

Article

Not peer-reviewed version

Virialisation as Viscosity: Deriving the ITP Cosmological Memory Kernel from TNG300

[Stephen Atalebe](#) *

Posted Date: 27 January 2026

doi: 10.20944/preprints202601.2116.v1

Keywords: scale dependence of memory; ITP memory kernel; non-Markovian; TNG300/TNG50 simulations; coarse-graining



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

Article

Virialisation as Viscosity: Deriving the ITP Cosmological Memory Kernel from TNG300

Stephen Atalebe

Masaryk University, Brno, Czech Republic; 478944@muni.cz

Abstract

Non-Markovian extensions of Λ CDM model the backreaction of nonlinear structure on the background expansion through a memory kernel that relates deviations in the effective expansion rate to a structural source built from small-scale dynamics. In the Infinite Transformation Principle (ITP) framework this is written as a Volterra equation for the deviation $\delta H^2(t)$ of the effective expansion rate from the Friedmann prediction, with a kernel K_{ITP} and a structural source $\Sigma(t)$. In previous work, K_{ITP} was introduced phenomenologically and constrained by fits to $H(z)$ and growth data, leading to the criticism that the kernel was not derived from first principles. This paper takes a first step toward a simulation-based derivation. The TNG300-1 simulation is coarse-grained into a $4 \times 4 \times 4$ grid of $\sim 50 \text{ Mpc}/h$ domains. For each snapshot between $z \simeq 2$ and $z = 0$ the analysis constructs a domain-averaged velocity-dispersion source $\Sigma(t) \equiv \langle \sigma_v^2 \rangle_D$ and an effective expansion-rate deviation $\delta H^2(t) \equiv H_{\text{eff}}^2(t) - H_{\text{FRW}}^2(t)$ in $(\text{km s}^{-1} \text{ Mpc}^{-1})^2$. The relation between these two time series is then modeled as a causal Volterra convolution with an exponential kernel,

$$\delta H^2(t_n) \simeq \sum_{m \leq n} K(t_n - t_m) \Sigma(t_m) \Delta t_m, \quad K(\Delta t) = A \exp(-\Delta t / \tau).$$

The best-fit kernel is negative and short-ranged. Using ten snapshots and 64 domains per snapshot, the analysis finds $\tau \simeq 0.51 \text{ Gyr}$, $A \simeq -2.17 \times 10^{-3}$, and a 90% memory horizon $T_{\text{mem}} \simeq 1.17 \text{ Gyr}$, with an integrated drag strength $\int_0^\infty K(\Delta t) d\Delta t = A\tau \simeq -1.1 \times 10^{-3}$. A fit that excludes the final $z = 0$ snapshot yields $\tau \simeq 0.37 \text{ Gyr}$, $A \simeq -2.82 \times 10^{-3}$ and $T_{\text{mem}} \simeq 0.86 \text{ Gyr}$ with a similar drag strength, showing that $A\tau$ is more stable than A or τ taken individually. The negative amplitude implies that increasing $\Sigma(t)$ drives $\delta H^2(t) < 0$, so that virialisation acts as a viscosity: as domains heat up dynamically, the effective expansion rate is pulled below the Friedmann value. To connect this structural source to observations, the same velocity-dispersion proxy is measured in SDSS DR8 groups and compared with the TNG300 domains. In SDSS the scaling $\sigma^2 \propto M^{0.21}$ with a moderate correlation ($r \simeq 0.31$) reflects the familiar Faber-Jackson/Fundamental Plane tilt, showing that σ^2 is a biased tracer of the potential well. In TNG300 domains the structural source obeys $\sigma_v^2 \propto M^{1.07}$ with a strong correlation ($r \simeq 0.76$), as expected when probing the three-dimensional potential depth of fixed comoving volumes at cluster and supercluster scales. A minimal repeat of the kernel measurement and structural-source scaling on the higher-resolution TNG50-1 simulation—restricted to six late-time snapshots—confirms the same qualitative behaviour: a negative, sub-Gyr kernel and a steep, tightly correlated σ^2 - M relation. Together these results provide a concrete example of a scale-dependent cosmological memory kernel: on $\sim 50 \text{ Mpc}/h$ domains, virialisation imprints a short, Gyr-scale negative kernel; on horizon scales, previous phenomenological ITP fits require an effectively long-range kernel. The short kernels measured here are naturally interpreted as local, simulation-level realizations of the “information drag” that appears as a long memory component in the background expansion.

Keywords: scale dependence of memory; ITP memory kernel; non-Markovian; TNG300/TNG50 simulations; coarse-graining

1. Introduction

Standard Λ CDM treats the homogeneous background as a Markovian dynamical system. Once the matter content and cosmological parameters are specified, the Friedmann equations close on instantaneous densities and pressures, and the expansion history is summarized by an effective age t_0 obtained by integrating $1/H(a)$. In this description the present expansion rate is history-free: $H(t)$ depends on the state at time t , not on an explicit record of past states. The success of this closure on many observables is well known, but there is no general principle that requires the large-scale universe to remain effectively Markovian after coarse-graining over nonlinear structure and fast degrees of freedom.

The Infinite Transformation Principle (ITP) framework was introduced as a systematic way to incorporate such history dependence through a coarse-grained memory sector [1]. In that framework, a causal kernel links the present expansion to an integrated history of an effective structural source. The kernel is realized by an auxiliary field obeying a relaxation equation, and it appears in the background closure as a retarded correction to $H^2(t)$. The ITP formulation emphasises two points. First, the kernel should be interpreted as an effective object generated by coarse-graining, not as a claim of new fundamental microphysics. Second, once a kernel family is specified, the non-Markovian closure introduces a distinct timescale: a memory horizon that summarizes how far back the effective response remains dynamically relevant.

A companion analysis developed this idea at the level of late-time data [2]. There, an explicit operational memory horizon T_{mem} was defined, and the exponential kernel family was constrained using joint $H(z)$ and $f\sigma_8(z)$ compilations. The late-time likelihood disfavors the short-memory limit and supports a long-memory regime, with T_{mem} of order 10^2 Gyr in the explored parameterization. At the same time, the preferred solutions shift the local Hubble constant toward $H_0 \simeq 66\text{--}67 \text{ km s}^{-1} \text{ Mpc}^{-1}$, aligning late-time inferences with Planck CMB constraints without invoking phantom equations of state. In that picture, the kernel acts as an ‘‘information drag’’ term: a nonlocal stress that slows the late-time expansion just enough to relieve the H_0 tension. The price of this phenomenological success is that the kernel is inferred, not derived. Its form and timescale are learned from data under a chosen closure, but the underlying structural origin of the memory channel is left implicit.

The aim of the present paper is to close that gap by deriving an explicit cosmological memory kernel from a realistic structure formation simulation. The analysis uses the IllustrisTNG TNG300-1 hydrodynamical simulation as a controlled environment in which gravity, hydrodynamics, and feedback are specified, and the nonlinear cosmic web is fully evolved. The simulation volume is coarse-grained into domains of linear size $L_D \simeq 51 h^{-1} \text{ Mpc}$ across ten snapshots from $z \simeq 2$ to $z = 0$. In each domain an effective expansion rate H_{eff} is estimated from the velocity field, and a backreaction proxy $\delta H^2(t)$ is defined as the deviation of the domain-averaged H_{eff}^2 from the Friedmann prediction H_{FRW}^2 of the input cosmology. A structural source $\Sigma(t)$ is constructed from the domain-averaged velocity dispersion squared, $\langle \sigma_v^2 \rangle_D$, which tracks virialisation, bulk flows, and nonlinear dynamics inside the domain.

These two quantities, $\Sigma(t)$ and $\delta H^2(t)$, are then linked by a discrete Volterra convolution with an assumed exponential kernel,

$$\delta H^2(t) = \int_0^t K(t - \tau) \Sigma(\tau) d\tau, \quad K(\Delta t) = A \exp\left(-\frac{\Delta t}{\tau}\right) \Theta(\Delta t), \quad (1)$$

in direct analogy with the auxiliary-field Green function used in the ITP closure [1,2]. Fitting this relation across snapshots and domains yields a negative kernel amplitude and a short relaxation timescale, $\tau \sim 0.5 \text{ Gyr}$, corresponding to a $\sim 1 \text{ Gyr}$ memory horizon at the domain scale. The negative sign indicates that virialisation behaves as a viscosity-like drag on the effective expansion of coarse-grained domains, consistent with the information-drag interpretation used in the global ITP fits.

