Boundary-Condition Quantum Mechanics IV_d: Bundle formation, graph–Darwinism, and inertial rigidity from shared edges

Peter M. Ferguson
Independent Researcher

31st December 2025

Abstract

BCQM IV_d is the bundle/clustering counterpart to the single–thread analysis of BCQM IV_c[1] within the Stage 1 inertial–noise programme of Boundary–Condition Quantum Mechanics (BCQM) as developed in BCQM IV[2]. BCQM IV_b[3] established a W_{coh}–blind control model with essentially constant inertial–noise amplitude $A_a(W_{\text{coh}})$ and fitted exponent $\beta \approx 0$ in $A_a \sim W_{\text{coh}}^{-\beta}$. BCQM IV_c then showed that single primitive threads equipped with a binary "soft–rudder" kernel and a natural W_{coh}–sensitive slip law live in a diffusive universality class, with $A_a(W_{\text{coh}}) \sim W_{\text{coh}}^{-1/2}$ and, more generally, $\beta \approx \alpha/2$ for slip laws $q(W_{\text{coh}}) \propto W_{\text{coh}}^{-\alpha}$.

BCQM IV_d takes those single–thread results as given and turns to bundles and clusters of threads. We define explicit toy bundle models on a pre–spacetime event graph, grounded in the BCQM primitives, and quantify how the centre–of–mass (COM) inertial noise of a bundle scales with the coherence horizon W_{coh} and with bundle size N. The bundles extend the IV_c soft–rudder dynamics by adding a shared–bias coupling of strength λ between threads, modelling the tendency of threads to reuse shared edges in a coherent patch.

Our numerical B–series runs show that independent bundles ($\lambda = 0$) exhibit the expected $N^{-1/2}$ suppression of COM noise and an effective exponent $\beta_{\text{COM}} \approx 0.4$ almost independent of N, i.e. the same diffusive universality class as IV_c up to finite–W_{coh} and observable effects. As the shared–bias coupling is increased through $\lambda = 0.25$, 0.5, and 0.75, bundles become more coherent and their COM inertial noise is further suppressed relative to the independent baseline. For small bundles at $\lambda = 0.75$, $\beta_{\text{COM}}(N)$ is pushed upwards towards the IV_c diffusive ceiling $\beta \approx 1/2$ while remaining below it and gently decreasing with N.

We conclude that modest, local bundle glue based on shared edges can stiffen bundle COM motion and suppress inertial noise beyond simple $N^{-1/2}$ averaging, but does not yield ballistic $\beta \approx 2$ behaviour without contrived single–thread slip laws. Single primitive threads remain in a robust diffusive universality class, and any effectively Newtonian "rigid" inertial response in BCQM must arise from stronger, more hierarchical bundle mechanisms to be explored in later papers.

1 Introduction and Stage 1 context

BCQM IV_d is conceived as the bundle paper of Stage 1 (kinematics / inertial noise) in the staged BCQM programme. Stage 1 focuses on kinematics and inertial noise on discrete event chains; later stages address emergent geometry (Stage 2) and gravity and gauge structure (Stage 3 and beyond). It starts from two established baselines:

- **BCQM IV_b (control).** A W_{coh}–blind, direction–free toy kernel on event chains produces an inertial–noise amplitude $A_a(W_{\text{coh}})$ that is essentially flat across the scanned W_{coh} values; the fitted exponent β in $A_a \sim W_{\text{coh}}^{-\beta}$ is consistent with $\beta \approx 0$ within errors.

IV_b's role is to validate the full numerical pipeline (event chains $\to$ acceleration PSD $\to$ amplitude $\to$ fitted β) while making no strong claims about physical inertia.

– **BCQM IV_c (single thread).** A one–dimensional binary "soft–rudder" kernel with internal state $v_n = \pm 1$, fixed step $x_{n+1} = x_n + v_n$, and a $W_{\rm coh}$–dependent slip probability $q(W_{\rm coh}) = 1 - p_{\rm stay}(W_{\rm coh})$ provides a $W_{\rm coh}$–sensitive test–bed. For the natural slip law $q(W_{\rm coh}) = k/W_{\rm coh}$ with $k = 2$, the measured amplitude obeys $A_a(W_{\rm coh}) \sim W_{\rm coh}^{-\beta}$ with $\beta \approx 0.5$. More generally, for a one–parameter family $q(W_{\rm coh}) \propto W_{\rm coh}^{-\alpha}$, numerics and random–telegraph analysis indicate $\beta \approx \alpha/2$.

