Boundary-Condition Quantum Mechanics IV_c: Diffusive inertia and the limits of binary hop kernels

Peter M. Ferguson
Independent Researcher

13th December 2025

Abstract

Boundary-Condition Quantum Mechanics (BCQM) introduces a finite coherence horizon W_{coh} at the level of quantum events and uses it to define an intrinsic inertial-noise spectrum. Papers IV [1] and IV_b [2] developed this formalism and a numerical pipeline based on Ornstein–Uhlenbeck–type control kernels which are deliberately blind to W_{coh}, leading to an amplitude scaling $A(W_{\mathrm{coh}}) \approx \mathrm{const}$ and fitted exponent $\beta \approx 0$. In this note we take the next step and study a family of simple, explicitly W_{coh}-dependent graph kernels in which a "soft rudder" rule makes the persistence of a binary hop direction increase with W_{coh}. Using the same inertial-noise pipeline as IV_b, we show that this entire class of binary, fixed-step models produces a robust amplitude scaling $A(W_{\mathrm{coh}}) \propto W_{\mathrm{coh}}^{-\beta}$ with $\beta \approx 1/2$, independent of the detailed slip law. We interpret this as evidence for a diffusive inertia universality class, limited by the shot noise of ± 1 steps, and argue that purely binary hop geometries cannot yield a smooth, high-inertia limit with $\beta \approx 2$. This motivates the richer graph primitives and clock-based curvature models developed in BCQM V.

Within the broader BCQM continuum-limit programme, this paper is a Stage 1 contribution: it concerns only the kinematic emergence of effective trajectories and their inertial-noise spectra from discrete hop rules. Questions of emergent metric structure (Stage 2), curvature from mass or energy clusters (Stage 3), and internal symmetry or gauge structure (Stage 3b) are deliberately left to future work, to be addressed in BCQM V and beyond.

1 Introduction

BCQM I–III [3][4][5] introduced the coherence horizon W_{coh} and a discrete event-based picture in which inertial motion and spacetime emerge from the statistics of collapse events. BCQM IV [1] formalised the notion of an intrinsic inertial noise spectrum associated with a probe thread and linked it to W_{coh} at the level of general principles. BCQM IV_b [2] then constructed an explicit numerical pipeline using an Ornstein–Uhlenbeck (OU) control kernel that is deliberately *blind* to W_{coh}, and showed that this control model yields an amplitude $A(W_{\mathrm{coh}})$ that is essentially independent of W_{coh} (fitted $\beta \approx 0$), together with the expected $N^{-1/2}$ suppression of centre-of-mass noise for clusters and a mapping into SI units.

The natural next question is what happens when W_{coh} enters a graph-level hop kernel in a controlled way. Can a purely local, W_{coh}-dependent rule reduce inertial noise and thereby realise the idea of emergent inertia from event statistics? If so, how far can such a rule be pushed before the discrete geometry itself becomes the limiting factor?

The purpose of this paper (BCQM IV_c) is to address these questions for a broad and analytically transparent class of one-dimensional "soft rudder" models in which W_{coh} only affects the persistence of a binary hop direction $v_n \in \{\pm 1\}$, while the step size remains fixed. Using the same spectral pipeline as IV_b, we show that these models produce a robust scaling $A(W_{\mathrm{coh}}) \propto W_{\mathrm{coh}}^{-\beta}$ with $\beta \approx 1/2$, essentially independent of the precise slip law used. We interpret this as evidence for a *diffusive inertia* universality class arising from the shot noise of binary steps, and we argue that a purely binary hop geometry cannot produce a smooth, high-inertia limit with $\beta \approx 2$. This motivates the richer graph primitives and clock-based curvature models to be explored in BCQM V.

Structure of the paper. Section 2 defines the binary soft-rudder hop kernel and its relation to the W-blind OU control model of BCQM IV_b. Section 3 summarises the numerical pipeline used to extract the inertial-noise amplitude scaling and recalls its validation in BCQM IV_b. Section 4 presents the amplitude scaling results and interprets them in terms of a diffusive inertia universality class and a no-go statement for binary fixed-step kernels. Section 5 summarises the lessons for the broader BCQM continuum-limit programme and for the planned BCQM V. Appendix A collects numerical details, parameter tables, and robustness checks.

