it lowers the weight of the observations below a prescribed tolerance, wherever that tolerance exceeds the gain floor set by the process noise, with divergence at weakly forced growing scales (Theorem 4). Without injected variance, it collapses ensemble spread at the rate d^2m^2 per cycle wherever $d\,m < 1$, against m^2 for a perfect model at decaying scales (Theorem 5). Four extensions delimit these results. Phase errors of the emulator are invisible to the filter yet increase the emulator-in-the-loop error. Scale-transfer errors enter at second order. Exact inflation budgets restore the optimal gain but repair the error only when the forecast mean is not emulator-based. For gridpoint ensemble filters, the collapse law is a rigorous bound without localization. With localization, it describes the moderately damped modes to within a few per cent at a suitable half-width and does not describe the deeply damped modes. The closed forms of Theorems 1–4 are verified symbolically and agree with 50-digit fixed-point computations to better than 10^{-42}, and the full matrix recursions agree with the closed forms to better than 10^{-10}. The practical outcome is a set of per-scale design rules, comprising a diagnostic, gain floors, inflation budgets, blind-spot thresholds and survival criteria, that turn a shared qualitative concern about ML emulators in data assimilation [3,6] into arithmetic. The rules are first-order guides derived in a linear setting; the Lorenz-96 experiment shows that in a nonlinear system the degradation extends to scales that the linear theory leaves untouched.

Author Contributions: Conceptualization, N.L., P.G., S.K., and M.Z.; methodology, N.L. and P.G.; software, N.L.; validation, N.L., P.G., S.K., and M.Z.; formal analysis, N.L.; investigation, N.L. and M.Z.; resources, R.K.G., F.F.B.H., D.G.K., and M.Z.; data curation, N.L.; writing—original draft preparation, N.L.; writing—review and editing, P.G., S.K., R.K.G., F.F.B.H., D.G.K., and M.Z.;

visualization, N.L.; supervision, P.G., S.K., and M.Z.; project administration, S.K. and M.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Princess Nourah bint Abdulrahman University Researchers Supporting Project number (PNURSP2026R675), Princess Nourah bint Abdulrahman University, Riyadh, Saudi Arabia.

Data Availability Statement: All results regenerate end to end from the code repository accompanying this manuscript, which contains the recursion solvers, the symbolic checks, the 50-digit certification, the spectral toy model, the spectral and gridpoint EnSRF implementations, the Lorenz-96 experiment, one standalone script per figure and table, the saved arrays and the verification log. The theorems are labelled T1, T2, T2′, T3 and T4 in the code and correspond to Theorems 1–5 here. The complete repository, with its scripts, logs and saved arrays, is openly deposited at Zenodo, https://doi.org/10.5281/zenodo.23530076. No observational data were used; all inputs are generated synthetically with fixed random seeds recorded in the scripts.

