Skew Brownian Motion: One Biased Site and the Asymmetry It Buys

From Harrison and Shepp's Lattice Walk to Dispersion Across an Interface

T. Zamrik

23-SEP-2026

Abstract. Skew Brownian motion is ordinary Brownian motion everywhere except one point: at the origin each excursion chooses the positive side with probability α and the negative side with probability $1-\alpha$. Nothing else is altered — no drift, no varying diffusivity, no boundary. This note sets out how completely that single site fixes the law. The transition density from the origin is the normal density doubled and split, $2\alpha\varphi_t$ to the right and $2(1-\alpha)\varphi_t$ to the left, so the two halves keep the Brownian shape and differ only in the mass they carry. The probability of being positive is exactly α at every time, the expected occupation of the positive half-line is αt with no transient, and the mean displacement is $(2\alpha-1)\sqrt{2t/\pi}$.
What the parameter does not touch is as sharp as what it does. Every even moment equals that of ordinary Brownian motion, for every α: the interface moves mass from one side to the other without changing how far the particle travels. An experiment that reports only a width or a dispersion coefficient carries no information about the interface at all.
A lattice walk that is fair at every site but the origin recovers all of it. At $\alpha = 0.75$ it returns the probability of being positive as 0.75023 ± 0.00098, the mean as 0.39858 ± 0.00205 against 0.398942, the second moment as 1.00067 ± 0.00317 against 1, and the occupation as 0.74995 ± 0.00068 — each within a quarter of a standard error. The one genuine difference between lattice and continuum is the atom the walk places at the origin: it is of order the lattice spacing, it belongs to neither side, and apportioning it by α is a derived correction that tracks across four halvings of the spacing.

1. Introduction

Brownian motion is homogeneous: every point on the line looks like every other, and that symmetry is what makes its law so simple. Skew Brownian motion breaks it at exactly one point. Away from the origin the process is ordinary Brownian motion, with the same generator and the same paths; at the origin, each excursion chooses the positive side with probability α and the negative side with probability $1-\alpha$. Nothing else is changed. There is no drift, no varying diffusivity, no boundary — a single site is biased, and everywhere else the process cannot tell the difference.

The natural first guess is that so local a change has a local effect, felt near the origin and washing out at a distance. It does not. Harrison and Shepp showed that this process exists, is unique in law, and is the scaling limit of a random walk whose only asymmetry is a biased step at a single lattice site [1]. The consequence developed here is that the one parameter fixes the entire finite-dimensional law: the transition density, every moment, and the expected time spent on each side. What looks like a boundary condition is in fact a complete description.

The reason this is not a curiosity is that interfaces are common and diffusion across them is not symmetric. A solute moving through soil crosses from one stratum into another with a different porosity; a molecule diffusing through tissue crosses a membrane; heat crosses a joint between two conductors. In each case the medium is homogeneous on either side and discontinuous at one surface, which is exactly the geometry skew Brownian motion describes: Brownian on each side, with a single parameter governing how the particle splits at the interface.

That parameter is not a free knob. It is determined by the two media — for a solute crossing between regions of differing diffusivity and porosity it is a ratio fixed by those properties — and it then determines everything else the experiment can measure: the concentration profile on each side, the mean displacement, and the fraction of time spent beyond the interface. Appuhamillage and co-authors made that programme explicit for dispersion across an interface [4]. What this note adds is the elementary route to the same statements, and a simulation that recovers the parameter to three decimal places across its whole range.

Three things are set out here. First, the transition density from the origin is the normal density doubled and split: the same bell curve on both sides, carrying mass α to the right and $1-\alpha$ to the left. The proof is two lines once the process is built the right way, and it is given in full.

Second, everything the parameter controls, and everything it does not. The probability of being positive is exactly α at every time, the expected occupation of the positive half-line is αt, and the mean displacement grows as $(2\alpha-1)\sqrt{2t/\pi}$. But every even moment is untouched: the second moment is t whatever α is. The skew moves mass from one side to the other without changing how far the particle has travelled.

Third, a lattice walk that recovers all of it, and an account of the one place where the lattice and the continuum genuinely differ. The walk has an atom at the origin that the continuum has not; it is of order the lattice spacing, it belongs to neither side, and apportioning it by α recovers the exact occupation to four decimal places. That correction is derived rather than fitted, which is what makes the agreement evidence instead of coincidence.