In particular, the natural $\alpha = 1$ case gives a diffusive single–thread exponent $\beta \approx 1/2$. Achieving $\beta = 2$ with a single binary fixed–step thread would require $\alpha \approx 4$, i.e. a highly contrived suppression of slips. The conservative lesson of IV_c is therefore that *single primitive threads live in a diffusive universality class*, and that if BCQM is to support an effectively Newtonian, "rigid" inertial response, it must arise from bundles or clusters of threads rather than from any single worldline.

BCQM IV_d develops this bundle picture in a deliberately minimalist way. It will:

– treat bundles as explicit many–thread constructions on a noisy event graph, not as axioms;

– define all notions (evaporation, bundle membership, rigidity, COM motion) in transparent graph or toy–model terms;

– measure how the COM inertial noise scales with $W_{\rm coh}$ and bundle size N for simple bundle kernels;

– report the actual exponents and flip statistics, even if they fall short of the most optimistic "ballistic" scenario.

As background for these bundle constructions, the next section sketches a minimal picture of the pre–spacetime foam we have in mind at Stage 1, after which we turn to threads, events, and bundles and their COM motion.

2 Pre–spacetime foam and a Rado-like background

At the primitive level, BCQM imagines a pre–spacetime "foam" of events and directed edges with no preferred locations or directions. A convenient mathematical toy model for such a maximally symmetric background is the classical Rado graph: a countable random graph obtained by connecting each pair of nodes independently with probability $p \in (0, 1)$. In the present paper we use this only as an *idealised reference* for what a featureless background might look like; we do not assume that the physical event graph is literally Rado.

For the purposes of IV_d, the only properties we need are that a Rado–like background:

– has no preferred nodes, directions, or length scales, so that finite neighbourhoods see statistically similar surroundings;

– contains copies of every finite motif somewhere in the web, so coherent patches and "dusty" regions are both allowed in principle;

– is compatible with the idea that decoherence, finite coherence horizon $W_{\rm coh}$, and hop–bounded propagators K_r act as a form of *graph Darwinism*, favouring some subgraphs while suppressing others.

In this picture, long–lived bundles of threads correspond to non–generic, symmetry–broken regions where certain edges have been repeatedly reinforced by the dynamics and function as coherent channels within an otherwise noisy foam.

3 Threads, events, and bundles

Given this Rado-like foam picture for the ambient event graph, we now recall BCQM's primitives and explain how bundles and their centre-of-mass motion are represented. BCQM's primitives can be summarised as:

- *Events* E — realised "ticks".
- *Directed edges* $E_i \to E_j$ with complex amplitudes $a(E_i \to E_j)$, with $|a| \leq 1$.
- *Threads* — one realised chain of events for a particle: $E_0^{(a)} \to E_1^{(a)} \to E_2^{(a)} \to \dots$.

Graphically, events are the nodes of the directed graph; there is no separate notion of "node" distinct from "event". A single primitive particle is represented by one realised thread. When many particles are present, the primitive picture is a family of threads, all embedded in the same ambient event graph.

In this language, there are several ways that two or more threads can have "shared history":

- they may literally share events (their realised paths intersect);
- they may traverse distinct events but follow many of the same directed edges;
- they may be confined, for many hops, to the same small subgraph $\mathcal{C}$ of the ambient graph, even if they do not cross exactly the same edges at the same times.

The last two cases motivate a working notion of *shared edges* as bundle glue: a bundle is not defined by a one–shot coincidence, but by many hops in which the participating threads reuse the same subset of strong edges in a coherent local patch.