Acknowledgments: The authors extend their appreciation to Princess Nourah bint Abdulrahman University Researchers Supporting Project number (PNURSP2026R675), Princess Nourah bint Abdulrahman University, Riyadh, Saudi Arabia. During the preparation of this manuscript, the authors used Claude Fable 5.1 (Anthropic, San Francisco, CA, USA) for the following purposes: writing and debugging the analysis and verification code (recursion solvers, computer-algebra checks, 50-digit certification, ensemble and Lorenz-96 experiments, and figure scripts). The mathematical statements, the proofs and the numerical results were directed, reviewed and verified by the authors, who take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bi, K.; Xie, L.; Zhang, H.; Chen, X.; Gu, X.; Tian, Q. Accurate medium-range global weather forecasting with 3D neural networks. *Nature* **2023**, *619*, 533–538. https://doi.org/10.1038/s41586-023-06185-3.
2. Lam, R.; Sanchez-Gonzalez, A.; Willson, M.; Wirnsberger, P.; Fortunato, M.; Alet, F.; Ravuri, S.; Ewalds, T.; Eaton-Rosen, Z.; Hu, W.; et al. Learning skillful medium-range global weather forecasting. *Science* **2023**, *382*, 1416–1421. https://doi.org/10.1126/science.adi2336.
3. Arcucci, R.; Healy, S.; Dance, S.; Lei, L.; Bach, E.; Weaver, A.T.; Miyoshi, T.; Dillon, M.E.; Draper, C.; Schneider, R.; et al. The convergence of machine learning and data assimilation in Earth system science. *npj Artif. Intell.* **2026**, *2*, 48. https://doi.org/10.1038/s44387-026-00107-0.
4. Bach, E.; Baptista, R.; Sanz-Alonso, D.; Stuart, A. Machine learning for inverse problems and data assimilation. *arXiv* **2025**, arXiv:2410.10523. https://doi.org/10.48550/arXiv.2410.10523.
5. Bonavita, M. On some limitations of current machine learning weather prediction models. *Geophys. Res. Lett.* **2024**, *51*, e2023GL107377. https://doi.org/10.1029/2023GL107377.
6. Smith, T.A.; Penny, S.G.; Platt, J.A.; Chen, T.-C. Temporal subsampling diminishes small spatial scales in recurrent neural network emulators of geophysical turbulence. *J. Adv. Model. Earth Syst.* **2023**, *15*, e2023MS003792. https://doi.org/10.1029/2023MS003792.
7. Chattopadhyay, A.; Hassanzadeh, P. Long-term instabilities of deep learning-based digital twins of the climate system: The cause and a solution. *arXiv* **2023**, arXiv:2304.07029. https://doi.org/10.48550/arXiv.2304.07029.
8. Molinaro, R.; Lanthaler, S.; Raonić, B.; Rohner, T.; Armegioiu, V.; Simonis, S.; Grund, D.; Ramic, Y.; Wan, Z.Y.; Sha, F.; et al. Generative AI for fast and accurate statistical computation of fluids. *arXiv* **2025**, arXiv:2409.18359. https://doi.org/10.48550/arXiv.2409.18359.
9. Subich, C.; Husain, S.Z.; Separovic, L.; Yang, J. Fixing the double penalty in data-driven weather forecasting through a modified spherical harmonic loss function. In Proceedings of the 42nd International Conference on Machine Learning (ICML 2025), Vancouver, BC, Canada, 13–19 July 2025; Volume 267, pp. 57191–57211.

