Supplementary Material for

Equation-Free Screening of Mittag–Leffler Compatibility from Multi-Horizon Relaxation Profiles

N'Gbo N'Gbo, Andrei Velichko, and Bruno Carpentieri

This document collects the optimization diagnostics, the complete results of the comparison of profile constructions and of the noisy specificity analysis, the stable sensitivity results, the numerical verification of the single-record example, and the superseded analysis of that record, retained for reproducibility. Equation, table, and figure numbers without the prefix "S" refer to the main text. Scripts, CSV files, and optimizer logs are archived at `https://doi.org/10.5281/zenodo.23494805`.

S1 Optimization diagnostics

All fits use the settings of Table 1. For the focused comparison of Table 5, the exponential fit uses one start, the Mittag–Leffler fit six starts, and the biexponential fit eighteen starts. The spread of the final Mittag–Leffler objectives across starts is 4.0×10^{-4}, 4.4×10^{-4}, and 1.6×10^{-20} for Cases A, B, and C; the corresponding biexponential spreads are 7.12×10^{-8}, 7.21×10^{-8}, and 7.88×10^{-10}, with 12/18, 9/18, and 3/18 starts terminating with `exitflag`$= 1$. In every case, the start with the lowest final objective is retained.

Table S1: Per-start biexponential optimizer diagnostics for Cases A and B at $\alpha = 0.8$ and Case C at $s = 3$ (main run). w: weight of the slow component in the starting amplitudes, $C_{1,0} = wG_1$, $C_{2,0} = (1 - w)G_1$. The objective is the mean squared residual on the fitting horizons; the retained start is shown in bold. Exit flag 1: converged; 0: `MaxIter` or `MaxFunEvals` reached.

Case	Start	w	$\lambda_{1,0}$	$(\lambda_2 - \lambda_1)_0$	Final objective	Exit flag	Iter. / F-evals
A	1	0.3	0.02	0.1	7.20×10^{-8}	1	566 / 970
	2	0.3	0.02	0.3	2.63×10^{-9}	0	2000 / 3254
	3	0.3	0.02	1.0	7.48×10^{-10}	1	763 / 1297
	4	0.3	0.05	0.1	7.20×10^{-8}	1	789 / 1364
	5	0.3	0.05	0.3	5.78×10^{-8}	0	2000 / 3260
	6	0.3	0.05	1.0	7.48×10^{-10}	1	734 / 1251
	7	0.3	0.1	0.1	4.67×10^{-8}	0	2000 / 3240
	8	0.3	0.1	0.3	5.29×10^{-8}	0	2000 / 3285
	9	0.3	0.1	1.0	7.48×10^{-10}	1	897 / 1534
	10	0.7	0.02	0.1	7.48×10^{-10}	1	1451 / 2428
	11	0.7	0.02	0.3	$\mathbf{7.48 \times 10^{-10}}$	1	1884 / 3163
	12	0.7	0.02	1.0	7.29×10^{-9}	1	363 / 648
	13	0.7	0.05	0.1	4.18×10^{-8}	0	2000 / 3271
	14	0.7	0.05	0.3	7.48×10^{-10}	1	1826 / 3036
	15	0.7	0.05	1.0	7.48×10^{-10}	1	753 / 1269
	16	0.7	0.1	0.1	7.48×10^{-10}	0	2000 / 3246
	17	0.7	0.1	0.3	7.48×10^{-10}	1	1621 / 2685
	18	0.7	0.1	1.0	7.48×10^{-10}	1	680 / 1153
B	1	0.3	0.02	0.1	6.33×10^{-9}	0	2000 / 3261
	2	0.3	0.02	0.3	1.13×10^{-9}	0	2000 / 3274
	3	0.3	0.02	1.0	7.99×10^{-10}	1	939 / 1602
	4	0.3	0.05	0.1	7.29×10^{-8}	1	618 / 1089
	5	0.3	0.05	0.3	4.33×10^{-8}	0	2000 / 3283
	6	0.3	0.05	1.0	7.99×10^{-10}	1	676 / 1153
	7	0.3	0.1	0.1	1.48×10^{-8}	0	2000 / 3279
	8	0.3	0.1	0.3	3.64×10^{-8}	0	2000 / 3274
	9	0.3	0.1	1.0	7.99×10^{-10}	1	843 / 1456
	10	0.7	0.02	0.1	1.56×10^{-8}	0	2000 / 3304
	11	0.7	0.02	0.3	7.99×10^{-10}	1	1860 / 3086
	12	0.7	0.02	1.0	7.99×10^{-10}	1	708 / 1214
	13	0.7	0.05	0.1	2.01×10^{-8}	0	2000 / 3234
	14	0.7	0.05	0.3	7.99×10^{-10}	1	1877 / 3158
	15	0.7	0.05	1.0	7.99×10^{-10}	1	694 / 1208
	16	0.7	0.1	0.1	7.99×10^{-10}	0	2000 / 3282
	17	0.7	0.1	0.3	5.47×10^{-9}	0	2000 / 3271
	18	0.7	0.1	1.0	$\mathbf{7.99 \times 10^{-10}}$	1	699 / 1196
C	1	0.3	0.02	0.1	2.70×10^{-10}	0	2000 / 3260
	2	0.3	0.02	0.3	4.90×10^{-10}	0	2000 / 3257
	3	0.3	0.02	1.0	7.20×10^{-10}	0	2000 / 3244
	4	0.3	0.05	0.1	9.87×10^{-11}	0	2000 / 3259
	5	0.3	0.05	0.3	2.07×10^{-10}	0	2000 / 3310
	6	0.3	0.05	1.0	7.87×10^{-10}	0	2000 / 3248
	7	0.3	0.1	0.1	2.65×10^{-10}	1	1963 / 3248
	8	0.3	0.1	0.3	1.75×10^{-10}	0	2000 / 3272
	9	0.3	0.1	1.0	7.88×10^{-10}	0	2000 / 3267
	10	0.7	0.02	0.1	8.19×10^{-11}	0	2000 / 3273
	11	0.7	0.02	0.3	1.02×10^{-13}	0	2000 / 3238
	12	0.7	0.02	1.0	7.72×10^{-10}	0	2000 / 3249
	13	0.7	0.05	0.1	2.70×10^{-11}	0	2000 / 3248
	14	0.7	0.05	0.3	1.12×10^{-11}	0	2000 / 3238
	15	0.7	0.05	1.0	7.88×10^{-10}	0	2000 / 3266
	16	0.7	0.1	0.1	4.88×10^{-14}	1	996 / 1690
	17	0.7	0.1	0.3	$\mathbf{4.88 \times 10^{-14}}$	1	1095 / 1838
	18	0.7	0.1	1.0	7.04×10^{-10}	0	2000 / 3250