3.1 Shared edges as bundle glue

Let $\mathcal{G}$ denote the ambient event graph and $\mathcal{C} \subset \mathcal{G}$ a small connected subgraph. Informally, we say that:

- $\mathcal{C}$ is a *coherent patch* if its internal edges carry relatively large amplitudes $|a(E_i \to E_j)|$ and are frequently traversed by realised threads;
- a set of threads forms a *bundle* over some time window if their realised events lie, for many hops, primarily inside $\mathcal{C}$ and repeatedly reuse that same subset of strong edges.

Because the hop–bounded propagator K_r favours paths with large composite amplitude $|A[\gamma]|$, a coherent patch $\mathcal{C}$ acts as a channel: remaining inside $\mathcal{C}$ reuses strong edges and so yields large $|K_r|$, while leaving $\mathcal{C}$ typically involves weaker or misaligned edges and hence smaller $|K_r|$.

BCQM IV_d will not impose a hard "no–flip rule" at the bundle level. Individual threads remain free to flip, wander, and evaporate. Instead, the bundle kernels will encode the *relative cost* of leaving shared edges and will measure the consequences for COM motion and bundle longevity.

In this sense, the shared–edge mechanism realises, at the bundle level, the phase–survival principle developed in BCQM III[4]: loops and repeated use of strong edges provide a form of mutual error correction via interference and can, in favourable patches, extend phase coherence over W_{coh} relative to isolated threads that sample the ambient graph more freely.

4 Toy bundle kernels for IV_d

For concreteness, IV_d will work with toy kernels that extend the single–thread soft–rudder dynamics of IV_c to bundles of N threads.

4.1 Single–thread soft–rudder recap

In one dimension, a single thread is described in discrete time by its position x_n and a velocity–like internal state $v_n \in \{\pm 1\}$, with fixed step

$$x_{n+1} = x_n + v_n. \tag{1}$$

At each hop the soft–rudder kernel either keeps or flips the direction,

$$v_{n+1} = \begin{cases} v_n, & \text{with probability } p_{\text{stay}}(W_{\text{coh}}), \\ -v_n, & \text{with probability } q(W_{\text{coh}}) = 1 - p_{\text{stay}}(W_{\text{coh}}), \end{cases} \tag{2}$$

with a slip law $q(W_{\text{coh}})$ that decays with the coherence horizon. The case $q(W_{\text{coh}}) = k/W_{\text{coh}}$ with $k = 2$ was found in IV_c to give $\beta \approx 1/2$ for the inertial–noise amplitude.

4.2 Many–thread bundles and COM motion

A bundle of N such threads is specified by the collection $\{x_n^{(i)}, v_n^{(i)}\}_{i=1}^{N}$ at each step n. The centre–of–mass (COM) position is

$$X_n = \frac{1}{N}\sum_{i=1}^{N} x_n^{(i)}, \tag{3}$$

with COM velocity and discrete acceleration defined by finite differences,

$$V_n = X_{n+1} - X_n, \qquad a_n = V_{n+1} - V_n. \tag{4}$$

IV_d will consider bundle kernels in which:

- each thread still obeys a soft–rudder update rule, but
- the probabilities p_{stay} and q are weakly coupled across the bundle, for example through a shared bias towards the current COM direction, or through correlations induced by shared edges in a local patch of the event graph.

The simplest cases interpolate between:

- *independent threads*, where each $v_n^{(i)}$ follows an independent soft–rudder process (yielding COM noise suppressed as $\sim N^{-1/2}$ but still diffusive in W_{coh});
- *strongly synchronised bundles*, where the $v_n^{(i)}$ tend to align and flip together when shared edges are disrupted, potentially producing a more rigid COM trajectory with rarer large accelerations.

IV_d will use these toy kernels to explore how the COM inertial–noise amplitude $A_{\text{COM}}(W_{\text{coh}}, N)$ scales with W_{coh} and N, and to what extent any effective exponent $\beta_{\text{COM}}(N)$ in fits of the form

$$A_{\text{COM}}(W_{\text{coh}}, N) \sim W_{\text{coh}}^{-\beta_{\text{COM}}(N)} \tag{5}$$

can exceed the single–thread value $\beta \approx 1/2$ without fine tuning.