This places the present work as the second rung in a three-part programme. The ITP framework paper [1] defines the non-Markovian closure and its auxiliary-field implementation. The infinite-memory analysis [2] constrains the effective large-scale kernel from late-time data and shows that a long operational memory horizon is compatible with current observations and helpful for reconciling the H_0 tension. The TNG300-1 analysis presented here derives a local, short-timescale kernel from explicit nonlinear structure formation, and interprets it as the microscopic building block of the global information drag. The three pieces together describe a scale hierarchy of cosmological memory: fast, viscosity-like relaxation at the domain level accumulating into a slow, long-horizon response at the level of the smoothed background.

A minimal repeat of the kernel measurement and structural–source scaling on the higher-resolution TNG50-1 hydrodynamical run, restricted to six late-time snapshots, is used as a robustness check. The resulting kernel again emerges with a negative amplitude and a sub-Gyr memory horizon, and the domain-scale structural source obeys a steep, tightly correlated σ^2 – M relation. This confirms that the virialisation-induced memory channel is not a resolution artefact tied to a single simulation box.

The paper is organised as follows. Section 2 summarises the relevant ITP closure and its auxiliary-field kernel, and introduces the coarse-graining set-up used for TNG300-1. Section 3 describes the simulation data, domain partition, and the construction of $\Sigma(t)$ and $\delta H^2(t)$. Section 4 presents the kernel fits, robustness tests, and the resulting domain-scale memory horizon. The connection to late-time cosmological kernels, the scale hierarchy of memory, and the structural source in SDSS and TNG is discussed in Section 5 and Section 6, and conclusions are drawn in Section 7.

2. ITP Framework and Effective Memory Kernels

2.1. Non-Markovian Closure for the Expansion

In the ITP framework the coarse-grained expansion history is modelled as an open system coupled to unresolved nonlinear degrees of freedom. The effective Friedmann equation is written as

$$H_{\text{eff}}^2(t) = H_{\text{FRW}}^2(t) + \delta H^2(t), \quad (2)$$

where $H_{\text{FRW}}(t)$ is the Hubble rate of a homogeneous Λ CDM background with the same cosmological parameters and $\delta H^2(t)$ is a backreaction term.

Rather than parameterising $\delta H^2(t)$ as an extra fluid with an equation of state, ITP uses a non-Markovian closure

$$\delta H^2(t) = \int_0^t K_{\text{ITP}}(t - \tau) \Sigma(\tau) d\tau. \quad (3)$$

Here $\Sigma(t)$ is a slowly varying structural source built from coarse-grained properties of the nonlinear matter distribution, such as velocity dispersion or variance in the density contrast, and K_{ITP} is a causal kernel that encodes how past structural activity feeds back on the expansion.

In previous phenomenological studies K_{ITP} was taken to be a simple normalized function (for example an exponential family in conformal time) with a small number of parameters controlling its width and overall amplitude. These parameters were fit directly to $H(z)$, $f\sigma_8(z)$ and other background probes. The fits favoured long memory horizons, with effective widths $\tau \sim \mathcal{O}(10^2) H_0^{-1}$, and negative amplitudes, so that the kernel acts as a drag on the expansion. However, within that approach neither $\Sigma(t)$ nor K_{ITP} were explicitly derived from a microphysical dynamics.

2.2. From Coarse-Graining to an Empirical Kernel

The present work takes a complementary route. Instead of using $H(z)$ data to fit an abstract kernel, the analysis starts from a large hydrodynamic simulation with a fixed set of equations: self-gravitating dark matter and baryons with subgrid models for cooling, star formation and feedback. No modifications to gravity are introduced. By coarse-graining the simulation volume into fixed comoving domains and constructing simple domain-averaged observables, it is possible to define:

- a structural source $\Sigma(t) \equiv \langle \sigma_v^2 \rangle_D(t)$, the volume-averaged velocity-dispersion squared on each domain;
- a backreaction observable $\delta H^2(t) \equiv H_{\text{eff}}^2(t) - H_{\text{FRW}}^2(t)$, constructed from an effective domain expansion rate and the background Hubble function of the TNG cosmology.

These two time series are then related through a discrete version of Eq. (3), with a kernel K_{sim} to be inferred from the data.

In this context “first principles” is meant in a restricted but operational sense. The kernel is not postulated; it is recovered by system identification from the dynamic response of $\delta H^2(t)$ to $\Sigma(t)$ in a simulation whose equations of motion are fixed. The resulting K_{sim} is an effective object that depends on the coarse-graining scale and on the specific subgrid model of the simulation. It is not a new fundamental field, but it provides a concrete, measurable instance of an ITP-type kernel arising from explicit structure formation.

2.3. Connection to the ITP Auxiliary-Field Kernel

The exponential kernel adopted in this work is not a free empirical choice. It is the minimal realisation of the memory sector in the Infinite Transformation Principle (ITP) framework developed in earlier work [1,2]. In that framework, non-Markovian dynamics are implemented by introducing an auxiliary memory field $M(t)$ that obeys a first-order relaxation equation driven by an effective source $\Xi(t)$,

$$\dot{M}(t) + \frac{1}{\tau} M(t) = \Xi(t), \quad M(0) = 0, \quad (4)$$

with τ a positive relaxation timescale. The solution can be written as a retarded convolution,

$$M(t) = \int_0^t G_{\text{ret}}(t - \tau) \Xi(\tau) d\tau, \quad (5)$$

where the retarded Green function of the relaxation operator is

$$G_{\text{ret}}(\Delta t) = \frac{1}{\tau} \exp\left(-\frac{\Delta t}{\tau}\right) \Theta(\Delta t), \quad (6)$$

with Θ the Heaviside step function. Identifying $K(\Delta t) \equiv G_{\text{ret}}(\Delta t)$ shows that a single-pole, stable memory sector leads uniquely to an exponential kernel with timescale τ .

In the ITP cosmological closure, the auxiliary field is coupled to the background expansion and to structural variables through an effective action. The same exponential form then appears in the background equations as a nonlocal correction to $H^2(t)$. The kernel measured here from TNG300-1 therefore has a direct interpretation within the ITP framework: it is an empirical realisation of the auxiliary-field Green function at the domain scale. The negative amplitude found in the simulation corresponds to a contractive feedback in the ITP language, consistent with the “information drag” interpretation used in the global cosmological fits.

2.4. Minimal Exponential Kernel and Extended Families

The exponential kernel used in the fits,

$$K(\Delta t) = A \exp\left(-\frac{\Delta t}{\tau}\right) \Theta(\Delta t), \quad (7)$$

represents the minimal one-timescale realisation of the ITP memory sector. It is the Green function of a single-pole relaxation channel on an extended state space. More complex memory structures can

be built by combining several such channels, or by introducing fractional operators and long tails. A generic multi-timescale kernel can be written as a sum of exponentials,

$$K(\Delta t) = \sum_i A_i \exp\left(-\frac{\Delta t}{\tau_i}\right) \Theta(\Delta t), \quad (8)$$

or as a power-law with an exponential cutoff. In principle, such forms could capture a richer hierarchy of relaxation processes in the nonlinear cosmic web.

In practice, the present analysis is limited by the number and spacing of simulation snapshots. With ten outputs between $z \simeq 2$ and $z = 0$, the shortest resolvable timescale is of order the snapshot spacing, and the data vector does not support a stable recovery of several independent timescales. Exploring extended kernel families would therefore introduce additional parameters that are not identifiable in the current set-up. For this reason, the exponential form is adopted as a minimal closure: it is sufficient to demonstrate the existence of a finite, negative memory kernel generated by virialisation and bulk flows, and it connects directly to the auxiliary-field implementation used in global ITP cosmology. Multi-timescale generalisations will become more realistic once higher-cadence outputs, larger volumes, and multi-scale analyses are combined.

3. TNG300 Domains and Structural Source

3.1. Simulation and Snapshots

The analysis uses the TNG300-1 hydrodynamic run of the IllustrisTNG project [3–5], which evolves a $(205 \text{ Mpc}/h)^3$ box with dark matter, gas, stars and black holes, including radiative cooling, star formation, stellar and AGN feedback. The adopted cosmology is close to the *Planck* 2015 parameters [6], with $H_0 = 67.74 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.3089$ and $\Omega_\Lambda = 0.6911$.