Table S2: Fitted parameters for the expanded candidate comparison of Table 5 (mean ± standard deviation over ten seeds). Exponential: $G = Ce^{-\lambda t}$; Mittag–Leffler: $G = CE_\alpha(-\lambda t^\alpha)$; biexponential: $G = C_1 e^{-\lambda_1 t} + C_2 e^{-\lambda_2 t}$, $\lambda_1 < \lambda_2$.

Case	Model	C or C_1	λ or λ_1	α_{fit}	C_2	λ_2
A	Exponential	1.0614 ± 0.0002	$0.1122 \pm 3 \times 10^{-10}$	–	–	–
	Mittag–Leffler	1.1116 ± 0.0002	$0.1500 \pm 4 \times 10^{-12}$	$0.8000 \pm 2 \times 10^{-11}$	–	–
	Biexponential	0.9901 ± 0.0002	$0.0887 \pm 4 \times 10^{-10}$	–	$0.1036 \pm 2 \times 10^{-5}$	$0.8842 \pm 2 \times 10^{-8}$
B	Exponential	1.0582 ± 0.0020	0.1078 ± 0.0030	–	–	–
	Mittag–Leffler	1.1128 ± 0.0037	0.1492 ± 0.0044	0.7789 ± 0.0010	–	–
	Biexponential	0.9810 ± 0.0012	0.0824 ± 0.0022	–	0.1109 ± 0.0037	0.8629 ± 0.0027
C	Exponential	1.0974 ± 0.0056	0.1565 ± 0.0068	–	–	–
	Mittag–Leffler	1.1181 ± 0.0069	0.1715 ± 0.0078	0.9330 ± 0.0027	–	–
	Biexponential	0.7357 ± 0.0287	0.0997 ± 0.0022	–	0.3734 ± 0.0346	0.3139 ± 0.0060

S2 Comparison of profile constructions

The kNN, constant training-mean, and direct ensemble-average profiles (Section 2.3) are passed through the same exponential-versus–Mittag–Leffler fitting and horizon-extrapolation pipeline in every noise-free configuration of Cases A–C, with the same ten seeds. Table S3 extends the representative Table 6 of the main text to all thirteen configurations. The three constructions give the same outcome in every seed of every configuration; in Case A at $\alpha = 1$, all ten seeds are ties under each construction.

Table S3: Comparison of profile constructions in all thirteen noise-free configurations (ten seeds each). α_{fit}: mean fitted order for the kNN, constant training-mean (Const), and direct ensemble-average (Direct) profiles. W/L/T: Mittag–Leffler wins, losses, and ties on the held-out horizons for the kNN profile. Agreement: seeds in which Const and Direct give the same outcome as kNN. E: mean held-out individual-trajectory error, Equation (2.12) of the main text, for the kNN and constant predictors; it is not defined for the direct average, which uses the observed futures.

		α_{fit}					E
Case	Parameter	kNN	Const	Direct	W/L/T	Agreement	kNN / Const
A	$\alpha = 0.6$	0.600	0.600	0.600	10/0/0	10/10	0.0026 / 0.348
A	$\alpha = 0.7$	0.700	0.700	0.700	10/0/0	10/10	0.0025 / 0.334
A	$\alpha = 0.8$	0.800	0.800	0.800	10/0/0	10/10	0.0024 / 0.318
A	$\alpha = 0.9$	0.900	0.900	0.900	10/0/0	10/10	0.0022 / 0.303
A	$\alpha = 1$	1.000	1.000	1.000	0/0/10	10/10	0.0021 / 0.287
B	$\alpha = 0.6$	0.583	0.580	0.580	10/0/0	10/10	0.0286 / 0.360
B	$\alpha = 0.7$	0.681	0.677	0.678	10/0/0	10/10	0.0336 / 0.349
B	$\alpha = 0.8$	0.779	0.775	0.775	10/0/0	10/10	0.0396 / 0.339
B	$\alpha = 0.9$	0.878	0.873	0.874	10/0/0	10/10	0.0466 / 0.329
B	$\alpha = 1$	0.977	0.972	0.973	10/0/0	10/10	0.0531 / 0.320
C	$s = 3$	0.933	0.924	0.926	10/0/0	10/10	0.0634 / 0.354
C	$s = 5$	0.847	0.832	0.836	10/0/0	10/10	0.0695 / 0.416
C	$s = 10$	0.668	0.659	0.667	10/0/0	10/10	0.0769 / 0.499