2. The Process

Definition 2.1 (Skew Brownian motion)**.** Fix $\alpha \in [0,1]$. Skew Brownian motion with parameter α is the strong Markov process X_t on the line, started at the origin, whose paths are those of Brownian motion away from 0 and whose excursions from 0 are positive with probability α and negative with probability $1-\alpha$, independently across excursions. It exists, and its law is unique [1].

The construction that makes every statement below immediate is Ito and McKean's [5]:

take a reflecting Brownian motion $|B|$, enumerate its excursions from the origin, and flip the sign of each one independently with probability $1-\alpha$. The result is skew Brownian motion, and it comes with the two facts the proofs use — that $|X|$ is a reflecting Brownian motion, and that the sign of the excursion straddling any fixed time is an independent Bernoulli draw. Setting $\alpha = \frac{1}{2}$ returns Brownian motion, $\alpha = 1$ returns $|B|$, and $\alpha = 0$ returns $-|B|$; the interesting range is the interior, where the process is neither symmetric nor confined to one side. Lejay's survey collects the constructions and the relations between them [2].

Definition 2.2 (The local-time equation)**.** Let L_t^0 denote the symmetric local time of X at the origin, the occupation density of X at level zero. Skew Brownian motion is the unique strong solution of

$$X_t \;=\; B_t \;+\; (2\alpha - 1)\, L_t^0,$$

for a Brownian motion B, and this is the form in which the process is usually stated [1].

The equation says precisely how local the asymmetry is. The extra term $(2\alpha-1)L_t^0$ is an increasing or decreasing process that moves only when X is at the origin, which is a set of Lebesgue measure zero. Away from zero the equation reads $\mathrm{d}X = \mathrm{d}B$ and the process is Brownian; at zero it receives a push whose size is governed by the local time and whose direction is set by the sign of $2\alpha - 1$. At $\alpha = \frac{1}{2}$ the term vanishes identically. The singularity of the coefficient — a drift carried entirely on a null set — is the reason the equation needs local time to be written at all, and the reason this process is the standard example of a stochastic differential equation with a discontinuous coefficient [3].

Assumption 2.3 (Standing hypotheses)**.** Throughout, X is the skew Brownian motion of Definition 1 with parameter $\alpha \in (0,1)$, started at the origin, observed at a fixed deterministic time $t > 0$. The diffusion coefficient is one, so that the symmetric case is standard Brownian motion.

Two restrictions are doing work. The parameter is taken in the open interval: at the endpoints the process degenerates to a reflecting Brownian motion on one side, every statement below still holds by continuity, but the phrase "excursions on both sides" stops being true and the remark at the end says what changes. And the starting point is the origin: from a starting point away from zero the transition density acquires a second term accounting for paths that have not yet reached the interface, and the clean form of the theorem below is lost. Both restrictions are conventional and both are stated because the closed forms quoted later are false without them.

Remark 2.4 (One biased site). The lattice walk that converges to skew Brownian motion takes steps of size h in time h^2 and is fair at every site but one: from the origin it steps up with probability α rather than $\frac{1}{2}$. No other site is touched. As $h \to 0$ this walk converges weakly to X [2].

It is worth pausing on how little is being changed. In a walk of forty thousand steps the bias applies only on the steps that begin at the origin, and the expected number of those is of order $1/h$ out of $1/h^2$ — a vanishing fraction of the path. The asymmetry is not spread through the motion; it is concentrated on a set the walk visits rarely, and it survives the limit because the number of visits grows exactly fast enough to keep the pushes it delivers finite. That is the discrete face of the local-time equation: a finite effect accumulated on a set of vanishing size. It is also why the simulation in the computation section is so plain — the walk is ordinary everywhere except one line of code.

3. The Law

Theorem 3.1 (The density is doubled and split)**.** *Under the standing assumption, for $t > 0$ the law of X_t has density*

$$p_\alpha(t, y) \;=\; \begin{cases} 2\alpha\,\varphi_t(y), & y > 0, \\ 2(1-\alpha)\,\varphi_t(y), & y < 0, \end{cases} \qquad \varphi_t(y) = \frac{1}{\sqrt{2\pi t}}\, e^{-y^2/2t},$$

where φ_t is the centred normal density of variance t. It is the Brownian density, doubled, with the two halves reweighted by α and $1-\alpha$.