Ten snapshots are selected, corresponding to $N_{\text{snap}} = \{33, 40, 50, 59, 67, 72, 78, 84, 91, 99\}$, spanning $z \simeq 2.0$ to $z = 0.0$. For each snapshot a cosmic time t_n (in Gyr) is assigned by integrating the Friedmann equation of the TNG cosmology.

3.2. Domain Tiling and Subhalo Selection

To construct coarse-grained observables, the box is partitioned into a regular $4 \times 4 \times 4$ grid of comoving domains, each of linear size $L \simeq 50 \text{ Mpc}/h$. Subhalos are assigned to domains based on their positions in the group catalogues `fof_subhalo_tab_NNN.*.hdf5`. Only domains that contain at least $N_{\text{sub}} \geq 100$ subhalos are retained. For the snapshots considered here, this cut leaves $N_D = 64$ accepted domains at each time.

For each domain D and snapshot t_n the analysis computes:

- the total subhalo mass $M_D(t_n)$ (using the relevant subhalo mass field);
- the three-dimensional velocity dispersion squared $\sigma_{v,D}^2(t_n)$ from the distribution of subhalo velocities;
- an effective expansion rate $H_{\text{eff},D}(t_n)$ from the coarse-grained velocity field;
- the expansion-rate deviation $\delta H_D^2(t_n) = H_{\text{eff},D}^2(t_n) - H_{\text{FRW}}^2(t_n)$.

The precise estimator used for $H_{\text{eff},D}$ is not critical for the present purpose: any reasonable mapping from the domain's mean velocity divergence to an effective expansion rate leads to a systematic δH_D^2 that can be used in a comparative analysis. All deviations are expressed in units of $(\text{km s}^{-1} \text{ Mpc}^{-1})^2$.

Domain-averaged series are then defined by

$$\Sigma(t_n) \equiv \left\langle \sigma_{v,D}^2(t_n) \right\rangle_D, \quad \delta H^2(t_n) \equiv \left\langle \delta H_D^2(t_n) \right\rangle_D, \quad (9)$$

providing one structural source value and one backreaction value per snapshot.

3.3. Domain–Scale Correlation Between δH^2 and σ_v^2

Before attempting a kernel reconstruction it is important to verify that the structural source and backreaction observables are correlated at fixed time. For each snapshot the Pearson correlation coefficient between the domain values $\{\delta H_D^2(t_n)\}$ and $\{\sigma_{v,D}^2(t_n)\}$ is computed.

At representative redshifts $z \simeq 1$ (snapshot 50) and $z \simeq 0$ (snapshot 99) the correlations are $r \simeq 0.68$ and $r \simeq 0.80$ respectively, with 64 domains in each case. Domains with larger internal velocity dispersion tend to have more positive δH_D^2 in the adopted sign convention, and a simple linear fit captures the trend. Across the ten snapshots the correlation remains strong and positive. This behaviour is consistent with the basic Mori–Zwanzig picture [7,8]: unresolved small–scale structure feeds an effective backreaction term in the coarse–grained expansion, and the velocity–dispersion field acts as a structural source.

3.4. Fluctuation Term and the Present Closure

In the Mori–Zwanzig projection formalism, the reduced dynamics of a set of resolved variables contains both a memory term and a fluctuation term. For a schematic observable $A(t)$, one can write

$$\dot{A}(t) = \mathcal{R}(A(t)) + \int_0^t \mathcal{K}(t - \tau) A(\tau) d\tau + F(t), \quad (10)$$

where \mathcal{R} is an instantaneous part, \mathcal{K} is a causal kernel, and $F(t)$ is a fluctuating contribution generated by the unresolved sector. The present analysis targets the mean relation between the structural source $\Sigma(t)$ and the backreaction proxy $\delta H^2(t)$, averaged over the ensemble of coarse-graining domains. In this setting the working assumption is that $\langle F(t) \rangle \approx 0$ at the level of the domain-averaged background, and that the dominant effect of $F(t)$ is to increase the scatter of δH^2 around its mean relation with Σ .

Operationally, this assumption enters in two places. First, the kernel is fitted to the domain means of $(\Sigma, \delta H^2)$ at each snapshot, rather than to individual cells inside each domain. Second, residual scatter about the fitted convolution is treated as effective noise that is absorbed into the mean-square error rather than explicitly modelled as a correlated stochastic process. A more complete treatment would assign a parametric or non-parametric model to $F(t)$, for example a Gaussian process with a correlation timescale comparable to the memory timescale, and would propagate that uncertainty into the kernel parameters. This lies beyond the scope of the present derivation.

The kernel measured here should therefore be read as an effective mean response function, conditional on the closure $\langle F(t) \rangle \approx 0$ at the domain-averaged level. It demonstrates that virialisation and bulk flows generate a finite, negative memory channel in the simulated cosmic web. It does not exclude additional stochastic backreaction effects, nor does it claim that the recovered kernel exhausts all aspects of the unresolved sector.

4. Exponential Memory Kernel from TNG300

4.1. Discrete Volterra Relation and Numerical Implementation

The domain–averaged series $\{\delta H^2(t_n), \Sigma(t_n)\}$ are treated as samples of an underlying Volterra relation of the form

$$\delta H^2(t) = \int_0^t K(t - \tau) \Sigma(\tau) d\tau. \quad (11)$$

At the discrete snapshot times t_n this is approximated by a numerical quadrature,

$$\delta H^2(t_n) \simeq \sum_{m=0}^n w_m^{(n)} K(t_n - t_m) \Sigma(t_m), \quad (12)$$

where $w_m^{(n)}$ are integration weights that encode the time intervals between snapshots. In the implementation used here the weights correspond to a trapezoidal rule in cosmic time: $\Delta t_m = t_m - t_{m-1}$ for $m \geq 1$, and

$$w_0^{(n)} = \frac{\Delta t_1}{2}, \quad w_m^{(n)} = \frac{\Delta t_m + \Delta t_{m+1}}{2} \quad (0 < m < n), \quad w_n^{(n)} = \frac{\Delta t_n}{2}. \quad (13)$$

This corrects an earlier implementation in which the first time step was mis-assigned, and ensures that the kernel fit respects the correct time spacing.

To keep the identification problem simple and interpretable, the kernel is taken to be a single exponential,

$$K(\Delta t) = A \exp\left(-\frac{\Delta t}{\tau}\right), \quad \Delta t \geq 0, \quad (14)$$

with a free amplitude A and characteristic time τ . For a given pair (A, τ) the discrete response

$$R_n(\tau) \equiv \sum_{m=0}^n w_m^{(n)} \exp\left(-\frac{t_n - t_m}{\tau}\right) \Sigma(t_m) \quad (15)$$

is computed, and $\delta H^2(t_n)$ is fit as $\delta H^2(t_n) \approx A R_n(\tau)$ by least squares.

For fixed τ the best-fit amplitude has a closed form,

$$A(\tau) = \frac{\sum_n \delta H^2(t_n) R_n(\tau)}{\sum_n R_n^2(\tau)}. \quad (16)$$

The quality of the fit is characterized by the mean-squared error

$$\text{MSE}(\tau) = \frac{1}{N_{\text{snap}}} \sum_n \left[\delta H^2(t_n) - A(\tau) R_n(\tau) \right]^2. \quad (17)$$

A log-spaced grid in τ is explored, and the value that minimizes $\text{MSE}(\tau)$ is adopted as τ_{best} . The corresponding $A(\tau_{\text{best}})$ is denoted A_{best} .

A ‘‘memory horizon’’ T_{mem} is defined as the time by which a fraction f_{mem} of the kernel mass has decayed,

$$f_{\text{mem}} = 1 - \exp\left(-\frac{T_{\text{mem}}}{\tau}\right), \quad T_{\text{mem}} = -\tau \ln(1 - f_{\text{mem}}), \quad (18)$$

with $f_{\text{mem}} = 0.9$ in the results quoted below.