S3 Noisy specificity analysis

All noisy settings use the noise model of Equation (2.19), the same fifty seeds, fits on horizons 1–20, and evaluation on horizons 21–30. In the binary comparison (Table S4), a seed is a tie when $|\Delta\text{RMSE}| < 10^{-8}$ or $|\alpha_{\text{fit}} - 1| < 10^{-6}$; boundary hits are counted separately, as seeds with $|\alpha_{\text{fit}} - 1| < 10^{-6}$, so that every boundary hit is a tie under the binary rule but not conversely.

The sign and Wilcoxon tests exclude ties, are two-sided, and are not adjusted for multiple comparisons; they are reported as supporting, not confirmatory, analyses. In the three-family comparison (Table 8 of the main text), the family with the lowest held-out RMSE is selected and families within 10^{-8} of the lowest RMSE are tied; the convention $|\alpha_{\text{fit}} - 1| < 10^{-6}$ applies only to the exponential/Mittag–Leffler pair. Table S5 reports all paired differences, and Table S6 the biexponential optimizer diagnostics.

Table S4: Binary exponential-versus–Mittag–Leffler comparison under observational noise ($\sigma_j = 0.03|x_{0,j}|$, fifty seeds). Single-rate: Case A generator. W/L/T: Mittag–Leffler wins, losses, and ties under the binary rule of Section 2.5 of the main text. Selection frequencies of the Mittag–Leffler model on the held-out horizons and by descriptive AIC, with 95% Wilson intervals. Exact sign test and Wilcoxon signed-rank test (normal approximation with tie correction) on the paired differences, excluding ties; two-sided, unadjusted. Boundary: seeds with $|\alpha_{\text{fit}} - 1| < 10^{-6}$.

Setting	W/L/T	ML, held-out	ML, AIC	Sign p	Wilcoxon p	Boundary
Single-rate, $\alpha = 0.6$	50/0/0	1.00 [0.93, 1.00]	1.00 [0.93, 1.00]	1.8×10^{-15}	7.6×10^{-10}	0
Single-rate, $\alpha = 0.7$	49/1/0	0.98 [0.90, 1.00]	1.00 [0.93, 1.00]	9.1×10^{-14}	8.0×10^{-10}	0
Single-rate, $\alpha = 0.8$	49/1/0	0.98 [0.90, 1.00]	1.00 [0.93, 1.00]	9.1×10^{-14}	8.0×10^{-10}	0
Single-rate, $\alpha = 0.9$	47/2/1	0.94 [0.84, 0.98]	0.84 [0.71, 0.92]	4.4×10^{-12}	2.3×10^{-8}	1
Single-rate, $\alpha = 1$	5/15/30	0.10 [0.04, 0.21]	0.04 [0.01, 0.13]	0.0414	0.0072	30
Classical B, $\alpha = 1$	30/10/10	0.60 [0.46, 0.72]	0.22 [0.13, 0.35]	0.0022	0.0248	10
Classical C, $s = 3$	46/4/0	0.92 [0.81, 0.97]	0.78 [0.65, 0.87]	4.5×10^{-10}	2.8×10^{-9}	0

Table S5: Paired held-out RMSE differences in the noisy three-family comparison (fifty seeds), mean with 95% Student t confidence interval, in units of 10^{-3}. $\Delta_{a,b} = \text{RMSE}_a^* - \text{RMSE}_b^*$; negative values favor family a. Boundary: seeds with $|\alpha_{\text{fit}} - 1| < 10^{-6}$.

Setting	$\Delta_{\text{Exp,ML}}$	$\Delta_{\text{Exp,Bi}}$	$\Delta_{\text{ML,Bi}}$	Boundary
Classical B, $\alpha = 1$	0.39 [−0.01, 0.79]	0.25 [−0.24, 0.74]	−0.14 [−0.36, 0.08]	10
Classical C, $s = 3$	3.45 [2.86, 4.03]	2.76 [2.02, 3.49]	−0.69 [−1.17, −0.22]	0
Fractional B, $\alpha = 0.8$	8.90 [8.25, 9.55]	7.33 [6.39, 8.28]	−1.56 [−2.16, −0.97]	0

Table S6: Biexponential optimizer diagnostics in the noisy three-family comparison (eighteen starts per seed, fifty seeds). Converged: starts terminating with `exitflag`$= 1$, median / minimum over seeds. No converged start: seeds in which no start terminated with `exitflag`$= 1$. At best: minimum over seeds of the number of starts reaching the best objective within a relative tolerance of 10^{-6}. Spread: difference between the largest and smallest final objectives, median / maximum over seeds. Degenerate: fits with a normalized component weight or relative rate gap below 10^{-3}.