2 Binary soft-rudder hop kernel

In this section we define the one-dimensional binary soft-rudder model that will serve as our test-bed for W_{coh}-dependent kernels. The aim is to keep the construction as simple and transparent as possible, while still capturing the essential feature of a direction persistence that increases with W_{coh}.

2.1 State variables and flat control model

We consider a discrete-time process with step index $n = 0, 1, 2, \ldots$, in which a probe has position $x_n \in \mathbb{Z}$ and a direction variable $v_n \in \{+1, -1\}$. Each hop updates the state via

$$x_{n+1} = x_n + v_{n+1}, \tag{1}$$

so that the step size is fixed to $|\Delta x| = 1$ and all non-trivial structure enters through the dynamics of v_n.

For comparison we recall the W-blind control model used in BCQM IV_b, in which an effective OU-type acceleration kernel is used to generate an inertial-noise process $a(t)$ that is independent of W_{coh}. In the present binary setting one can define an analogous W-blind control process in which v_n flips with a fixed probability p_0 at each step, independent of W_{coh}. This yields a simple symmetric random walk and serves as the baseline against which we compare the W-dependent soft-rudder kernels below.

2.2 Soft-rudder rule and W_{coh}-dependence

We now introduce a family of W-dependent "soft rudder" rules in which the direction variable v_n is updated according to

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

The key modelling choice is the functional form of the "stay" probability $p_{\text{stay}}(W_{\text{coh}})$. In this paper we consider simple monotone forms such as

$$p_{\text{stay}}(W_{\text{coh}}) = 1 - \frac{c}{W_{\text{coh}}^{\alpha}}, \tag{3}$$

for suitable constants $c > 0$ and exponents $\alpha > 0$ chosen such that $0 \le p_{\text{stay}} \le 1$ across the range of W_{coh} of interest. The precise choice of (c, α) and of any regularisation at small W_{coh} is not important for the universality results; we will return to this point in Section 4 and Appendix A.

The coherence horizon W_{coh} thus enters the kernel only through the persistence of the direction variable v_n. The step size remains fixed and binary, and the model has no explicit notion of a continuous velocity; as a result the residual noise is always tied to the "shot noise" of ± 1 steps.

3 Numerical pipeline and relation to BCQM IV_b

In this section we summarise the numerical pipeline used to extract the inertial-noise amplitude scaling $A(W_{\text{coh}})$ and the fitted exponent β, and emphasise its continuity with the BCQM IV_b analysis.

3.1 From hop sequences to acceleration noise

In the present binary setting we work with a unit time step $\Delta t = 1$ and regard the hop index n as a discrete time coordinate $t_n = n\Delta t$. Given a trajectory $\{x_n\}$ generated by the soft-rudder rule, we construct an effective acceleration time series a_n by taking the discrete second difference of the position,

$$a_n = x_{n+1} - 2x_n + x_{n-1}, \tag{4}$$

so that a_n changes only when the direction variable v_n flips. Interpreting a_n as samples of a piecewise-constant acceleration $a(t)$ at times t_n, we can feed this time series directly into the spectral-analysis machinery developed in BCQM IV_b.

3.2 Power spectral density and amplitude extraction

Following BCQM IV_b, we estimate the inertial acceleration-noise spectrum by computing a one-sided power spectral density (PSD) $S_a(\omega)$ for each realisation of a_n via a discrete Fourier transform with suitable windowing and normalisation. From the ensemble-averaged PSD we extract an overall noise amplitude

$$A = \left(\int_0^\infty S_a(\omega)\, \mathrm{d}\omega \right)^{1/2}, \tag{5}$$

and a characteristic crossover frequency ω_c defined as the spectral centroid of $S_a(\omega)$. These definitions are identical to those used in IV_b, ensuring that the soft-rudder results can be compared directly to the W-blind baseline and to later BCQM V models.