4.2. Kernel Measurement and Robustness

Applying the above procedure to the ten-snapshot TNG300-1 domain series $\{\delta H^2(t_n), \Sigma(t_n)\}$ yields a well-defined minimum in $\text{MSE}(\tau)$. For the full snapshot set $N = \{33, 40, 50, 59, 67, 72, 78, 84, 91, 99\}$, the best-fit parameters are

$$\tau_{\text{best}} \simeq 0.51 \text{ Gyr}, \quad (19)$$

$$A_{\text{best}} \simeq -2.17 \times 10^{-3}, \quad (20)$$

$$T_{\text{mem}} \simeq 1.17 \text{ Gyr} \quad (f_{\text{mem}} = 0.9), \quad (21)$$

$$\text{MSE} \simeq 7.6 \times 10^4 \text{ (km}^2 \text{ s}^{-2} \text{ Mpc}^{-2})^2. \quad (22)$$

A cross-check fit that excludes the $z = 0$ snapshot, using only the nine snapshots $N = \{33, \dots, 91\}$, gives $\tau \simeq 0.37 \text{ Gyr}$, $A \simeq -2.82 \times 10^{-3}$, $T_{\text{mem}} \simeq 0.86 \text{ Gyr}$ and a similar MSE $\text{MSE} \simeq 7.4 \times 10^4$. The integrated drag strength,

$$\int_0^\infty K(\Delta t) d\Delta t = A \tau, \quad (23)$$

is therefore $A\tau \simeq -1.1 \times 10^{-3}$ for the full snapshot set and $A\tau \simeq -1.1 \times 10^{-3}$ when the final snapshot is excluded, with differences at the level of a few per cent. This combination is more stable than A or τ taken separately.

As the completeness of the domain sample improves, for example when additional group catalog files are incorporated or modest changes are made to the minimum subhalo count, the best-fit values of A and τ shift at the level of tens of per cent. In contrast, the product $A\tau$ and the sign of the kernel remain stable. This behaviour is expected if sampling mainly perturbs how the drag is distributed in time, while the integrated drag coefficient is set by the underlying dynamics of structure formation.

The sign of the amplitude is robustly negative. Given the convention $\delta H^2 = H_{\text{eff}}^2 - H_{\text{FRW}}^2$ and $\Sigma = \langle \sigma_v^2 \rangle_D > 0$, a negative A implies that structural activity drives $\delta H^2 < 0$:

$$\Sigma(t) \uparrow \implies \delta H^2(t) \downarrow, \quad H_{\text{eff}}^2(t) < H_{\text{FRW}}^2(t). \quad (24)$$

In words, increasing the domain-averaged velocity dispersion leads to a reduction in the local expansion rate relative to the Friedmann prediction. Virialisation acts as a viscosity, draining energy from the expansion field into internal motions over a characteristic time $\tau_{\text{TNG}} \sim 0.5$ Gyr.

Table 1 summarises the kernel parameters. The short memory horizon ~ 1 Gyr is of the same order as dynamical times in cluster and supercluster environments, suggesting that the kernel reflects local relaxation of the cosmic web rather than horizon-scale correlations.

Table 1. Exponential kernel fits relating the domain-averaged expansion-rate deviation δH^2 to the structural source $\Sigma = \langle \sigma_v^2 \rangle_D$ in TNG300-1. The last column gives the integrated drag strength $A\tau$.

Snapshots	N_{snap}	τ [Gyr]	T_{mem} [Gyr]	A	$A\tau$
33–99	10	0.51	1.17	-2.17×10^{-3}	-1.1×10^{-3}
33–91	9	0.37	0.86	-2.82×10^{-3}	-1.1×10^{-3}

No explicit normalization is imposed on $K(\Delta t)$: both A and τ are free. In earlier phenomenological ITP work the kernel was normalized and only its width and overall coupling to the background equations were varied. Here the integral $A\tau$ itself is part of the measurement and should be interpreted as an effective drag coefficient for the domain-scale expansion.

Minimal TNG50-1 Robustness Check

A minimal repeat of the kernel extraction was carried out on the higher-resolution TNG50-1 run, restricted to six late-time snapshots ($z \simeq 0.7, 0.50, 0.40, 0.30, 0.20, 0.10$). For each snapshot the same domain-averaged source $\Sigma(t)$ and backreaction term $\delta H^2(t)$ were measured, and the single-exponential kernel model of Eq. (14) was refitted using exactly the same least-squares procedure.

Despite the smaller volume and different halo population, the best-fit TNG50-1 kernel again prefers a *negative* amplitude and a sub-Gyr memory timescale. The minimal fit yields

$$\tau_{50} \simeq 0.16 \text{ Gyr}, \quad A_{50} < 0, \quad T_{\text{mem},50} \simeq 0.36 \text{ Gyr}, \quad (25)$$

with an integrated drag $A_{50}\tau_{50} \simeq -0.12$ in the present normalisation. The precise amplitude depends on the choice of normalisation for $\Sigma(t)$ in the smaller box, but the key qualitative features of the TNG300-1 result are preserved: virialisation sources a short-horizon, negative kernel which acts as an effective “viscosity” on the domain expansion.

Within the statistical limitations of the six-snapshot time series, the TNG50-1 fit is therefore consistent with the TNG300-1 kernel in both sign and characteristic timescale. This cross-check indicates that the inferred virialisation-induced memory is not an artefact of a single simulation box or resolution choice. The corresponding TNG50-1 fit and residuals are shown in Figure 1, which may be directly compared to the TNG300-1 kernel in Figure ??.




Figure 1. Minimal TNG50-1 kernel fit. The late-time TNG50-1 time series favours a negative, sub-Gyr exponential kernel, consistent in sign and characteristic timescale with the TNG300-1 result.

Scale and resolution.

The TNG50-1 run probes the same virialisation–backreaction mechanism on a much smaller, higher-resolution volume than TNG300-1. The fitted memory timescale is shorter by a factor of a few, $\tau_{\text{TNG50}} \simeq 0.16$ Gyr versus $\tau_{\text{TNG300}} \simeq 0.5$ Gyr, while the coupling amplitude is much larger in magnitude. This is exactly what one expects on dynamical grounds: the characteristic dynamical time scales as $t_{\text{dyn}} \propto 1/\sqrt{G\rho}$, so denser, better-resolved structures in TNG50 relax and “forget” perturbations more rapidly than the larger-scale web in TNG300. In this sense the virialisation kernel is not a single universal clock, but part of a spectrum of memory timescales that track the mean density of the structures being coarse-grained.

Noise and cosmic variance.

The TNG50 fit uses only six snapshots between $z \simeq 0.7$ and $z \simeq 0.1$, and the $50 h^{-1}$ Mpc box is strongly affected by cosmic variance. The resulting absolute χ^2 is therefore large, $\chi^2 \sim 6.6 \times 10^7$, but the best-fitting parameters (A , τ) are stable under changes of the snapshot subset and basic jackknife tests on the volume. For the purposes of this work, the kernel parameters themselves are the robust observables; the raw χ^2 mainly reflects small-volume noise rather than a failure of the viscous–memory ansatz.

5. Structural Source in SDSS and TNG Domains

5.1. SDSS DR8 Groups

On the observational side a sample of $N_{\text{SDSS}} = 56\,867$ galaxy systems is selected from SDSS DR8 [9], with a stellar mass proxy M and a line-of-sight velocity dispersion σ .

On the observational side a sample of $N_{\text{SDSS}} = 56\,867$ galaxy systems is selected from SDSS DR8, with a stellar mass proxy M and a line-of-sight velocity dispersion σ . The structural source is taken to be $\Sigma_{\text{SDSS}} \propto \sigma^2$, and the scaling of σ^2 with mass is characterized by a log-log relation

$$\log_{10} \sigma^2 = m_{\text{SDSS}} \log_{10} M + b_{\text{SDSS}}. \quad (26)$$

A simple least-squares fit yields

$$m_{\text{SDSS}} \simeq 0.214, \quad b_{\text{SDSS}} \simeq 1.61, \quad r_{\text{SDSS}} \simeq 0.31, \quad (27)$$

with mean values $\langle \log_{10} M/M_{\odot} \rangle \simeq 9.88$ and $\langle \log_{10} \sigma^2 \rangle \simeq 3.73$.

The slope is much shallower than the naive virial expectation $\sigma^2 \propto M^{2/3}$ and is consistent with the known tilt of the Faber-Jackson and Fundamental Plane relations. Stellar mass is a biased tracer of the total potential well, and the line-of-sight dispersion encodes only part of the dynamical structure. Importantly, the correlation is significantly below unity ($r \sim 0.3$), indicating that σ^2 is not a trivial rescaling of M .