Setting	Converged	No converged start	At best	Spread	Degenerate
Classical B, $\alpha = 1$	6 / 0	1	1	3.8×10^{-8} / 2.9×10^{-6}	9
Classical C, $s = 3$	5.5 / 0	2	1	8.0×10^{-8} / 2.3×10^{-6}	2
Fractional B, $\alpha = 0.8$	12 / 5	0	4	1.1×10^{-7} / 6.3×10^{-6}	0

The diagnostics in Table S6 are reported without discarding any fit. A start that stops at the iteration or function-evaluation limit (`exitflag`$= 0$) does not by itself make the retained fit or the resulting ranking incorrect, and a small spread of final objectives does not by itself establish convergence. In the two classical settings, a few seeds have no start terminating with `exitflag`$= 1$, and several biexponential fits are degenerate, in which case the biexponential

coincides numerically with a single exponential; these features should be taken into account when interpreting the biexponential selections in those settings.

S4 Stable sensitivity

Table S7: One-at-a-time sensitivity analysis on Case B at $\alpha = 0.8$. Each row changes a single setting relative to the baseline ($K = 3$, $m_{\mathrm{hist}} = 7$, fitting horizons 1–20, $\Delta t = 0.1$). $\Delta\mathrm{RMSE} = \mathrm{RMSE}^*_{\mathrm{exp}} - \mathrm{RMSE}^*_{\mathrm{ML}}$ on the held-out horizons, in units of 10^{-3}; positive values favor the Mittag–Leffler model. $\Delta\mathrm{AIC} = \mathrm{AIC}_{\mathrm{ML}} - \mathrm{AIC}_{\mathrm{exp}}$ on the fitting horizons. Mean $\pm$ standard deviation over the same ten seeds in every configuration.

Setting	Value	α_{fit}	ΔRMSE	ΔAIC	ML wins
Baseline	–	0.7789 ± 0.0010	11.89 ± 0.31	-198.3 ± 2.0	10/10
K	1	0.7780 ± 0.0017	11.94 ± 0.30	-196.8 ± 3.1	10/10
	5	0.7791 ± 0.0015	11.88 ± 0.29	-198.7 ± 2.7	10/10
	10	0.7790 ± 0.0011	11.88 ± 0.30	-198.5 ± 2.0	10/10
m_{hist}	4	0.7823 ± 0.0010	14.19 ± 0.34	-191.3 ± 2.4	10/10
	10	0.7756 ± 0.0017	10.28 ± 0.29	-203.4 ± 3.0	10/10
Fitting horizons	1–15	0.7806 ± 0.0009	15.12 ± 0.39	-154.5 ± 1.5	10/10
	1–25[a]	0.7774 ± 0.0010	9.46 ± 0.26	-241.0 ± 2.7	10/10
Δt[b]	0.05	0.7792 ± 0.0008	12.09 ± 0.32	-396.4 ± 3.3	10/10
	0.2	0.7782 ± 0.0011	11.45 ± 0.29	-99.6 ± 1.2	10/10

[a] Only five held-out horizons (26–30).

[b] Physical history, forecast, and fitting windows are held fixed (Section 2.6): $m_{\mathrm{hist}} = 13$, $H_{\max} = 60$, fitting horizons 1–40 at $\Delta t = 0.05$; $m_{\mathrm{hist}} = 4$, $H_{\max} = 15$, fitting horizons 1–10 at $\Delta t = 0.2$. ΔAIC scales with the number of fitted points and is comparable only between rows with the same fitting range.

S5 Numerical verification of the single-record example

System (2.25) with $a = -2$, $b = -6.4$, $c = 1$, $d = 1$, and $\mathbf{x}_0 = (0.2, 0.2, 0.2)^{\mathsf{T}}$ is integrated to $T = 200$ with Garrappa's `fde12` implementation of the Adams–Bashforth–Moulton method, in predictor-only mode and in predictor–corrector (PECE) mode (Section 2.7). The diagnostics refer to the x-component: standard deviations over $t \geq 40$ and over $180 \leq t \leq 200$, and the final distance from the nearest nonzero equilibrium $E_{\pm} = (0, \pm\sqrt{6.4}, 6.4)$.

Table S8: Reproduction of the original explicit runs by the predictor-only mode of `fde12`: post-transient standard deviation of x at $\alpha = 0.916$.

Δt	Original explicit scheme	Predictor-only `fde12`	Difference
0.0100	2.3454	2.345426	2.6×10^{-5}
0.0050	2.3529	2.352849	5.1×10^{-5}
0.0025	1.4168	1.416763	3.7×10^{-5}

Table S9: Predictor–corrector (PECE) refinement at $\alpha = 0.916$. RMS difference: root-mean-square difference of x between the PECE solutions at Δt and $\Delta t/2$ on their common grid over $40 \le t \le 200$. Iterated corrector: corrector iterated to convergence at $\Delta t = 0.005$. Control: PECE at $\alpha = 0.93$, $\Delta t = 0.005$.