5 Numerical observables and scans

To make the bundle story sharp and testable, IV_d will specify a numerical pipeline closely parallel to IV_b and IV_c:

- generate ensembles of bundle trajectories for grids of W_{coh} and N;
- compute the COM acceleration time series a_n for each run;
- estimate the acceleration power spectral density for COM motion and extract the amplitude $A_{\mathrm{COM}}(W_{\mathrm{coh}}, N)$ at a chosen characteristic frequency window;
- fit amplitude scaling laws of the form $A_{\mathrm{COM}}(W_{\mathrm{coh}}, N) \sim W_{\mathrm{coh}}^{-\beta_{\mathrm{COM}}(N)}$.

In addition, the simulations will track several diagnostic statistics:

- the distribution of single–thread flips within a bundle;
- the frequency of multi–thread or whole–bundle flips;
- bundle "evaporation" events in which a significant fraction of threads leave the coherent patch $\mathcal{C}$ and cease to share edges;
- characteristic evaporation times as functions of W_{coh} and N.

These observables are designed to separate three effects:

(a) suppression of COM noise simply by averaging independent diffusive threads;

(b) additional suppression arising from shared–edge synchronisation;

(c) failure modes where bundles evaporate too quickly to act as meaningful inertial carriers.

5.1 Bundle B-series runs: A2, B1, B2, B3

To go beyond the design stage, we implemented a concrete "B-series" of bundle runs using the soft–rudder kernel and a shared–bias bundle coupling. All runs share a common numerical pipeline, closely parallel to BCQM IV_b and IV_c: ensembles of bundle trajectories are generated on a grid of coherence horizons W_{coh} and bundle sizes N, the COM acceleration time series are extracted, and the inertial–noise amplitude $A_{\mathrm{COM}}(W_{\mathrm{coh}}, N)$ is estimated from the COM acceleration power spectral density in a fixed frequency band. Log–log fits of $A_{\mathrm{COM}}(W_{\mathrm{coh}}, N)$ vs W_{coh} then define effective exponents $\beta_{\mathrm{COM}}(N)$ via

$$A_{\mathrm{COM}}(W_{\mathrm{coh}}, N) \sim W_{\mathrm{coh}}^{-\beta_{\mathrm{COM}}(N)}. \tag{6}$$

In addition, we record alignment and flip statistics — the mean instantaneous alignment S_v, an effective flip coherence parameter κ_{eff}, and a bundle lifetime based on an alignment threshold — as diagnostic observables.

The B-series keeps the single–thread dynamics fixed to the IV_c "natural" soft–rudder law $q(W_{\mathrm{coh}}) = 2/W_{\mathrm{coh}}$ and varies only the bundle coupling strength λ:

- **A2** (independent baseline): $\lambda = 0$, no shared bias between threads;
- **B1** (weak coupling): $\lambda = 0.25$;
- **B2** (intermediate coupling): $\lambda = 0.5$;
- **B3** (stronger coupling): $\lambda = 0.75$.

All four runs use the same numerical grid and analysis choices:

- $W_{coh} \in \{20, 50, 100\}$;
- bundle sizes $N \in \{1, 2, 4, 8, 16, 32\}$;
- identical ensemble size and trajectory length for each (W_{coh}, N) pair;
- a common PSD window and frequency band for extracting A_{COM}.

This ensures that differences between A2, B1, B2, and B3 can be attributed to bundle correlations rather than numerical details.

Before turning to bundles, we also reimplemented the IV_c single–thread soft–rudder model within the `bcqm_bundles` pipeline as a one–thread[5] "A1" regression run with $N = 1$ only. Using the same COM acceleration–PSD observable and the same W_{coh} grid as in the B-series, the fitted single–thread exponent comes out slightly below the ideal telegraph value $\beta \approx 1/2$ (we find $\beta_{single} \approx 0.4$), which we trace to the change of observable and to the limited W_{coh} range rather than to a different universality class. We therefore continue to treat $\beta \approx 1/2$, as established analytically and numerically in BCQM IV_c[1], as the appropriate single–thread diffusive ceiling when interpreting the bundle exponents.