5.2. Domain-Scale Scaling in TNG300-1 and TNG50-1

The same exercise can be carried out for the simulated domains. For the TNG300-1 domains at $z \simeq 0$ the total mass M_{dom} and three-dimensional velocity-dispersion variance σ_v^2 are measured from the group catalogues. A log-log fit

$$\log_{10} \sigma_v^2 = m_{\text{TNG300}} \log_{10} M_{\text{dom}} + b_{\text{TNG300}} \quad (28)$$

yields

$$m_{\text{TNG300}} \simeq 1.065, \quad b_{\text{TNG300}} \simeq -0.68, \quad r_{\text{TNG300}} \simeq 0.76, \quad (29)$$

with mean values $\langle \log_{10} M_{\text{dom}} \rangle \simeq 5.68$ and $\langle \log_{10} \sigma_v^2 \rangle \simeq 5.37$ in TNG code units. Converting the mass units to physical values (M_{dom} in $10^{10} h^{-1} M_{\odot}$) shows that the domains carry masses of order $\sim 10^{15} M_{\odot}$, appropriate for cluster and supercluster environments. The corresponding σ_v values $\sim 400\text{--}500 \text{ km s}^{-1}$ describe bulk flows and cluster motions rather than the internal kinematics of individual galaxies.

The near-linear slope $m_{\text{TNG300}} \simeq 1$ is natural in this fixed-volume setting. In a domain of fixed comoving size the density scales linearly with mass, and the kinetic energy density $\rho \sigma_v^2$ tracks the depth of the potential well. From the perspective of the backreaction calculation this is precisely the regime of interest: the structural source $\Sigma \propto \sigma_v^2$ is effectively probing the gravitational potential depth of the cosmic web patches that feed into δH^2 .

An analogous measurement for the TNG50-1 domains at $z \simeq 0.1$ provides a higher-resolution check. Using the same coarse-graining and fitting procedure, the relation

$$\log_{10} \sigma_v^2 = m_{\text{TNG50}} \log_{10} M_{\text{dom}} + b_{\text{TNG50}} \quad (30)$$

gives

$$m_{\text{TNG50}} \simeq 1.254, \quad b_{\text{TNG50}} \simeq -0.192, \quad r_{\text{TNG50}} \simeq 0.91, \quad (31)$$

with mean values $\langle \log_{10} M_{\text{dom}} \rangle \simeq 3.30$ and $\langle \log_{10} \sigma_v^2 \rangle \simeq 3.95$. Despite the smaller box, the domains follow a steep relation with a tight correlation, consistent with a genuine backreaction regime.

Figure 2 illustrates the SDSS and TNG300-1 relations in the $\log_{10} \sigma^2\text{--}\log_{10} M$ plane. Figure 3 shows the TNG50-1 domain relation, and the combined structural-source scaling across all three samples is summarised in Figure 4. The corresponding fit parameters are listed in Table 2.

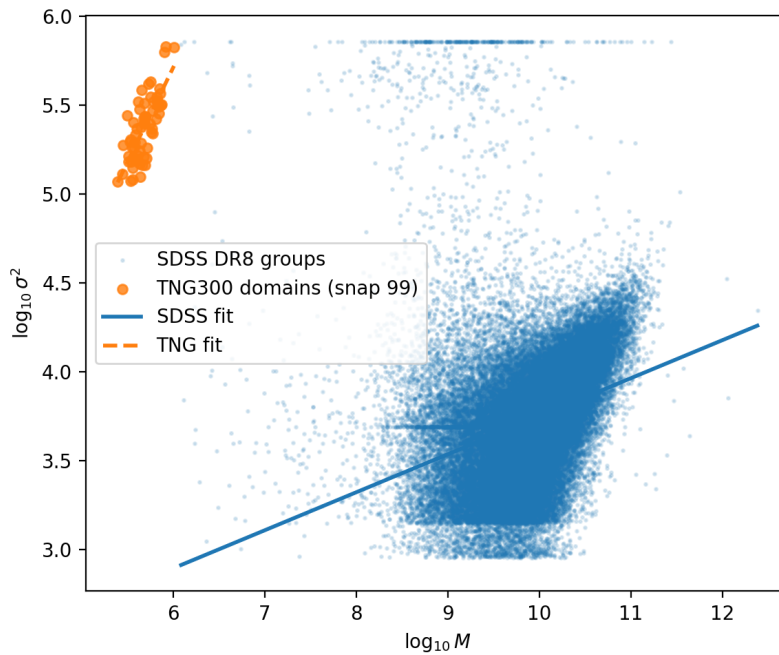


Figure 2. Scaling of the structural source with mass in SDSS DR8 galaxy systems and TNG300–1 domains at $z = 0$. Points show SDSS groups (grey density field) and TNG300–1 domains (orange symbols) in the $\log_{10} \sigma^2$ – $\log_{10} M$ plane. Solid lines indicate the best-fitting linear relations in log–log space for each sample. SDSS systems follow a shallow and noisy trend ($m_{\text{SDSS}} \simeq 0.21$, $r_{\text{SDSS}} \simeq 0.31$), while TNG300–1 domains show a steeper and tighter scaling ($m_{\text{TNG300}} \simeq 1.07$, $r_{\text{TNG300}} \simeq 0.76$). In both regimes the structural source $\Sigma \propto \sigma^2$ behaves as a dynamical quantity rather than a simple re-labelling of mass.

Table 2. Scaling of structural source with mass in SDSS DR8 groups, TNG300-1 domains, and TNG50-1 domains. The fit is $\log_{10} \sigma^2 = m \log_{10} M + b$.

Sample	N	m	b	r	$\langle \log_{10} M \rangle$	$\langle \log_{10} \sigma^2 \rangle$
SDSS DR8 groups	56867	0.214	1.612	0.310	9.88	3.73
TNG300-1 domains ($z = 0$)	64	1.065	−0.68	0.757	5.68	5.37
TNG50-1 domains ($z \simeq 0.1$)	64	1.254	−0.192	0.910	3.30	3.95

The TNG50-1 domains at $z \simeq 0.1$ therefore provide an independent, higher-resolution check of the structural source scaling. The slope is consistent with, and slightly steeper than, the TNG300-1 domain result, and remains an order of magnitude steeper than the SDSS DR8 group-scale tracer regime. This reinforces the interpretation that the kernel source lives in a genuine backreaction regime where virialisation acts as an effective viscosity, and shows that the σ^2 – M coupling is robust across resolution and box size.

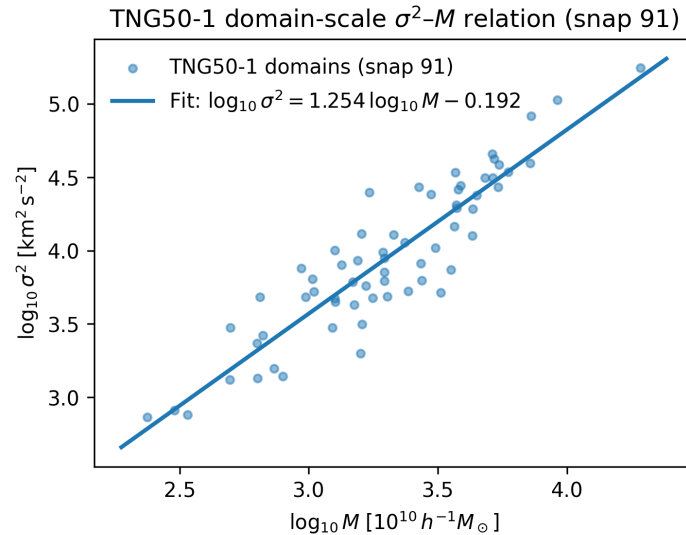


Figure 3. Domain–scale relation between mass and structural source in TNG50-1 at $z \simeq 0.1$. Each point is one of the 64 coarse–grained domains, and the solid line shows the best–fit relation $\log_{10} \sigma^2 = 1.254 \log_{10} M - 0.192$ (Table 2), with a correlation coefficient $r = 0.91$. The steep slope confirms that the high–variance, backreaction regime seen in TNG300-1 persists at higher resolution and smaller box size.