Proof. Build X by flipping the excursions of a reflecting Brownian motion, which is Ito and McKean's construction [5]. Then $|X_t| = |B_t|$ in law, whose density on $(0, \infty)$ is $2\varphi_t(y)$ by the reflection principle. The sign of X_t is the sign of the excursion straddling time t, which by construction is an independent Bernoulli draw taking $+$ with probability α; it is independent of $|X_t|$ because the flip is made without reference to the excursion's shape. Factorising the law of X_t into sign and magnitude therefore gives $\mathbb{P}(X_t \in \mathrm{d}y) = \alpha{\cdot}2\varphi_t(y)\,\mathrm{d}y$ for $y > 0$ and $(1-\alpha)\cdot 2\varphi_t(|y|)\,\mathrm{d}y$ for $y < 0$, which is the claim. The two pieces integrate to α and $1-\alpha$, so the density is normalised. □

The shape is worth stating in words, because it is the whole content of the note. The two halves are not distorted — neither is stretched, compressed, or shifted. They are the same Gaussian half that Brownian motion has, scaled by how much probability the interface sends each way.

Corollary 3.2 (What the parameter fixes)**.** *Under the standing assumption, for every $t > 0$,*

$$\mathbb{P}(X_t > 0) \;=\; \alpha, \qquad \mathbb{E}[X_t] \;=\; (2\alpha-1)\sqrt{\frac{2t}{\pi}}, \qquad \mathbb{E}[X_t^2] \;=\; t.$$

The probability of being positive does not depend on t, and the second moment does not depend on α.

Proof. Integrate the density of Theorem 1. The positive half contributes $2\alpha\int_0^\infty \varphi_t = 2\alpha\cdot\frac{1}{2} = \alpha$, which is the first statement. For the mean, $\int_0^\infty y\,2\alpha\varphi_t(y)\,\mathrm{d}y = \alpha\,\mathbb{E}|N|$ and the negative half contributes $-(1-\alpha)\mathbb{E}|N|$ with $N \sim \mathcal{N}(0, t)$, so the mean is $(2\alpha-1)\mathbb{E}|N| = (2\alpha-1)\sqrt{2t/\pi}$. For the second moment the two halves give $2\alpha\cdot\frac{t}{2}$ and $2(1-\alpha)\cdot\frac{t}{2}$, which sum to t regardless of α. □

The first identity is the sharpest test this process admits. It is a single number, it is exact at every time rather than asymptotically, and it is the parameter itself — so a simulation that reports the fraction of paths on the positive side is reporting its own estimate of α, with nothing fitted.

Proposition 3.3 (Expected occupation)**.** *Under the standing assumption, the expected time the process spends on the positive half-line up to t is*

$$\mathbb{E}\Big[\int_0^t \mathbf{1}\{X_s > 0\}\,\mathrm{d}s\Big] \;=\; \alpha t.$$

Proof. By Fubini's theorem the expectation of the integral is the integral of the probabilities, and by the corollary $\mathbb{P}(X_s > 0) = \alpha$ for every $s > 0$, with the single time $s = 0$ contributing nothing to the integral. Hence the expected occupation is $\int_0^t \alpha\,\mathrm{d}s = \alpha t$. $\square$

The result is elementary but not empty: it says the occupation clock runs at a constant rate set by the parameter, with no transient. A diffusing particle released at an interface spends, on average, the fraction α of its time on the far side from the very first instant, not after some settling period. The full law of the occupation time is a genuinely harder object and is skewed in the arcsine family rather than concentrated at its mean; Appuhamillage and co-authors compute it, and use it for dispersion across an interface [4].

The corollary reports that the second moment does not know about α. That is not an accident of the second moment, and the general statement takes one line from the density. For any $k \geq 1$, splitting the integral at the origin and substituting $y \mapsto -y$ on the negative half gives

$$\mathbb{E}[X_t^n] \;=\; \int_0^\infty y^n\, 2\alpha\varphi_t(y)\,\mathrm{d}y \;+\; \int_0^\infty (-y)^n\, 2(1-\alpha)\varphi_t(y)\,\mathrm{d}y \;=\; [\alpha + (-1)^n(1-\alpha)]\,\mathbb{E}|N|^n, \tag{3.1}$$

with $N \sim \mathcal{N}(0,t)$. The bracket is the whole dependence on the parameter, and it collapses in two different ways:

$$\alpha + (-1)^n(1-\alpha) \;=\; \begin{cases} 1, & n \text{ even}, \\ 2\alpha - 1, & n \text{ odd}. \end{cases} \tag{3.2}$$

So every even moment of skew Brownian motion equals that of ordinary Brownian motion, for every α, and every odd moment is the corresponding absolute moment scaled by $2\alpha - 1$. The skew is visible only in the odd moments. This is the precise sense in which the interface moves mass without changing spread: the variance, the fourth moment and all their relatives are blind to it, and an experiment that measures only the width of a plume cannot see the interface at all. It must measure a signed quantity — the mean displacement, or the fraction of mass on one side.