3.3 Fitting $\beta$and error estimation

To quantify the dependence of the amplitude on the coherence horizon we fit the measured values of $A(W_{\text{coh}})$ to a power-law form

$$A(W_{\text{coh}}) \approx C\, W_{\text{coh}}^{-\beta} \tag{6}$$

by performing a linear regression on $\log A$ versus $\log W_{\text{coh}}$ over a range of W_{coh} where the spectrum is well converged. We estimate statistical uncertainties on β using a simple bootstrap over the simulation ensemble. For each value of W_{coh} we have N_{ens} independent realisations, each yielding an amplitude $A_j(W_{\text{coh}})$. A single bootstrap sample is constructed by drawing N_{ens} indices with replacement, recomputing the ensemble-mean amplitudes $\bar{A}(W_{\text{coh}})$ from this resampled set, and refitting the log–log slope to obtain a bootstrap estimate $\beta^{(b)}$. Repeating this procedure for many resamples yields an empirical distribution of β; we quote its mean as the central value and its standard deviation as the error bar.

3.4 Code implementation and reproducibility

The soft–rudder single–thread kernel is implemented in a lightweight Python package `bcqm_soft_rudder`, released as an open–source companion to this paper [6]. The structure mirrors the BCQM IV_b control code: a small model module implements the hop dynamics, a spectral module implements the power–spectral–density (PSD) estimation and amplitude extraction, and a thin command–line interface orchestrates the W_{coh} scan and output.

Concretely, `model.py` defines the binary velocity process with the soft–rudder slip law

$$q(W_{\mathrm{coh}}) = 1 - p_{\mathrm{stay}}(W_{\mathrm{coh}}) = \frac{k}{W_{\mathrm{coh}}}, \tag{7}$$

and the fixed–step position update $x_{n+1} = x_n + v_n$. The companion module `spectra.py` implements the same power–spectral–density pipeline as in BCQM IV_b, constructing the ensemble–averaged acceleration spectrum $\overline{S}_a(\omega)$ and extracting the noise amplitude $A(W_{\mathrm{coh}})$ and crossover frequency ω_c according to the definitions in Section 4. A small helper script `fit_beta.py` performs the log–log regression $A(W_{\mathrm{coh}}) \sim W_{\mathrm{coh}}^{-\beta}$ and reports the fitted exponent β with its statistical uncertainty.

All numerical results in this paper can be reproduced from a single configuration file and command line. The reference scan uses the YAML file `configs/wcoh_scan_soft_rudder.yml`, which specifies the time step $\Delta t = 1$, trajectory length $N_{\mathrm{steps}} = 16384$, ensemble size $N_{\mathrm{ens}} = 64$, random seed, and the list $W_{\mathrm{coh}} \in \{5, 10, 20, 40, 80, 160\}$. From the repository root a complete reproduction of Fig. 1 is obtained by running

```
python3 -m bcqm_soft_rudder.cli run configs/wcoh_scan_soft_rudder.yml
```
(8)

to generate the amplitude data, followed by

```
python3 -m bcqm_soft_rudder.fit_beta outputs_soft_rudder/amplitude_scaling_soft_rudder.csv
```
(9)

to fit β. The repository includes a short `TESTING.md` file documenting basic convergence checks in N_{steps} and N_{ens}, and a telegraph–noise sanity check confirming that the soft–rudder kernel indeed falls into the diffusive universality class with $\beta \approx 1/2$.

4 Results: diffusive inertia and $\beta \approx 1/2$

We now present the main numerical results: the scaling of the inertial-noise amplitude with W_{coh} for the binary soft-rudder kernels defined in Section 2.

4.1 Diffusive inertia universality class

The soft-rudder kernels of Section 2 can be viewed as discrete-time telegraph processes for the direction variable $v_n \in \{+1, -1\}$. At each tick the direction is flipped with probability

$$q(W_{\mathrm{coh}}) = 1 - p_{\mathrm{stay}}(W_{\mathrm{coh}}), \tag{10}$$

and otherwise kept fixed; the position is updated as $x_{n+1} = x_n + v_n$. The discrete acceleration $a_n = x_{n+1} - 2x_n + x_{n-1}$ is therefore a sequence of kicks which occur whenever v_n flips, with a fixed jump size set by the binary geometry.