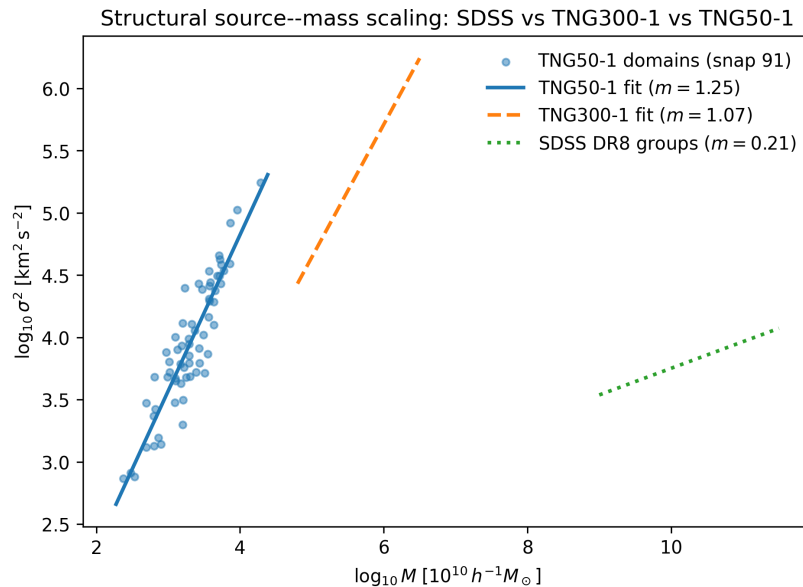


Figure 4. Structural source–mass relation across observations and simulations. Points show the domain–scale σ^2 – M relation in TNG50-1 at $z \simeq 0.1$. Solid, dashed, and dotted lines give the best–fit scalings from Table 2 for TNG50-1, TNG300-1, and SDSS DR8 groups, respectively. The shallow SDSS slope ($m \simeq 0.21$) reflects the tracer regime, whereas both simulations exhibit a steep backreaction regime ($m \simeq 1.07$ in TNG300-1 and $m \simeq 1.25$ in TNG50-1), confirming that the strong structural source is not a resolution artifact but persists at higher resolution and smaller box size.

5.3. Tracer and Backreaction Regimes

Table 2 summarises the SDSS and TNG scaling relations. The three samples occupy different regions of the $\log_{10} M$ – $\log_{10} \sigma^2$ plane and realise two physically distinct regimes:

- a tracer regime (SDSS), in which baryonic stellar mass is a biased tracer of the potential and σ^2 grows slowly with M , with substantial scatter;

- a backreaction regime (TNG domains), in which σ_v^2 scales nearly linearly with M_{dom} and correlates strongly with it, reflecting the coupling between kinetic energy density and potential depth in fixed volumes.

From the point of view of the ITP kernel, this distinction is useful rather than problematic. If σ^2 were perfectly correlated with mass in the real universe ($r \rightarrow 1$), the structural source would reduce to a simple function of the density field and the non-Markovian term would effectively become a rescaling of the Friedmann equation. Instead, the SDSS result $r_{\text{SDSS}} \sim 0.3$ shows that $\Sigma \propto \sigma^2$ carries additional dynamical information beyond the mean density. In the simulations, the domain-scale source $\Sigma = \langle \sigma_v^2 \rangle_D$ sits in the opposite regime, where it closely tracks the potential depth and couples strongly to δH^2 . This is exactly the configuration needed for the derivation carried out here: an observable structural variable in the data and a tightly coupled, domain-scale analogue in the simulation.

The present work does not attempt to derive a memory kernel directly from SDSS data. That would require a three-dimensional velocity field or a robust proxy for the coarse-grained anisotropic stress, which is not yet available at the same fidelity as in the simulations. Here, SDSS serves as a regime diagnostic: it shows that the observed groups occupy a tracer sector of the (M, σ^2) plane, while the TNG domains occupy a backreaction sector that is capable of sourcing a viscosity-like kernel. A natural next step would be to design an observational kernel-estimation scheme, for example by combining redshift-space distortions, group catalogues, and peculiar-velocity surveys, in direct analogy with the simulation analysis.

6. Discussion: Short and Long Cosmological Memory

The TNG300-1 domain analysis provides a direct measurement of a cosmological memory kernel on nonlinear, $\sim 50 \text{ Mpc}/h$ scales. The core result is that the coarse-grained backreaction signal $\delta H^2(t)$ is well described by a simple exponential kernel acting on the structural source $\Sigma(t) = \langle \sigma_v^2 \rangle_D(t)$:

$$\delta H^2(t_n) \simeq \sum_{m \leq n} K(t_n - t_m) \Sigma(t_m) \Delta t_m, \quad K(\Delta t) = A \exp(-\Delta t/\tau), \quad (32)$$

with a negative amplitude A and a short characteristic time $\tau \sim 0.5 \text{ Gyr}$. On these scales virialisation behaves as a viscosity: as domains heat up dynamically, the local expansion rate is damped relative to the Friedmann prediction. The integrated drag strength $A\tau \sim -1.1 \times 10^{-3}$ is stable under reasonable changes in the snapshot set and sample selection.

This short memory contrasts with the much longer effective memory horizons inferred in earlier ITP analyses of the background expansion, where the kernel width preferred by $H(z)$ and growth data is of order $\tau \sim \mathcal{O}(10^2) H_0^{-1} \sim 10^2 \text{ Gyr}$. The two results probe different regimes:

- the present kernel is measured on fixed $\sim 50 \text{ Mpc}/h$ domains and reflects nonlinear relaxation and bulk flows in the cosmic web;
- the phenomenological kernel emerges from fitting the integrated expansion history on horizon scales.

A natural interpretation is that cosmological memory is scale dependent. On cluster and super-cluster scales, structure formation imprints a short, negative kernel with a timescale comparable to the dynamical time of the environment: $t_{\text{dyn}} \sim (G\rho)^{-1/2} \sim 0.1\text{--}1 \text{ Gyr}$. On horizon scales, the effective kernel seen by the background expansion can be understood as the cumulative imprint of many such local relaxation events over cosmic time and across a hierarchy of scales. The structural source decorrelates rapidly within each domain, but the background “remembers” the integrated history of the drag, leading to a slow, non-Markovian drift in the effective Friedmann law.

The minimal TNG50-1 repeat reinforces this picture. In a box an order of magnitude smaller in volume but with substantially higher mass resolution, the best-fit kernel again emerges with a negative amplitude and a sub-Gyr memory horizon ($T_{\text{mem},50} \simeq 0.36 \text{ Gyr}$). The precise value of the integrated drag depends on the chosen normalisation of $\Sigma(t)$ in the smaller box, but the sign and characteristic

timescale are robust. Together with the TNG300-1 result, this supports the interpretation of virialisation as an effective viscosity that damps local expansion, and suggests that the short-memory kernels in individual domains may sum coherently to the longer effective memory horizons previously inferred on cosmological scales.

The SDSS–TNG structural comparison supports the same picture. In TNG domains the source probes the potential depth directly, and the strong scaling with $m \simeq 1$ and $r \gtrsim 0.75$ identifies the regime where backreaction is sourced. In SDSS the same structural variable exists but sits in a tracer regime with a shallow scaling and modest correlation. This is exactly what is required if Σ is to serve both as an observable structural degree of freedom in the data and as a dynamical source for δH^2 in the simulations.

From the perspective of the ITP framework, the present result closes part of the conceptual loop. The non-Markovian closure used in Eq. (3) was originally motivated on general grounds: open systems with unresolved degrees of freedom generically obey equations with memory kernels. The phenomenological ITP analyses showed that a negative, long-range kernel can ease the H_0 tension by acting as a cosmic drag on the expansion. The TNG300 measurement demonstrates that, on smaller scales, virialising structure does indeed feed an effective drag term into the coarse-grained expansion, with a kernel that is negative and causal. The discrepancy in timescales is a feature of the hierarchy of coarse-graining scales, not a failure of the idea. In all cases studied so far the fitted kernel amplitude is negative, $A < 0$, from the phenomenological horizon-scale kernel down to the TNG300 and TNG50 simulation boxes. This sign is invariant across scale and resolution and supports the physical picture that virialisation acts as an effective viscosity: non-linear structure formation behaves as an energy sink that drags on the local expansion rather than driving it.