Run	SD, $t \ge 40$	SD, $180 \le t \le 200$	Final distance	RMS difference
PECE, $\Delta t = 0.01$	0.92869	0.27639	0.49166	0.0336 (0.01 vs 0.005)
PECE, $\Delta t = 0.005$	0.92554	0.27416	0.48882	0.0084 (0.005 vs 0.0025)
PECE, $\Delta t = 0.0025$	0.92474	0.27359	0.48800	0.0022 (0.0025 vs 0.00125)
PECE, $\Delta t = 0.00125$	0.92453	0.27345	0.48778	–
Iterated corrector, $\Delta t = 0.005$	0.92443	0.27340	0.48765	–
Control, $\alpha = 0.93$	2.44263	2.41959	–	–

At $E_{\pm}$, the Jacobian of System (2.25) has eigenvalues $\lambda_1 \approx -3.481547$ and $\lambda_{2,3} \approx 0.240773 \pm 1.902250\,\mathrm{i}$, so that $\alpha_c = (2/\pi)\,|\arg \lambda_2| \approx 0.9198474$. The equilibria are locally asymptotically stable for $\alpha < \alpha_c$ and unstable for $\alpha > \alpha_c$. The ratios of successive RMS differences (about 4.0 and 3.9) are finite-step numerical evidence of consistent refinement over the simulated interval, not a statement about an asymptotic order of convergence; the local stability result does not determine the global dynamics.

S6 Superseded analysis of the single-record example

Status of this section. The following analyses were performed on trajectories generated by the original explicit discretization. Subsequent independent numerical verification (Section S5) showed that the observed sustained regime is not robust under refinement. These results are retained for reproducibility and are not used to support conclusions about the continuous fractional system.

S6.1 Original protocol

The original record was generated with the explicit product-integration scheme at $\Delta t = 0.01$ and observed as the x-component at $\Delta t_{\mathrm{obs}} = 0.05$ after the first 40 time units. Delay vectors with embedding dimension $m = 8$ and lag $\tau_{\mathrm{embed}} = 2$ were formed, and direct-horizon kNN predictions ($K = 3$) were made for $h = 1, \ldots, 600$, that is, $\tau = h\,\Delta t_{\mathrm{obs}} \le 30$. Training and test reference vectors were separated by one cut in physical time (70% training allocation) with an embargo of $H_{\max} + (m-1)\tau_{\mathrm{embed}} = 614$ samples, so that no training target fell inside any test history. The forecast-error growth profile $G(\tau)$ was the geometric mean of the absolute forecast errors over the test references, normalized by its value at $\tau = \Delta t_{\mathrm{obs}}$. The pre-transition interval $\tau \le \tau^*$ was selected by a two-segment linear representation of $\log G$, scored by the mean absolute residual, with at least 20 profile points on each side. Over this interval, the exponential $C\exp(\lambda\tau)$, Mittag–Leffler $CE_\alpha(\lambda\tau^\alpha)$, and logistic $C/\{1+\exp[-k(\tau-\tau_0)]\}$ growth laws were fitted by minimizing $Q_{\log} = \frac{1}{n}\sum_i [\log G_i - \log \widehat{G}(\tau_i)]^2$, with the settings of Table S10.

Table S10: Optimization settings of the original single-record fits (same `fminsearch` options as Table 1 of the main text).

Model	Parameterization and admissible region	Initialization actually used
Exponential	$C = e^{p_1}$, $\lambda = e^{p_2}$; $C, \lambda > 0$ by construction.	One start: $C_0 = \max(G_1, 10^{-8})$, $\lambda_0 = 0.1$.
Mittag–Leffler	$C = e^{p_1}$, $\lambda = e^{p_2}$, $\alpha = e^{p_3}$. The log transform enforces only positivity; $0.05 \leq \alpha \leq 1$ is imposed separately by objective penalty 10^{100}.	Six paired starts, all with $C_0 = \max(G_1, 10^{-8})$: $(\lambda_0, \alpha_0) = (0.05, 0.2), (0.10, 0.5), (0.15, 0.7), (0.20, 0.85), (0.25, 0.95), (0.30, 1.0)$.
Logistic	$C = e^{p_1}$, $k = e^{p_2}$, $\tau_0 = e^{p_3}$; thus $C, k, \tau_0 > 0$ by construction, with no additional explicit bounds.	Four starts, with $G_{\max} = \max_i G_i$ and $\tau_{\max}$ the last pre-transition fitting time: $(G_{\max}, 0.5, 0.50\tau_{\max})$, $(G_{\max}, 1, 0.75\tau_{\max})$, $(2G_{\max}, 1, \tau_{\max})$, $(G_{\max}, 2, 0.25\tau_{\max})$.

S6.2 Split and profile of the original record

The original record is generated from the fractional system of Equation (2.25) with $\alpha = 0.916$, integrated to $T = 200$ with solver step $\Delta t = 0.01$. After removal of the first 40 time units, the x-component is sampled at $\Delta t_{\text{obs}} = 0.05$. Training and test reference vectors are separated by a single cut in physical time with an embargo of $H_{\max} + L_{\text{emb}} = 614$ observed samples (Section S6.1). Table S11 lists the resulting index ranges and counts for the baseline 70% allocation.