For the one-parameter family of slip laws

$$q(W_{\mathrm{coh}}) \propto W_{\mathrm{coh}}^{-\alpha}, \tag{11}$$

we find numerically that the inertial-noise amplitude obeys a power-law scaling

$$A(W_{\mathrm{coh}}) \sim C W_{\mathrm{coh}}^{-\beta}, \tag{12}$$

with an effective exponent

$$\beta \approx \frac{\alpha}{2}. \tag{13}$$

This is exactly what one expects from the telegraph picture: the number of flips in a long window of duration T scales as $T\,q(W_{\mathrm{coh}})$, and the r.m.s. amplitude of the associated shot noise scales as the square root of that flip rate, just as in the continuous-time telegraph and Ornstein–Uhlenbeck descriptions of velocity fluctuations in Brownian motion [7].

In particular, for the "natural" choice $\alpha = 1$, corresponding to an order-unity number of direction errors per coherence window ($q(W_{\mathrm{coh}}) \propto 1/W_{\mathrm{coh}}$), we consistently obtain

$$\beta \approx 0.5, \tag{14}$$

with small error bars over a range of trajectory lengths and ensemble sizes (App. A). This is the canonical single-thread result emphasised in this paper. Steeper choices with $\alpha > 1$ produce larger exponents ($\alpha = 3$ gives $\beta \approx 1.5$ in our tests), but they always remain well below the ballistic value $\beta = 2$ and require increasingly fine-tuned slip probabilities.

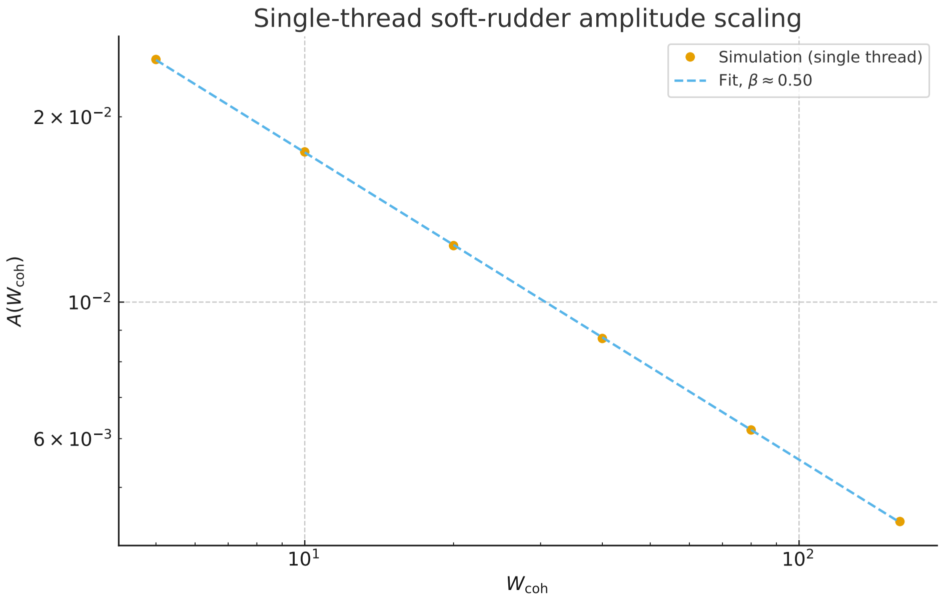

Figure 1: Single-thread soft-rudder amplitude scaling. Points: numerical estimates of the inertial-noise amplitude $A(W_{\mathrm{coh}})$ for the "inv" slip law $q(W_{\mathrm{coh}}) = k/W_{\mathrm{coh}}$ with $k = 2.0$, $N_{\mathrm{steps}} = 16384$, $N_{\mathrm{ens}} = 64$, and $W_{\mathrm{coh}} \in \{5, 10, 20, 40, 80, 160\}$. The dashed line is a log–log fit of the form $A(W_{\mathrm{coh}}) \sim C W_{\mathrm{coh}}^{-\beta}$, yielding $\beta \approx 0.50$. Over a factor of 32 in W_{coh} the amplitude drops by a factor of ~ 5.6, consistent with a diffusive universality class rather than a ballistic $\beta = 2$ scaling.