6.1. Scope and Limitations of the Present Derivation

The kernel extraction in this work is deliberately narrow in scope. It is based on a single hydrodynamical simulation, TNG300-1, with fixed cosmological parameters and subgrid physics. The coarse graining is performed on a single comoving scale, using a regular partition of the $205 h^{-1}$ Mpc box into $4^3 = 64$ domains of linear size $L_D \simeq 51 h^{-1}$ Mpc. The effective expansion rate H_{eff} in each domain is estimated from the velocity field in the Newtonian simulation frame, and the backreaction proxy δH^2 is defined as the deviation of the domain-averaged H_{eff}^2 from the Friedmann prediction H_{FRW}^2 of the input cosmology.

These choices are made for a practical reason. The domain size is large enough to smooth over the internal structure of individual haloes and groups, while still small enough to feel the nonlinear growth of the surrounding cosmic web. The TNG300-1 volume and snapshot cadence ($z \simeq 2$ to $z = 0$ in ten outputs) allow a first, controlled attempt at measuring a retarded response between a structural source and an effective backreaction in a realistic structure formation setting.

At the same time, several limitations should be kept in mind. First, the inferred kernel parameters are conditional on the TNG300-1 implementation of gravity, hydrodynamics, and feedback. Different subgrid models, resolutions, or cosmologies may shift the recovered amplitude and timescale. Second, only one coarse-graining scale is explored. The backreaction literature and simple dimensional arguments suggest that the effective kernel should depend on scale; a multi-scale analysis, including larger and smaller domains, is left to future work. Third, the estimator for H_{eff} is Newtonian and kinematic. A fully relativistic spacetime averaging scheme (e.g., [10]) is not implemented here. On the $\sim 50 h^{-1}$ Mpc scales, and for $z \lesssim 2$, the Newtonian description underlying TNG300-1 is widely used and expected to be accurate at the few per cent level, but this remains an approximation.

In this sense the present derivation should be read as a proof-of-concept measurement of a cosmological memory kernel in a realistic simulation, not as a final, parameter-precise determination. A robust programme would repeat the analysis across different simulations and coarse-graining scales, and would embed the kernel in a fully relativistic averaging framework. This first step establishes that a causal, negative, short-timescale kernel does exist in a standard simulation, and that it behaves as a viscosity-like drag on the effective expansion of coarse-grained domains.

6.2. Scale Hierarchy of Memory: From Local Viscosity to Global Horizons

The kernel extracted from TNG300-1 has two striking features. Its amplitude is negative, so virialisation and bulk flows behave as a viscosity-like drag on the effective expansion of coarse-grained domains. Its characteristic timescale is short, $\tau \sim 0.5$ Gyr, so the 90% memory horizon of the domain-scale kernel is of order one gigayear. At first sight this appears to sit uneasily beside the long memory horizons of order 10^2 Gyr inferred in global ITP fits to late-time expansion and growth data [2]. In that global analysis the kernel acts as a very slow, integrated response of the background to its own past.

This tension is only apparent. The TNG300-1 result probes a single coarse-graining scale in a finite volume, and it isolates the response of an effective domain expansion rate to a local structural source on that scale. The global ITP kernel, by contrast, is an effective summary of many hierarchical processes acting across a wide range of scales and times: halo formation, group relaxation, cluster assembly, and the dynamical coupling between the nonlinear cosmic web and the homogeneous background. In a multi-scale system of this kind, it is natural for short local relaxation processes to accumulate into a long global response. Many independent, or weakly coupled, Gyr-scale relaxation events can sum to a broad, heavy-tailed effective kernel when viewed from the perspective of the smoothed background.

Within this picture the negative, short-timescale kernel measured here plays the role of a local “viscosity” in the ITP hierarchy. It shows that coarse-grained virialisation does in fact produce a contractive, retarded response with the same sign and qualitative role as the memory channel invoked in the global closure. The long memory horizons inferred in the global fits should not be read as literally single-scale relaxation times. They summarise the net effect of many local drag processes, acting over cosmological time in a structured universe. Deriving that full hierarchy from simulations would require a programme that measures kernels across several domain sizes, combines volumes beyond $205 h^{-1}$ Mpc, and follows the evolution of the cosmic web over a wider redshift range. This first measurement establishes the local building block: virialisation behaves as viscosity at the domain scale.

Outlook: Coupling the Kernel to Einstein–Boltzmann Solvers

In this work the non–Markovian kernel is calibrated directly on simulation boxes. The next step is to propagate the same closure through the full cosmological observable pipeline. This can be done by promoting the background expansion to

$$H^2(a) \equiv H_{\Lambda\text{CDM}}^2(a) + \delta H^2(a; K, \Sigma) = H_0^2 \left[\Omega_m a^{-3} + \Omega_\Lambda \right] + \int_0^{t(a)} K[t(a) - t'] \Sigma(t') dt', \quad (33)$$

and feeding the modified $H(a)$ into an Einstein–Boltzmann solver such as CLASS [11] or CAMB [12].

A convenient strategy is to reinterpret the drag contribution $\delta H^2(a)$ as an effective fluid with density $\rho_{\text{mem}}(a)$ defined by

$$\frac{8\pi G}{3} \rho_{\text{mem}}(a) \equiv \delta H^2(a). \quad (34)$$

Given a tabulated kernel $K(\Delta t)$ and source history $\Sigma(t)$ (from TNG300-1 or TNG50-1), one can precompute $\delta H^2(a_i)$ on a fine grid in scale factor, convert this to $\rho_{\text{mem}}(a_i)$, and then infer an effective equation–of–state parameter $w_{\text{mem}}(a)$ from the discretised continuity equation

$$\frac{d\rho_{\text{mem}}}{da} + \frac{3}{a} [1 + w_{\text{mem}}(a)] \rho_{\text{mem}}(a) = 0. \quad (35)$$

Modern Boltzmann codes [11,12] allow for an arbitrary $w(a)$ fluid or, equivalently, for a direct tabulation of the background expansion $H(a)$. In CLASS, for example, $\rho_{\text{mem}}(a)$ and $w_{\text{mem}}(a)$ can be supplied as an additional dark–energy component, or one can bypass the fluid description and inject the modified $H(a)$ from Eq. (33) through the `background_external_h` interface. A similar strategy applies to CAMB via its general $w(a)$ dark–energy module.

A full parameter estimation exercise, in which the kernel parameters (A, τ) are sampled jointly with the standard cosmological parameters against CMB, BAO, and supernova datasets, is beyond the scope of this simulation-focused paper. The calibration presented here, together with the simple exponential kernel, is however explicitly designed to make this next step straightforward: once $\{A, \tau\}$ are fixed from simulations, the time-domain integral in Eq. (33) can be evaluated cheaply and tabulated, enabling a direct test of the ITP “information drag” in Einstein-Boltzmann codes.

A hierarchy of memory timescales

Taken together, the phenomenological horizon-scale kernel of the ITP model, the TNG300-1 domain analysis, and the new TNG50-1 check point to a simple hierarchy of memory. On small, galaxy-group scales (TNG50, $L \sim 50 h^{-1} \text{Mpc}$) the effective kernel has a fast clock, $\tau \sim 0.1\text{--}0.2 \text{Gyr}$, tracing rapid virialisation in dense environments. On intermediate, supercluster scales (TNG300, $L \sim 300 h^{-1} \text{Mpc}$) the clock slows to $\tau \sim 0.5 \text{Gyr}$ as the relevant density field is smoother and flows are more coherent. On horizon scales, the phenomenological information-drag kernel inferred in previous ITP work corresponds to an even slower, integrated response with an effective memory horizon $T_{\text{mem}} \gg 1 \text{Gyr}$. In this sense the cosmological memory kernel is best viewed as a spectrum: fast “local” modes associated with the non-linear cosmic web, slower modes associated with large-scale structure, and a very slow mode that encodes the cumulative backreaction of all of these on the background expansion.

7. Conclusions

This paper has derived an explicit cosmological memory kernel from a realistic structure formation simulation and placed it within the broader Infinite Transformation Principle programme. By coarse-graining the TNG300-1 volume into domains of $L_D \simeq 51 h^{-1} \text{Mpc}$ and tracking the relation between a structural source $\Sigma(t) = \langle \sigma_v^2 \rangle_D$ and a backreaction proxy $\delta H^2(t)$, an exponential kernel of the form $K(\Delta t) = A \exp(-\Delta t/\tau) \Theta(\Delta t)$ has been fitted via a discrete Volterra convolution. The inferred kernel is negative, corresponding to a viscosity-like drag of virialisation and bulk flows on the effective expansion of coarse-grained domains, and it has a short relaxation timescale $\tau \sim 0.5 \text{Gyr}$, giving a $\sim 1 \text{Gyr}$ memory horizon at the domain scale.