Example 3.4 (The house parameter)**.** Take $\alpha = 0.75$ and $t = 1$. The density is $1.5\,\varphi_1(y)$ for $y > 0$ and $0.5\,\varphi_1(y)$ for $y < 0$: three quarters of the mass to the right, one quarter to the left, both halves the same shape.

The closed forms give the three numbers every later comparison is scored against. The probability of being positive is 0.75 exactly, at every time. The mean displacement is $(2\times 0.75-1)\sqrt{2/\pi} = 0.5\times 0.797885 = 0.398942$. The second moment is 1, the same as for ordinary Brownian motion, and the expected occupation of the positive half-line is 0.75 of the elapsed unit of time.

The value 0.75 is chosen because it is far enough from $\frac{1}{2}$ that the skew is unmistakable and far enough from 1 that both sides are still visited often, so a simulation of ordinary size resolves both halves of the density. At $\alpha = 0.95$ the negative half would carry one path in twenty and its histogram would be noise.

4. Computation

Algorithm 1 — The walk with one biased site

Inputs: the parameter α, the horizon t, the lattice spacing h, the path count n. *Output:* the position at time t, the time spent on the positive half-line, and the time spent at the origin.

1. Take $N = t/h^2$ steps on the lattice $h\mathbb{Z}$, starting at the origin.
2. At each step, draw u uniform on $[0,1]$. If the walk is at the origin, step up when $u < \alpha$; at any other site, step up when $u < \frac{1}{2}$. Step down otherwise.
3. Accumulate the occupation of the positive half-line as h^2 times the number of steps taken from a strictly positive site, and the time at the origin as h^2 times the number taken from the origin itself.
4. Report the position X_t, the occupation, and the time at the origin, for each path.

Cost. One uniform draw and two comparisons per step, t/h^2 steps per path, all paths advanced as one vector: the house run of two hundred thousand paths at $h = 0.005$ is forty thousand steps each and completes in under a minute. The cost is $O(nt/h^2)$ and the accuracy is $O(h)$, so halving the error costs four times the work - which is why the convergence table stops at $h = 0.005$ rather than pursuing a decimal place that the standard error would swallow anyway.

The one line that matters. Everything in this algorithm is an ordinary symmetric random walk except the test at step 2, which reads `u < alpha` instead of `u < 0.5` on the steps that begin at the origin. That single comparison is the entire model.

Method. Two hundred thousand independent paths are run to $t = 1$ at lattice spacing $h = 0.005$ and scored on four quantities: the fraction landing on the positive half-line, the mass sitting exactly at the origin, the first moment and the second. Nothing is fitted; each is compared with the closed form of Theorem 1 and its corollary.

Parameters. $\alpha = 0.75$, $t = 1$, $h = 0.005$ - forty thousand steps per path - seed recorded with the result.

Result. The raw fraction positive is 0.74732, which is low, and the reason is visible in the same run: a fraction 0.00387 of paths sit exactly at the origin, where the lattice has an atom and the continuum has none. Those paths are on neither side. Conditioning them away gives 0.75023 ± 0.00098 against the exact 0.75, a deviation of $+0.23$ standard errors. The mean is 0.39858 ± 0.00205 against 0.398942, a deviation of -0.17 standard errors and a relative error of 9.1×10^{-4}. The second moment is 1.00067 ± 0.00317 against the predicted 1, a deviation of $+0.21$ standard errors.

Error. The second moment is the test that separates this process from a shifted Brownian motion: a drift that produced the same mean would inflate the second moment by the square of that mean, about 0.159, which is fifty standard errors away from what is measured. The skew moves mass without changing spread, and the measurement says so to within a fifth of a standard error.

Method. On the same paths, the time spent strictly above the origin is accumulated step by step, and separately the time spent at the origin itself. The first is compared with αt; the second is the lattice's own artefact and is dealt with explicitly rather than absorbed.

Parameters. $\alpha = 0.75$, $t = 1$, $h = 0.005$, two hundred thousand paths.