We treat the IV_c single–thread value $\beta \approx 1/2$ as a *diffusive ceiling*: the B-series asks how close bundle COM exponents $\beta_{COM}(N)$ can approach this ceiling, and under what conditions, without changing the single–thread law or introducing contrived glue. The independent baseline A2 confirms that bundles *without* correlations do not change the W_{coh}–scaling exponent: A_{COM} follows the expected $N^{-1/2}$ suppression, and $\beta_{COM}(N) \approx 0.4$ for all N, i.e. the same diffusive universality class as IV_c up to finite–W_{coh} and observable effects. Any enhancement $\beta_{COM} > \beta_{single}$ must therefore come from genuine inter–thread correlations, not from bundle size alone.

As λ is increased from 0 through 0.25 and 0.5 to 0.75, the B-series exhibits a clear and monotone "bundle stiffening" pattern:

- for $\lambda = 0$ (A2), the ratio $A_{COM}(W, N)/[A_{single}(W)/\sqrt{N}]$ stays close to 1 for all N;
- for weak coupling (B1), this ratio falls slightly below 1 for $N \geq 2$, especially at larger W_{coh} and N;
- for intermediate coupling (B2), the extra suppression becomes clearly visible and $\beta_{COM}(N)$ is pushed above the independent value for small N;
- for $\lambda = 0.75$ (B3), the suppression is strongest: small bundles show ratios as low as ~ 0.6 at $W_{coh} = 100$, and $\beta_{COM}(N)$ comes close to the diffusive ceiling for $N = 2, 4$, while remaining below $\beta \approx 1/2$ and gently decreasing with N.

At the same time, the diagnostics show that multi–thread bundles become more coherent as λ increases: κ_{eff} shifts towards positive values and the mean alignment $\langle S_v \rangle$ increases for $N > 1$, especially at small N. Bundle lifetimes grow but remain finite, and multi–thread bundles still evaporate under the chosen alignment threshold. Even at $\lambda = 0.75$ we do not impose any $1/N^2$–style "extreme glue" or engineered $N(W_{coh})$ scaling: the shared–bias rule is modest and local, so the observed stiffening reflects genuine bundle dynamics rather than a hard–wired rigidity constraint.

The main IV_d figures built from these runs are Figures 1 to 3, which together show the suppression of bundle centre–of–mass inertial noise relative to independent averaging and the associated alignment and flip coherence diagnostics.

6 Universality classes and expectations

The working universality–class picture going into IV_d can be summarised as:

- single primitive threads with natural soft–rudder slip laws sit in a robust *diffusive* class, with $\beta \approx 1/2$;
- bundles of many such threads can suppress COM noise by a factor $\sim N^{-1/2}$ purely by averaging;
- additional shared–edge structure may further suppress COM noise and reduce the frequency of large COM flips, producing an effectively more rigid inertial response.

IV_d will quantify how far this suppression can go for simple bundle kernels, and will provide a first set of diagnostics for how rigidity scales with bundle size N.

Alongside A_{COM} and β_{COM}, it is useful to have a dimensionless measure of combinatorial rigidity. One simple choice is

$$\kappa_{\mathrm{eff}} = \log \frac{P(0\,\text{flips})}{P(1\,\text{flip})}, \tag{7}$$

where $P(k)$ denotes the measured probability that exactly k threads in the bundle flip in a given update step. Large positive κ_{eff} indicates that no-flip steps are much more common than single-thread flips, consistent with a rigid bundle; values near zero correspond to a much softer bundle in which individual threads frequently change direction.

As a working hypothesis, motivated by the growth of potential pairwise constraints with bundle size, we will test whether combinatorially rigid bundles exhibit a scaling of the form $\kappa_{\mathrm{eff}} \sim \log(N^2)$. This is not assumed a priori; rather, IV_d will treat any such trend in $\kappa_{\mathrm{eff}}(N)$ as an empirical output of the simulations, and will also examine the full distribution $P(k)$ for $k > 1$ to characterise large multi-thread flips.

and will *not* assume or enforce a ballistic $\beta = 2$ a priori. In particular, it will:

- treat any apparent $\beta_{\mathrm{COM}} \gtrsim 1/2$ as a measurable output of the bundle dynamics, not as a target imposed on the model;
- explicitly document parameter regimes where bundles fail to form, evaporate rapidly, or show only trivial $N^{-1/2}$ averaging;
- avoid baking in rigidity by forbidding multi–thread flips or by imposing artificial constraints that have no clear graph–level interpretation.