Table S11: Physical-time split of the original record (baseline, 70% training allocation). Indices refer to the observed samples x_k retained after the transient.

Quantity	Value
Observed samples after transient	$N = 3201$
Embedding dimension / lag	$m = 8$, $\tau_{\text{embed}} = 2$
Embedding span L_{emb}	14 samples
Maximum horizon $H_{\max}$	600 samples (30 time units)
Embargo	614 samples (30.7 time units)
Cut time t_{cut} / index k_{cut}	139.70 / 1995
Training reference indices	15–1395 (1381 vectors)
Test reference indices	2010–2601 (592 vectors)
Transition τ^* (fitted points)	5.20 (104)

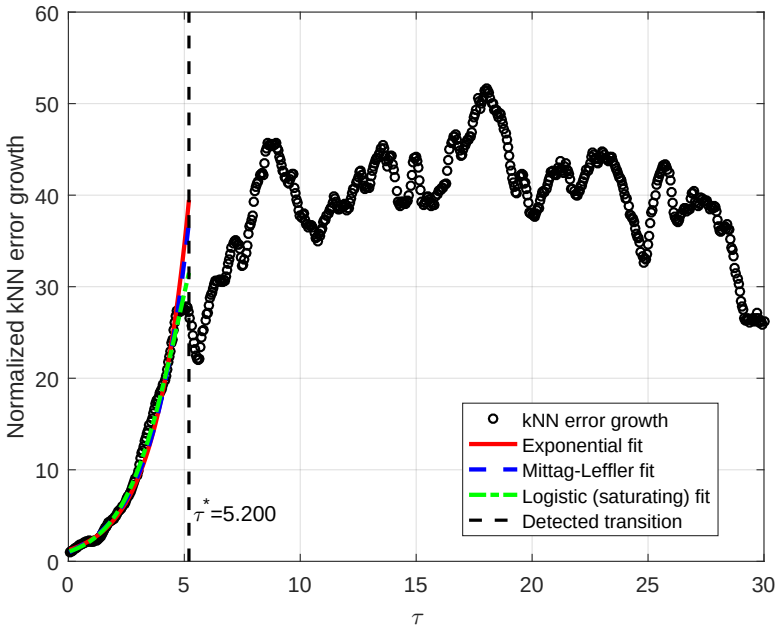


Figure S1: kNN forecast-error growth profile of the original record (baseline split) with the exponential, Mittag–Leffler, and logistic laws fitted over the pre-transition interval $\tau \leq \tau^* = 5.20$ (104 points).

Table S12: Pre-transition model comparison for the original kNN forecast-error growth profile using the physical-time split with embargo. All models are fitted over $\tau \leq \tau^* = 5.20$ (104 points). Here ΔAIC is reported relative to the lowest-AIC model, so the best model has $\Delta\text{AIC} = 0$ and larger values are worse; this differs from the pairwise signed convention used in the stable-case tables. The criterion is used descriptively.

Model	$\text{RMSE}_{\log}$	RMSE	ΔAIC
Exponential	0.114	2.439	93.7
Mittag–Leffler	0.101	1.926	69.4
Logistic	0.072	0.847	0

Table S13: Optimizer diagnostics for the original pre-transition fits at the three solver steps.

Solver step Δt	Model	Starts	Converged	Spread of final objectives
0.0100	Mittag–Leffler	6	6/6	3.5×10^{-18}
	Logistic	4	4/4	4.3×10^{-18}
0.0050	Mittag–Leffler	6	6/6	6.9×10^{-18}
	Logistic	4	$0/4^a$	2.2×10^{-9}
0.0025	Mittag–Leffler	6	6/6	1.4×10^{-17}
	Logistic	4	$4/4^b$	9.7×10^{-17}

[a] All four logistic starts stopped at the iteration limit (`MaxIter`$= 2000$) at practically the same degenerate solution, with final objectives differing by at most 2.2×10^{-9}: $C \approx (6\text{–}8) \times 10^6$, $\tau_0 \approx 22$, far outside the fitting window $\tau \leq \tau^* = 4.35$, and $k \approx 0.697$, equal to the exponential rate. Over the fitting window the logistic law therefore reduces to the exponential.

[b] All four starts converged to the same objective with $\tau_0 \to 0$ ($\tau_0 < 10^{-9}$), at the boundary of the log-parameterization, so that the fitted curve reduces to $C/(1 + e^{-k\tau})$.