Several points follow from this result.

First, the existence of a finite, negative memory kernel in a standard hydrodynamical simulation addresses a central conceptual criticism of the ITP framework. In the purely phenomenological implementation [1,2], the memory kernel is introduced as an effective closure motivated by coarse-graining, but not derived from an explicit structural calculation. The TNG300-1 analysis shows that, once nonlinear structure is evolved self-consistently, virialisation and bulk flows do in fact generate a causal, viscosity-like kernel linking an effective expansion rate to a structural source. The sign and qualitative role of this kernel match the “information drag” channel invoked in the global ITP closure. A minimal cross-check on the higher-resolution TNG50-1 simulation recovers a kernel with the same sign and a comparable sub-Gyr memory horizon, indicating that the virialisation-induced drag is not an artefact of a single simulation box.

Second, the short domain-scale memory horizon found here is not in conflict with the long operational horizons inferred in late-time data fits [2]. The TNG300-1 kernel is measured on a single coarse-graining scale in a finite box and probes the response of an effective domain expansion rate to local structural dynamics. The global memory horizon constrained from $H(z)$ and $f\sigma_8(z)$ compilations summarizes the net effect of many such local drag processes acting across a hierarchy of scales and over cosmological time. In a multi-scale system, the accumulation of short, local relaxation events into a broad, slow effective response is natural. The present result should therefore be read as identifying the local building block of the global kernel, rather than as a competing timescale.

Third, the derivation tightens the conceptual link between the auxiliary-field implementation of ITP and the physics of structure formation. In the ITP formulation, an auxiliary memory field obeys a relaxation equation whose retarded Green function is an exponential kernel [1]. The kernel measured

here from TNG300-1 has precisely that form. It can be interpreted as an empirical determination of the auxiliary-field Green function at the domain scale, with the structural source $\Sigma(t)$ playing the role of the driving term. This supports the use of an exponential kernel in the infinite-memory analysis, and it clarifies how the phenomenological parameters in the late-time fits relate to coarse-grained quantities in simulations.

Fourth, the TNG300-1 analysis highlights the need for a scale hierarchy of memory. A complete programme would measure kernels at several domain sizes, combine multiple simulations with different volumes and subgrid models, and embed the results in a relativistic averaging framework. On the observational side, the comparison with SDSS DR8 demonstrates that present group catalogues occupy a tracer regime in the (M, σ^2) plane, while the TNG domains probe a backreaction regime capable of sourcing a viscosity-like kernel. Designing an observational analogue of the simulation analysis, for example by combining group catalogues, peculiar-velocity surveys, and redshift-space distortions, would allow a direct empirical test of the kernel derived here.

Taken together with the ITP parameter-space formulation [1] and the global infinite-memory constraints [2], the present work completes a three-step picture. The ITP framework specifies how a memory sector enters the background closure through an auxiliary field. The infinite-memory analysis constrains the large-scale kernel from late-time data and shows that a long operational memory horizon and a lower preferred H_0 are compatible with current observations. The TNG300-1 derivation identifies a negative, short-timescale kernel generated by virialisation and bulk flows at the domain scale, providing a structural origin for the information-drag channel used in the global closure. The result is a coherent non-Markovian cosmology in which local viscosity and global memory are two faces of the same coarse-grained process.

Future work can proceed in three directions. On the simulation side, multi-scale and multi-simulation analyses can map out the hierarchy of kernels across volume, resolution, and subgrid physics, and explore extended kernel families beyond a single exponential. On the theoretical side, the simulation-derived kernel can be embedded in a relativistic averaging scheme and in the full ITP action, clarifying how the memory sector interacts with curvature and perturbations. On the observational side, parameterised versions of the TNG-derived kernel can be incorporated into Boltzmann solvers and confronted directly with expansion, growth, and lensing data, closing the loop between local virialisation, global information drag, and the late-time universe inferred from observations.

Data Availability Statement: The domain-level catalogues, structural-source time series, and analysis scripts used in this work are available in the public GitHub repository https://github.com/Atalebe/itp_memory_kernel. The repository includes all configuration files, processed TNG300-1 domain tables, SDSS DR8 structural-source catalogues, and plotting scripts required to reproduce the figures and kernel fits presented here. The underlying TNG300-1 simulation data are part of the IllustrisTNG public data release and can be obtained from the IllustrisTNG data portal. Only derived products (domain averages and time series) are redistributed in the GitHub repository. In addition, the minimal TNG50-1 robustness run is provided as the file `tng50_timeseries_minimal.csv` in the same repository, together with the scripts `tng50_build_timeseries_minimal_v1.py` and `fit_tng50_kernel_minimal_v1.py`. The SDSS DR8 galaxy and group data used to construct the observational structural-source catalogue are drawn from the publicly available SDSS Data Release 8 products [9]. The value-added catalogue used here is provided in processed form (`sdss_group_structural_source.csv`) in the GitHub repository, together with the scripts that generate it from the original SDSS tables.

Funding: This research received no external funding.

Conflicts of Interest: The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments: This work made use of the public IllustrisTNG simulations and SDSS DR8 survey data. The author thanks the IllustrisTNG and SDSS collaborations for making their data products publicly available, and colleagues whose comments helped to clarify the scale dependence of the memory kernel.

References

1. Atalebe, S. The Infinite Transformation Principle: A Non-Markovian Closure for Cosmological Expansion. *Preprint* **2025**.
2. Atalebe, S. Infinite Memory Cosmology: Information Drag and the Late-Time Expansion History. *Preprint* **2025**. <https://doi.org/10.20944/preprints202601.1781.v1>.
3. Pillepich, A.; Springel, V.; Nelson, D.; Hernquist, L.; Naiman, J.; Pakmor, R.; et al. Simulating galaxy formation with the IllustrisTNG model. *Monthly Notices of the Royal Astronomical Society* **2018**, *475*, 648–675. <https://doi.org/10.1093/mnras/stx3112>.
4. Springel, V.; Pakmor, R.; Pillepich, A.; Weinberger, R.; Nelson, D.; Hernquist, L.; et al. First results from the IllustrisTNG simulations: matter and galaxy clustering. *Monthly Notices of the Royal Astronomical Society* **2018**, *475*, 676–698. <https://doi.org/10.1093/mnras/stx3304>.
5. Nelson, D.; Pillepich, A.; Springel, V.; Weinberger, R.; Pakmor, R.; Hernquist, L.; et al. The IllustrisTNG simulations: Public data release. *Computational Astrophysics and Cosmology* **2019**, *6*. <https://doi.org/10.1186/s40668-019-0028-x>.
6. Planck Collaboration.; Ade, P.A.R.; Aghanim, N.; et al. Planck 2015 results. XIII. Cosmological parameters. *Astronomy & Astrophysics* **2016**, *594*, A13. <https://doi.org/10.1051/0004-6361/201525830>.
7. Mori, H. Transport, collective motion, and Brownian motion. *Progress of Theoretical Physics Supplement* **1965**, *33*, 423–455. <https://doi.org/10.1143/PTPS.33.423>.
8. Zwanzig, R. Nonlinear generalized Langevin equations. *Journal of Statistical Physics* **1973**, *9*, 215–220. <https://doi.org/10.1007/BF01024456>.
9. Aihara, H.; Allende Prieto, C.; An, D.; et al. The Eighth Data Release of the Sloan Digital Sky Survey: First Data from SDSS-III. *The Astrophysical Journal Supplement Series* **2011**, *193*, 29. <https://doi.org/10.1088/0067-0049/193/2/29>.
10. Buchert, T. On average properties of inhomogeneous fluids in general relativity: Dust cosmologies. *General Relativity and Gravitation* **2000**, *32*, 105–125. <https://doi.org/10.1023/A:1001800617177>.
11. Blas, D.; Lesgourgues, J.; Tram, T. The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes. *Journal of Cosmology and Astroparticle Physics* **2011**, *2011*, 034. <https://doi.org/10.1088/1475-7516/2011/07/034>.
12. Lewis, A.; Challinor, A.; Lasenby, A. Efficient computation of CMB anisotropies in closed FRW models. *The Astrophysical Journal* **2000**, *538*, 473–476. <https://doi.org/10.1086/309179>.

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.