10. Rahaman, N.; Baratin, A.; Arpit, D.; Draxler, F.; Lin, M.; Hamprecht, F.; Bengio, Y.; Courville, A. On the spectral bias of neural networks. In Proceedings of the 36th International Conference on Machine Learning (ICML 2019), Long Beach, CA, USA, 9–15 June 2019; Volume 97, pp. 5301–5310.
11. Selz, T.; Craig, G.C. Can artificial intelligence-based weather prediction models simulate the butterfly effect? *Geophys. Res. Lett.* **2023**, *50*, e2023GL105747. https://doi.org/10.1029/2023GL105747.
12. Kotsuki, S.; Shiraishi, K.; Okazaki, A. Ensemble data assimilation to diagnose AI-based weather prediction models: A case with ClimaX version 0.3.1. *Geosci. Model Dev.* **2025**, *18*, 7215–7225. https://doi.org/10.5194/gmd-18-7215-2025.
13. Slivinski, L.C.; Whitaker, J.S.; Frolov, S.; Smith, T.A.; Agarwal, N. Assimilating observed surface pressure into ML weather prediction models. *Geophys. Res. Lett.* **2025**, *52*, e2024GL114396. https://doi.org/10.1029/2024GL114396.
14. Tian, X.; Holdaway, D.; Kleist, D. Evaluating machine learning weather models for data assimilation: Fundamental limitations in tangent linear and adjoint properties. *Geophys. Res. Lett.* **2026**, *53*, e2025GL119402. https://doi.org/10.1029/2025GL119402.
15. Carrassi, A.; Bocquet, M.; Bertino, L.; Evensen, G. Data assimilation in the geosciences: An overview of methods, issues, and perspectives. *WIREs Clim. Change* **2018**, *9*, e535. https://doi.org/10.1002/wcc.535.
16. Bannister, R.N. A review of operational methods of variational and ensemble-variational data assimilation. *Q. J. R. Meteorol. Soc.* **2017**, *143*, 607–633. https://doi.org/10.1002/qj.2982.
17. Adrian, M.; Sanz-Alonso, D.; Willett, R. Data assimilation with machine learning surrogate models: A case study with FourCastNet. *Artif. Intell. Earth Syst.* **2025**, *4*, e240050. https://doi.org/10.1175/AIES-D-24-0050.1.
18. Sanz-Alonso, D.; Waniorek, N. Long-time accuracy of ensemble Kalman filters for chaotic dynamical systems and machine-learned dynamical systems. *SIAM J. Appl. Dyn. Syst.* **2025**, *24*, 2246–2286. https://doi.org/10.1137/24M1719232.
19. Kalman, R.E. A new approach to linear filtering and prediction problems. *J. Basic Eng.* **1960**, *82*, 35–45. https://doi.org/10.1115/1.3662552.
20. Gaspari, G.; Cohn, S.E. Construction of correlation functions in two and three dimensions. *Q. J. R. Meteorol. Soc.* **1999**, *125*, 723–757. https://doi.org/10.1002/qj.49712555417.
21. Jazwinski, A.H. Stochastic Processes and Filtering Theory. In *Mathematics in Science and Engineering*; Academic Press: New York, NY, USA, 1970; Volume 64.
22. Gelb, A. (Ed.) *Applied Optimal Estimation*; MIT Press: Cambridge, MA, USA, 1974.
23. Evensen, G. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.* **1994**, *99*, 10143–10162. https://doi.org/10.1029/94JC00572.
24. Whitaker, J.S.; Hamill, T.M. Ensemble data assimilation without perturbed observations. *Mon. Weather Rev.* **2002**, *130*, 1913–1924.
25. Olver, F.W.J.; Lozier, D.W.; Boisvert, R.F.; Clark, C.W. (Eds.) *NIST Handbook of Mathematical Functions*; Cambridge University Press: New York, NY, USA, 2010; Section 8.19.
26. Tippett, M.K.; Anderson, J.L.; Bishop, C.H.; Hamill, T.M.; Whitaker, J.S. Ensemble square root filters. *Mon. Weather Rev.* **2003**, *131*, 1485–1490.
27. Houtekamer, P.L.; Mitchell, H.L. Data assimilation using an ensemble Kalman filter technique. *Mon. Weather Rev.* **1998**, *126*, 796–811. https://doi.org/10.1175/1520-0493(1998)126<0796:DAUAEK>2.0.CO;2.
28. Houtekamer, P.L.; Mitchell, H.L. A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Weather Rev.* **2001**, *129*, 123–137. https://doi.org/10.1175/1520-0493(2001)129<0123:ASEKFF>2.0.CO;2.
29. Hamill, T.M.; Whitaker, J.S.; Snyder, C. Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Weather Rev.* **2001**, *129*, 2776–2790. https://doi.org/10.1175/1520-0493(2001)129<2776:DDFOBE>2.0.CO;2.
30. Anderson, J.L.; Anderson, S.L. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Mon. Weather Rev.* **1999**, *127*, 2741–2758. https://doi.org/10.1175/1520-0493(1999)127<2741:AMCIOT>2.0.CO;2.
31. Lorenz, E.N. Predictability: A problem partly solved. In *Proceedings of the ECMWF Seminar on Predictability, Reading, UK, 4–8 September 1995*; ECMWF: Reading, UK, 1996; Volume 1, pp. 1–18.
32. Lorenz, E.N.; Emanuel, K.A. Optimal sites for supplementary weather observations: Simulation with a small model. *J. Atmos. Sci.* **1998**, *55*, 399–414. https://doi.org/10.1175/1520-0469(1998)055<0399:OSFSWO>2.0.CO;2.
33. Price, I.; Sanchez-Gonzalez, A.; Alet, F.; Andersson, T.R.; El-Kadi, A.; Masters, D.; Ewalds, T.; Stott, J.; Mohamed, S.; Battaglia, P.; et al. Probabilistic weather forecasting with machine learning. *Nature* **2025**, *637*, 84–90. https://doi.org/10.1038/s41586-024-08252-9.

34. Allen, A.; Markou, S.; Tebbutt, W.; Requeima, J.; Bruinsma, W.P.; Andersson, T.R.; Herzog, M.; Lane, N.D.; Chantry, M.; Hosking, J.S.; et al. End-to-end data-driven weather prediction. *Nature* **2025**, *641*, 1172–1179. https://doi.org/10.1038/s41586-025-08897-0.
35. Hodyss, D.; Morzfeld, M. Using diffusion models to do data assimilation. *Mon. Weather Rev.* **2026**, *154*, 165–182. https://doi.org/10.1175/MWR-D-25-0125.1.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.