S6.3 Original split sensitivity

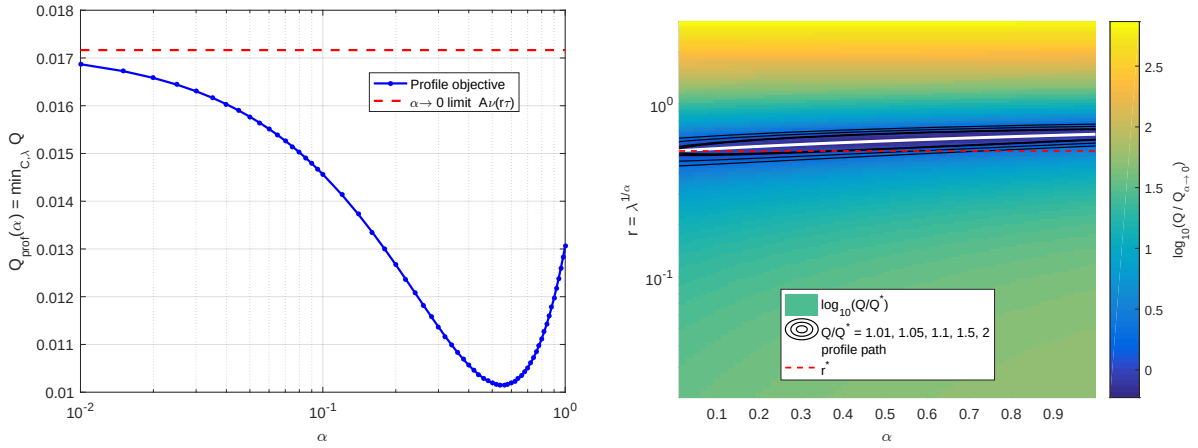
Table S14: Sensitivity of the original pre-transition comparison to the training allocation of the physical-time split (embargo of 614 samples, solver step $\Delta t = 0.01$). All models minimize $Q_{\log}$; $\mathrm{RMSE}_{\log}$ is computed on $\log G$ and RMSE on G over $\tau \le \tau^*$. ΔAIC is relative to the best model in each column.

Training allocation	60%	70%	80%
Cut time t_{cut}	129.85	139.70	149.55
Training / test vectors	1184 / 789	1381 / 592	1578 / 395
Transition τ^* (fitted points)	5.40 (108)	5.20 (104)	4.55 (91)
Exponential			
C, λ	1.404, 0.561	1.220, 0.669	1.187, 0.692
$\mathrm{RMSE}_{\log}$ / RMSE	0.128 / 2.157	0.114 / 2.439	0.106 / 1.753
ΔAIC	71.9	93.7	33.2
Mittag–Leffler			
C, λ	0.860, 0.724	0.825, 0.767	0.925, 0.749
α_{fit}	0.463	0.550	0.665
$\mathrm{RMSE}_{\log}$ / RMSE	0.110 / 1.740	0.101 / 1.926	0.099 / 1.542
ΔAIC	42.7	69.4	22.7
Logistic			
C, k, τ_0	39.106, 0.715, 4.827	55.809, 0.810, 4.879	53.120, 0.808, 4.807
$\mathrm{RMSE}_{\log}$ / RMSE	0.091 / 1.074	0.072 / 0.847	0.087 / 1.091
ΔAIC	0	0	0

The logistic curve is the lowest-AIC family for all three allocations. The fitted Mittag–Leffler order remains interior but varies between 0.46 and 0.67, increasing as the detected transition τ^* decreases.

S6.4 Identifiability of the original fitted order

For the baseline fit of the original record ($\tau \le \tau^* = 5.20$, 104 points), the profile objective $Q_{\mathrm{prof}}(\alpha) = \min_{C,\lambda} Q_{\log}(C, \lambda, \alpha)$ is computed by eliminating C in closed form and re-optimizing λ at each α. The profile has an interior minimum at $\alpha \approx 0.55$ ($Q_{\mathrm{prof}} = 0.0102$) and increases toward both limits of the Mittag–Leffler family: the exponential model at $\alpha = 1$ ($Q = 0.0131$) and the limit $\alpha \to 0^+$, in which $\alpha E_\alpha(x^\alpha)$ tends to the Volterra function $\nu(x)$ [1], giving the model $A\,\nu(r\tau)$ with $A = C/\alpha$ and $r = \lambda^{1/\alpha}$ ($Q = 0.0172$). The fitted order is therefore an interior optimum and does not depend on the lower bound imposed on α (Figure S2a). The objective map in (α, r) shows a narrow valley in which r varies only between about 0.54 and 0.67 across the whole range of α (Figure S2b): the data constrain the transformed rate parameter r much more tightly than the order, which is therefore only weakly localized.



(a) (b)

Figure S2: Identifiability of the Mittag–Leffler order for the baseline original fit ($\tau \leq \tau^* = 5.20$, 104 points). (a) Profile objective $Q_{\mathrm{prof}}(\alpha)$, with the $\alpha \to 0^+$ Volterra limit. (b) Objective map in (α, r), $r = \lambda^{1/\alpha}$.

S6.5 Solver-step diagnostics

Table S15: Solver-step comparison for the original fractional system ($\alpha = 0.916$, $T = 200$, $\Delta t_{\mathrm{obs}} = 0.05$, same split and pre-transition fitting procedure). Mean, SD, skewness, and kurtosis refer to the marginal distribution of the observed signal $x(t)$; λ_1 is the early slope of $\log G(\tau)$ (first segment of the two-segment representation); ΔAIC is relative to the lowest-AIC model.