7 Outlook towards BCQM V

Within the staged BCQM programme, IV_d is still a Stage 1 paper: it works entirely at the level of discrete event chains and toy kernels. Its role is to justify, or constrain, the idea that massive particles and macroscopic bodies correspond to bundles or knots of primitive threads whose COM inertial noise is strongly suppressed relative to that of a single primitive thread.

If the bundle kernels studied here show robust COM suppression and rare large accelerations without fine tuning, they will provide a concrete, testable basis for treating bundles as effective inertial carriers in the emergent spacetime picture of BCQM V. If they do not, IV_d will still play an important role by narrowing the space of viable bundle mechanisms and by recording null results that future refinements must overcome.

In keeping with its Stage 1 remit, IV_d deliberately restricts attention to modest, local glue rules: the shared–bias bundles studied here introduce only weak correlations between threads and avoid any hard–wired $1/N^2$–style constraints or engineered $N(W_{\text{coh}})$ scalings. In separate lab–note work we have explored much stronger, hierarchical glue — for example nested bundles with near–lockstep flips — which can drive the COM persistence length far beyond W_{coh} and produce an effectively rigid response. Those extreme kernels belong to the emergent–spacetime and mass–generation story of BCQM V rather than to the Stage 1 inertial–noise baseline recorded here.

Either way, BCQM IV_d will leave untouched the Stage 2 and Stage 3 questions of how a metric, curvature, and gravity emerge from the event graph. Those topics belong to later papers. The present work stays narrowly focused on what can and cannot be achieved with bundles and shared edges at the primitive event level.

Figures

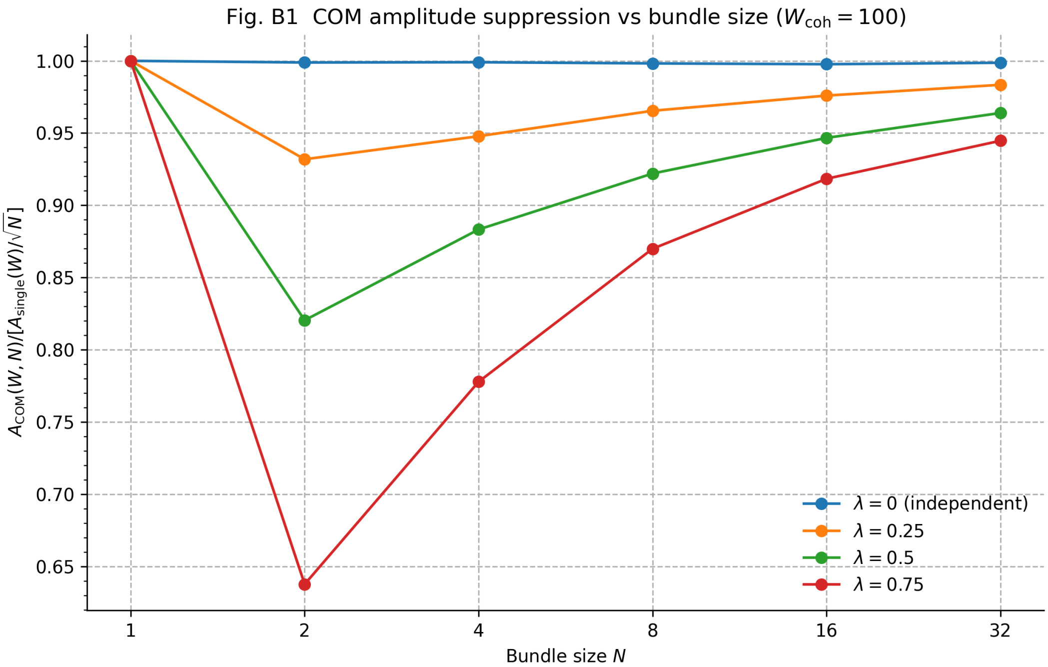

Figure 1: Suppression of bundle centre–of–mass inertial–noise amplitude relative to independent averaging. The curves show the ratio $A_{\text{COM}}(W_{\text{coh}}, N, \lambda)/[A_{\text{single}}(W_{\text{coh}}; 0)/\sqrt{N}]$ as a function of bundle size N at $W_{\text{coh}} = 100$ for the independent baseline ($\lambda = 0$) and shared–bias bundles with $\lambda = 0.25, 0.5, 0.75$. Independent bundles stay near unity, while increasing λ produces a smooth, monotone suppression for $N > 1$.