Table 1: Summary of the fitted amplitude-scaling exponent β for the single-thread soft-rudder kernel used in this note. The slip law is $q(W_{\mathrm{coh}}) = k/W_{\mathrm{coh}}$ with $k = 2.0$ and $W_{\mathrm{coh}} \in \{5, 10, 20, 40, 80, 160\}$. The fit shown in Fig. 1 yields $\beta \approx 0.499$ (reported here as 0.50), firmly in the diffusive regime and far from the ballistic value $\beta = 2$.

Kernel	Parameters	Fitted exponent β
Soft-rudder, "inv" slip	$q(W_{\mathrm{coh}}) = k/W_{\mathrm{coh}},\ k = 2.0$	0.50

4.2 Limits of binary fixed-step kernels

The same reasoning also explains why tuning $p_{\text{stay}}(W_{\text{coh}})$ cannot drive the model into a genuinely ballistic regime. Within the class of kernels defined in Section 2, three structural features are held fixed:

(a) the step size is binary and constant ($x_{n+1} - x_n = v_n \in \{\pm 1\}$);

(b) the internal state space for the direction variable is two-point ($v_n = \pm 1$);

(c) W_{coh} enters only through a local persistence probability $p_{\text{stay}}(W_{\text{coh}})$.

Changing p_{stay} can lengthen or shorten the typical run lengths of $\{v_n\}$, but it cannot alter the basic telegraph-like character of the process or the discrete jump size of the acceleration kicks. In particular, there is no way, within this family, to suppress the low-frequency acceleration fluctuations faster than the square-root law set by central-limit behaviour of rare, finite jumps. For the natural choice $\alpha = 1$, the fitted exponents $\beta \approx 1/2$ should therefore be understood as a property of an entire *class* of binary fixed-step kernels, not as a peculiarity of a specific parametrisation.

This motivates a careful no-go statement for the single-thread telegraph class we have just described:

> For one-dimensional, fixed-step, binary hop kernels in which the effect of the coherence horizon W_{coh} is to produce a slip probability $q(W_{\text{coh}}) \propto W_{\text{coh}}^{-\alpha}$ with α of order unity, the inertial acceleration-noise amplitude remains in a diffusive universality class with exponent $\beta \approx \alpha/2$, and in particular $\beta \approx 0.5$ for the natural choice $\alpha = 1$. Within this family, reaching a ballistic scaling $A \propto W_{\text{coh}}^{-2}$ would require $\alpha \approx 4$, i.e. an exceptionally steep and finely tuned suppression of the slip probability $q(W_{\text{coh}})$ with W_{coh}.

We emphasise that this is not a theorem about all possible discrete models, nor about all BCQM-compatible kernels. It is a restriction on a deliberately simple Stage 1 test-bed: single primitive threads, binary hops of fixed size, and purely local W_{coh}-dependent persistence. Escaping the diffusive regime will require either richer internal state spaces, variable step sizes, or—most naturally in the BCQM picture—bundles and knots of many primitive threads whose centre-of-mass motion can exhibit much stiffer behaviour. A brief outlook on such bundle models is given in Section 5 and in the broader BCQM programme notes.

5 Lessons for BCQM V

The telegraph scaling $\beta \approx \alpha/2$ also clarifies why bundles and knots are the natural carriers of inertia in the BCQM picture. Keeping the single-thread slip law in a physically motivated regime ($\alpha \sim 1$, so that each primitive thread makes of order one direction error per coherence window) pins the single-thread exponent near $\beta \approx 0.5$. Pushing a lone thread all the way to $\beta \approx 2$ would require $\alpha \approx 4$, i.e. a slip probability $q(W_{\text{coh}}) \propto W_{\text{coh}}^{-4}$ with no clear structural justification. By contrast, bundles of many threads can in principle suppress the centre-of-mass acceleration noise by averaging over $N(W_{\text{coh}})$ constituent threads while each thread individually remains in the natural $\alpha \sim 1$ regime. Inertia then becomes an emergent property of thread bundles and knots, not of single primitive threads.