Δt	Mean	SD	Skew.	Kurt.	τ^*	λ_1	α_{fit}	$\Delta\mathrm{AIC}_{\mathrm{exp}}$ / $\Delta\mathrm{AIC}_{\mathrm{ML}}$	Best
0.0100	−0.0085	2.3454	0.0311	2.4662	5.20	0.6686	0.5499	93.65 / 69.41	Logistic
0.0050	−0.0109	2.3529	0.0134	2.4501	4.35	0.6970	0.9063	0.00 / 0.58	Exponential
0.0025	−0.0101	1.4168	0.2052	2.4323	8.70	0.0431	0.7413	0.00 / 1.90	Exponential

At the two refined solver steps, the exponential has the lowest descriptive AIC, although its separation from the Mittag–Leffler fit is small (0.58 and 1.90 AIC units, respectively). These differences are reported descriptively and are not interpreted as formal evidence of model separation. At both steps the logistic fit is degenerate (Table S13, notes a and b).

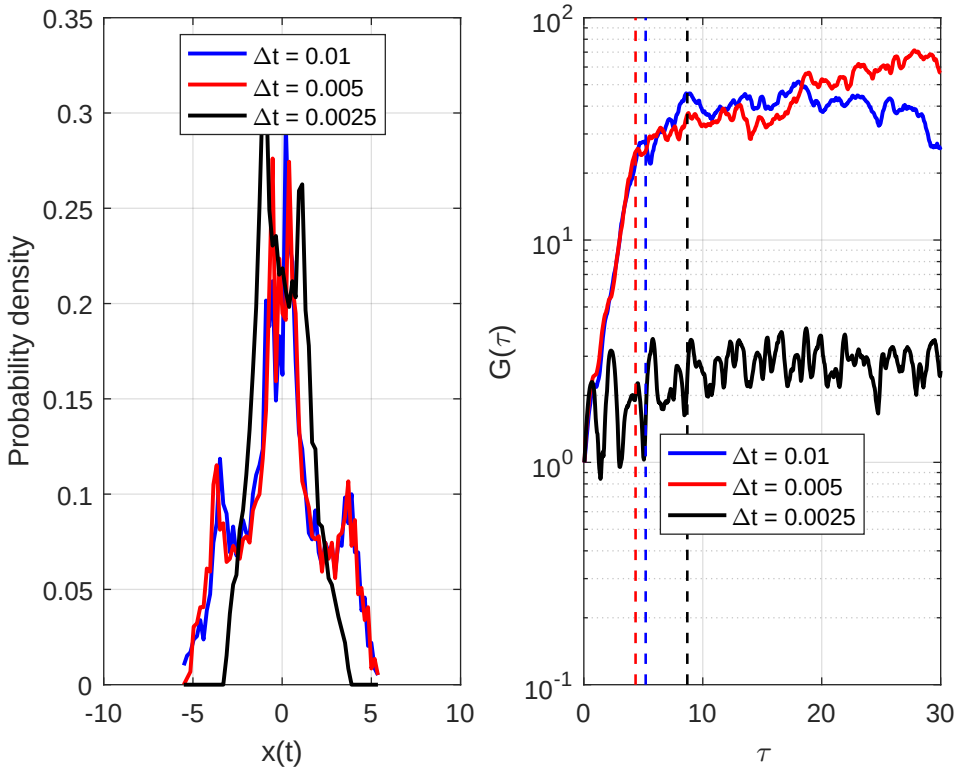


Figure S3: Solver-step comparison for the original fractional system ($\alpha = 0.916$, $T = 200$, $\Delta t_{\text{obs}} = 0.05$) for $\Delta t = 0.01$, 0.005, and 0.0025. Left: marginal distribution of the observed signal $x(t)$. Right: normalized forecast-error growth profile $G(\tau)$; the dashed vertical lines mark the detected transition times τ^*.

S6.6 Fitting-window and error-metric sensitivity

The original fits were also repeated on fixed fitting windows $\tau \leq \tau_{\text{end}}$, $\tau_{\text{end}} = 4, 5, 6$, and on a profile built from the mean absolute forecast error instead of the geometric mean (Table S16). The fitted order moves from 0.96 to the lower bound 0.05, and the selected family changes, as the window or the error metric changes.

Table S16: Original record, $\Delta t = 0.01$: sensitivity of the pre-transition fits to the fitting window and to the error metric. Geometric: geometric-mean forecast error (baseline profile); MAE: mean absolute forecast error. ΔAIC relative to the best model (log domain). *At the lower bound $\alpha_{\min} = 0.05$.

Profile	Window	n	α_{fit}	$\text{RMSE}_{\text{ML}}/\text{RMSE}_{\text{exp}}$	ΔAIC (Exp/ML/Log)	Best (AIC)
Geometric	$\tau^* = 5.2$ (baseline)	104	0.550	0.790	93.7 / 69.4 / 0	Logistic
	$\tau_{\text{end}} = 4$	80	0.961	0.983	0 / 1.7 / 2.0	Exponential
	$\tau_{\text{end}} = 5$	100	0.651	0.794	62.9 / 44.8 / 0	Logistic
	$\tau_{\text{end}} = 6$	120	0.050*	0.749	187.0 / 143.2 / 0	Logistic
MAE	own $\tau^* = 3.7$	74	0.363	0.703	63.0 / 25.4 / 0	Logistic
	$\tau^* = 5.2$	104	0.050*	0.681	258.9 / 185.9 / 0	Logistic

References

[1] Alexander Apelblat. *Integral Transforms and Volterra Functions*. Nova Science Publishers, New York, 2010.