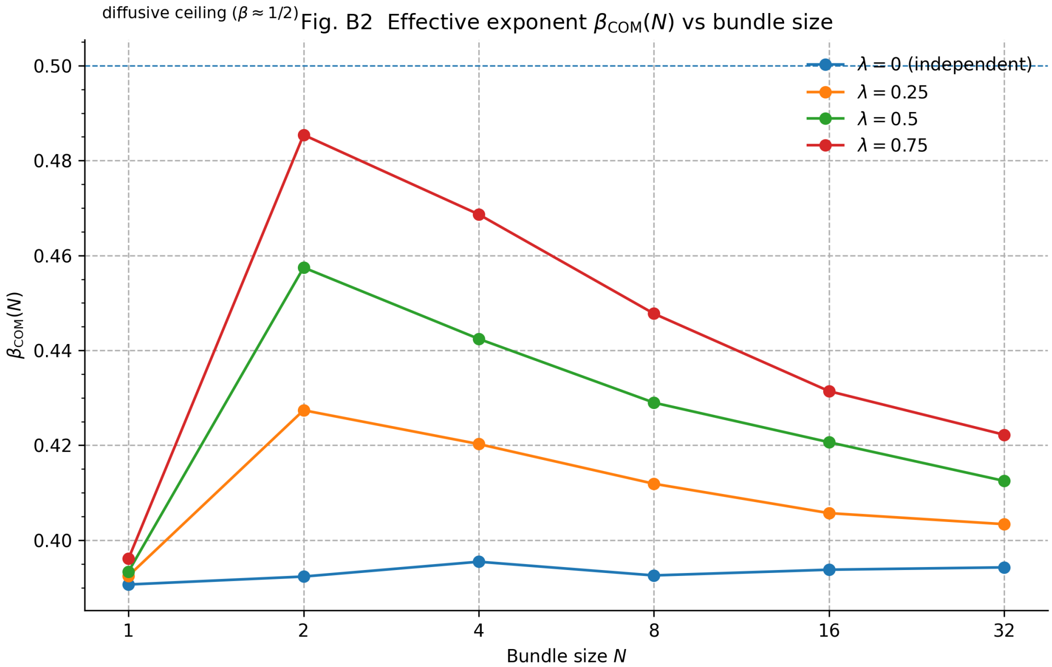

Figure 2: Effective bundle centre–of–mass exponent $\beta_{\rm COM}(N,\lambda)$ extracted from $A_{\rm COM}(W_{\rm coh},N,\lambda) \sim W_{\rm coh}^{-\beta_{\rm COM}(N,\lambda)}$ across $W_{\rm coh} \in \{20, 50, 100\}$. For the independent baseline $(\lambda = 0)$ one finds $\beta_{\rm COM} \approx 0.4$ almost independent of N, matching the IV__c diffusive universality class up to finite–$W_{\rm coh}$ and observable effects. As the shared–bias coupling is increased to $\lambda = 0.5$ and $\lambda = 0.75$, the exponents for small bundles are driven upwards towards the IV__c diffusive ceiling $\beta \approx 1/2$ while remaining below this value and decreasing gently with bundle size.