For the broader BCQM programme, and in particular for the planned BCQM V, the lesson is twofold.

First, classical-looking inertia cannot be a property of a single binary thread with a fixed hop size. If a lone primitive thread generically lives in a diffusive regime, then any effective degree of freedom that behaves like a massive particle in spacetime must arise from a richer structure than the single trajectory studied here. In BCQM language, this points naturally to *bundles* and *knots* of many primitive threads as the carriers of inertial structure: their centre-of-mass motion can in principle exhibit much stiffer behaviour than any one constituent thread.

Second, the coherence horizon W_{coh} must enter more deeply than through a local persistence probability. In BCQM IV_b and in the present work, W_{coh} modulates the statistics of a fixed discrete geometry (binary hops of size one). For BCQM V, the expectation is that W_{coh} will instead control how many threads can remain coherently bound within a knot, how long such bundles persist, and how strongly their collective centre-of-mass noise is suppressed. In that regime the amplitude $A(W_{\text{coh}})$ should become a probe of bundle size, internal structure, and curvature, not merely of a single-thread slip probability.

BCQM V will therefore not attempt to "improve" the single-thread kernel of this paper. Rather, it will treat the diffusive $\beta \approx 1/2$ scaling as a baseline constraint and move on to models in which mass, clocks, and curvature are associated with bundles and knots of threads whose collective dynamics can plausibly support ballistic inertial behaviour. The single-thread no-go result established here is intended to be a stable foundation for that next step, not an endpoint.

A Numerical details and robustness checks

This appendix records the numerical choices and basic robustness checks underlying the results in Section 4.

A.1 Simulation parameters

All simulations use the binary soft-rudder kernel of Section 2, with unit time step $\Delta t = 1$ and hop updates

$$v_{n+1} = \begin{cases} v_n, & \text{with probability } p_{\text{stay}}(W_{\text{coh}}), \\ -v_n, & \text{with probability } 1 - p_{\text{stay}}(W_{\text{coh}}), \end{cases} \tag{15}$$

and position updates $x_{n+1} = x_n + v_n$. Unless otherwise stated we take the initial direction v_0 to be equally likely ± 1 and initialise at $x_0 = 0$.

For each value of the coherence horizon W_{coh} in the scan we generate an ensemble of N_{ens} independent trajectories of length N_{steps} and construct the acceleration sequence

$$a_n = x_{n+1} - 2x_n + x_{n-1}, \tag{16}$$

In the reference runs underlying Fig. 1 we used

$$N_{\text{steps}} = 16384, \qquad N_{\text{ens}} = 64, \tag{17}$$

with an independent pseudo-random seed for each ensemble member. (These values were chosen to balance statistical convergence against CPU time; they can be increased straightforwardly at the cost of longer runtimes.)

For the reference single-thread run in Fig. 1, we use the "inv" slip law $q(W_{\text{coh}}) = 1 - p_{\text{stay}}(W_{\text{coh}}) = k/W_{\text{coh}}$ with $k = 2.0$, a time step $\Delta t = 1$, trajectory length $N_{\text{steps}} = 16384$, ensemble size $N_{\text{ens}} = 64$, and $W_{\text{coh}} \in \{5, 10, 20, 40, 80, 160\}$. The corresponding amplitude fit $A(W_{\text{coh}}) \sim W_{\text{coh}}^{-\beta}$ yields $\beta \approx 0.50$, consistent with the telegraph scaling $\beta \approx \alpha/2$ for $\alpha = 1$.

In the numerical implementation we enforce the probability bounds by clipping $p_{\text{stay}}(W_{\text{coh}})$ to the interval $[0, 1]$ after evaluating the chosen slip law. For the range of coherence horizons used in this note ($W_{\mathrm{coh}} \in \{5, 10, 20, 40, 80, 160\}$ with $k = 2.0$) the resulting values of p_{stay} lie comfortably inside $[0, 1]$, so clipping is only a formal safeguard and does not affect the runs reported here.