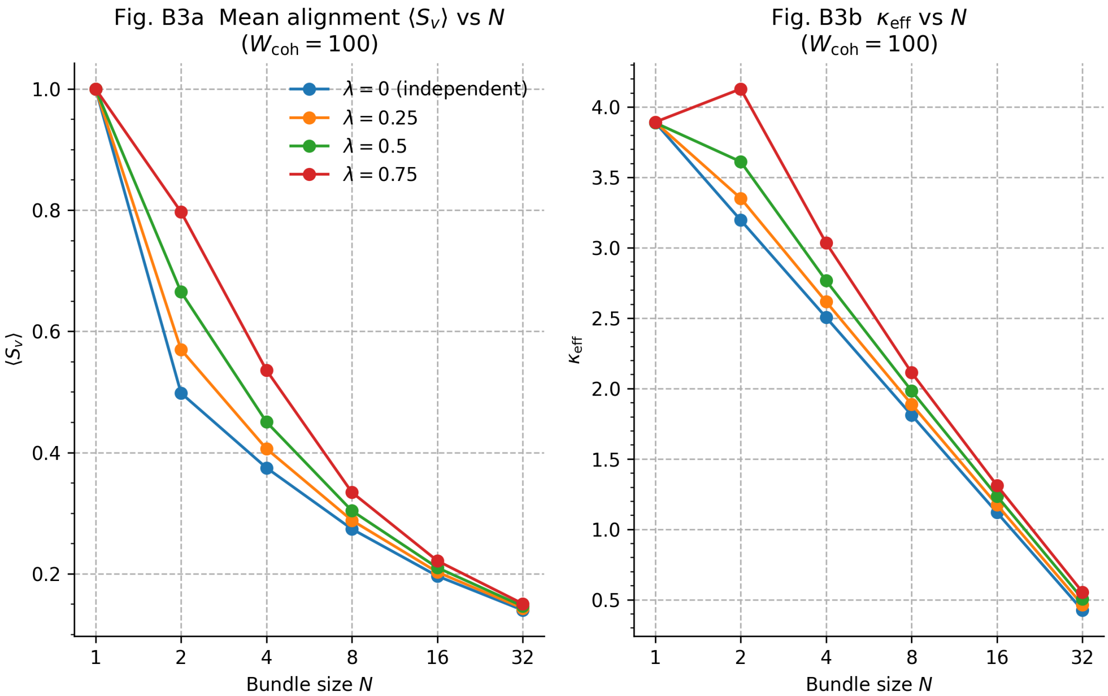

Figure 3: Bundle coherence diagnostics for the B-series at $W_{\text{coh}} = 100$. Left: mean instantaneous alignment $\langle S_v \rangle$ as a function of bundle size N for the independent baseline and shared–bias bundles with $\lambda = 0.25, 0.5, 0.75$. Right: effective flip coherence parameter $\kappa_{\text{eff}}(N, \lambda)$ for the same runs. Both diagnostics move in the "more coherent" direction as λ increases, confirming that the spectral stiffening seen in Figures 1 and 2 is tied to genuine inter–thread coherence rather than to a hidden $1/N^2$–style glue.

References

[1] Peter M. Ferguson. *Boundary-Condition Quantum Mechanics IV_c: Diffusive inertia and the limits of binary hop kernels*. 2025. DOI: 10.5281/zenodo.17921862. URL: https://doi.org/10.5281/zenodo.17921862.

[2] Peter M. Ferguson. *Boundary-Condition Quantum Mechanics IV: Inertial Noise and the Emergent Action*. 2025. DOI: 10.5281/zenodo.17650149. URL: https://doi.org/10.5281/zenodo.17650149.

[3] Peter M. Ferguson. *Boundary-Condition Quantum Mechanics IV_b: Baseline inertial-noise spectra form a W_coh-blind event-chain control model*. 2025. DOI: 10.5281/zenodo.17815888. URL: https://doi.org/10.5281/zenodo.17815888.

[4] Peter M. Ferguson. *Boundary-Condition Quantum Mechanics III: A Stochastic Growth Model for Causal Event Chains and the Emergence of Inertia*. 2025. DOI: 10.5281/zenodo.17632453. URL: https://doi.org/10.5281/zenodo.17632453.

[5] Peter M. Ferguson. *BCQM IV_d companion code v1.0.0*. GitHub+Zenodo snapshot of the BCQM_IV_d simulation and analysis code. 2025. DOI: 10.5281/zenodo.18108766. URL: https://doi.org/10.5281/zenodo.18108766.