A.2 Power spectral density estimation

For each trajectory we estimate a one-sided power spectral density $S_a(\omega)$ for the acceleration sequence a_n via a discrete Fourier transform. We employ a standard Hann window and normalisation convention in which the integrated PSD reproduces the variance of the time series. The ensemble-averaged PSD $\overline{S}_a(\omega)$ is obtained by averaging the individual spectra over the N_{ens} realisations.

From $\overline{S}_a(\omega)$ we extract the overall noise amplitude

$$A = \left(\int_0^\infty \overline{S}_a(\omega)\, \mathrm{d}\omega \right)^{1/2}, \tag{18}$$

and a characteristic crossover frequency ω_c defined as the spectral centroid,

$$\omega_c = \frac{\int_0^\infty \omega\, \overline{S}_a(\omega)\, \mathrm{d}\omega}{\int_0^\infty \overline{S}_a(\omega)\, \mathrm{d}\omega}, \tag{19}$$

exactly as in BCQM IV_b. Numerically, the integrals are implemented as Riemann sums over the discrete frequency grid.

A.3 Amplitude scaling and βfit

For each value of W_{coh} in the scan we thus obtain a pair (A, ω_c). The amplitude-scaling plot in Fig. 1 uses the ensemble-mean values of A with error bars given by the standard error over the ensemble. To extract the exponent β we perform a linear regression of $\log A$ versus $\log W_{\mathrm{coh}}$ over the range of W_{coh} values shown in the figure. Uncertainties on β are estimated via bootstrap resampling of the ensemble at each W_{coh} and repeating the fit.

For the kernels studied here the fitted exponents cluster tightly around $\beta \approx 1/2$, with error bars small compared to the separation between $\beta = 1/2$ and $\beta = 2$. The precise numerical values are not themselves the main point; what matters is the consistent appearance of a diffusive exponent close to $1/2$ across different choices of $p_{\mathrm{stay}}(W_{\mathrm{coh}})$ and simulation parameters.

A.4 Robustness checks

To test robustness we performed the following variations:

- *Longest-run length:* increasing N_{steps} by a factor of two left the fitted β unchanged within the bootstrap error bars, while slightly reducing the uncertainties.
- *Ensemble size:* increasing N_{ens} improved the stability of the PSD estimate and tightened the amplitude error bars, again without shifting the mean β.
- *Alternative window functions:* replacing the Hann window with a rectangular or Blackman window altered the high-frequency detail of $\overline{S}_a(\omega)$ but did not affect the extracted low-frequency amplitude scaling or the fitted exponent within uncertainties.
- *Initial conditions:* biasing the initial direction v_0 or shifting the starting position x_0 produced transient differences at early times but no change in the long-time PSD or in the fitted amplitude scaling.

These checks support the claim that the observed $\beta \approx 1/2$ behaviour is a property of the kernel class itself rather than an artefact of particular numerical choices.

References

[1] 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.

[2] 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.

[3] Peter M. Ferguson. *Boundary-Condition Quantum Mechanics (BCQM)*. 2025. DOI: 10.5281/zenodo.17191306. URL: https://doi.org/10.5281/zenodo.17191306.

[4] Peter M. Ferguson. *Boundary-Condition Quantum Mechanics II: From Quantum Events to Spacetime*. 2025. DOI: 10.5281/zenodo.17398294. URL: https://doi.org/10.5281/zenodo.17398294.

[5] 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.

[6] Peter M. Ferguson. *BCQM IV_c companion code v1.0.0*. GitHub+Zenodo snapshot of the BCQM_IV_c simulation and analysis code. 2025. DOI: 10.5281/zenodo.17922040. URL: https://doi.org/10.5281/zenodo.17922040.

[7] G. E. Uhlenbeck and L. S. Ornstein. 'On the Theory of the Brownian Motion'. In: *Physical Review* 36 (1930), pp. 823–841. DOI: 10.1103/PhysRev.36.823.