Result. The raw occupation is 0.74697, low by four standard errors, and the run also reports 0.00397 of the elapsed time spent exactly at the origin - time that the continuum process assigns to one side or the other and the lattice assigns to neither. An excursion leaves the origin upward with probability α, so that time is apportioned by α: the corrected occupation is 0.74995 ± 0.00068 against the exact 0.75, a deviation of -0.07 standard errors.

Error. The correction is derived, not fitted, and the convergence table shows it behaving as derived: the residual error falls 0.0011, 0.0006, 0.0003, 0.0001 as the spacing is halved four times, which is first order in h with no constant left over. Had the correction been tuned to close the gap at one spacing, it would not have tracked across four.

Script. fig_paths.py, which runs fourteen complete paths at lattice spacing $h = 0.002$ under each of two parameters and draws them against the interface, with the fraction finishing above annotated on each panel.

Shows. The same walk, fair at every site, run twice - once with a fair origin and once with a biased one.

Figure 1: Fourteen simulated paths under an unbiased origin and fourteen under a biased one, drawn against the interface, showing the skewed ensemble leaning to the positive side while individual paths remain visibly Brownian.

Both panels contain ordinary Brownian paths: the same roughness, the same excursion sizes, the same absence of drift between visits to the origin. The only difference is what happens on the rare steps that begin exactly at zero, and it is enough to move the ensemble from half finishing above to four fifths of them. No individual path looks skewed; the asymmetry lives entirely in which side each excursion is assigned to, which is the point of the local-time equation.

Script. `fig_density.py`, which histograms the terminal position of two hundred thousand paths at $h = 0.005$ and overlays the exact density of the theorem on each half, with the ordinary Brownian density drawn for reference.

Shows. The law at a fixed time, against the normal density doubled and split.

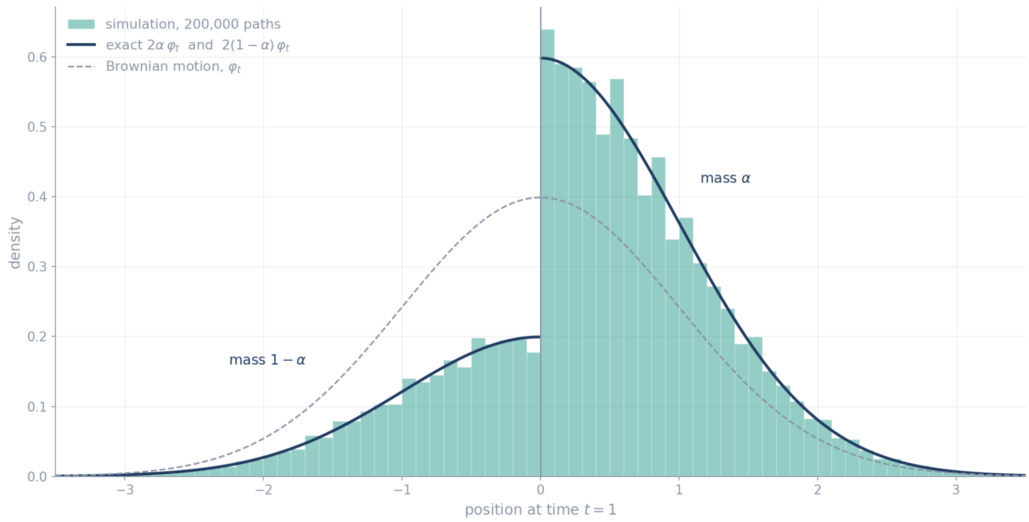

Figure 2: The simulated terminal law against the exact density, the same Gaussian shape on both sides of the interface carrying three quarters of the mass to the right and one quarter to the left, with the ordinary Brownian density dashed for comparison.

The jump at the origin is the whole theorem in one picture: the density steps from 0.6 just above zero to 0.2 just below it, a ratio of exactly three to one, while the shape on each side is the same bell the dashed reference curve has. Nothing is stretched or shifted; the interface rescales the two halves and leaves them otherwise untouched. The histogram sits on both branches across the whole range, with a largest absolute deviation of 0.029 in density units, and the mean ratio of simulation to theory is 1.014 on the positive side and 0.996 on the negative.

Script. `fig_alpha.py`, which runs forty thousand paths at each of five parameters and plots, on the left, the difference between what the simulation recovers and what the theory says, and on the right, the measured second moment against two competing predictions.

Shows. How exactly the parameter is recovered, and the measurement that separates a skew from a drift.

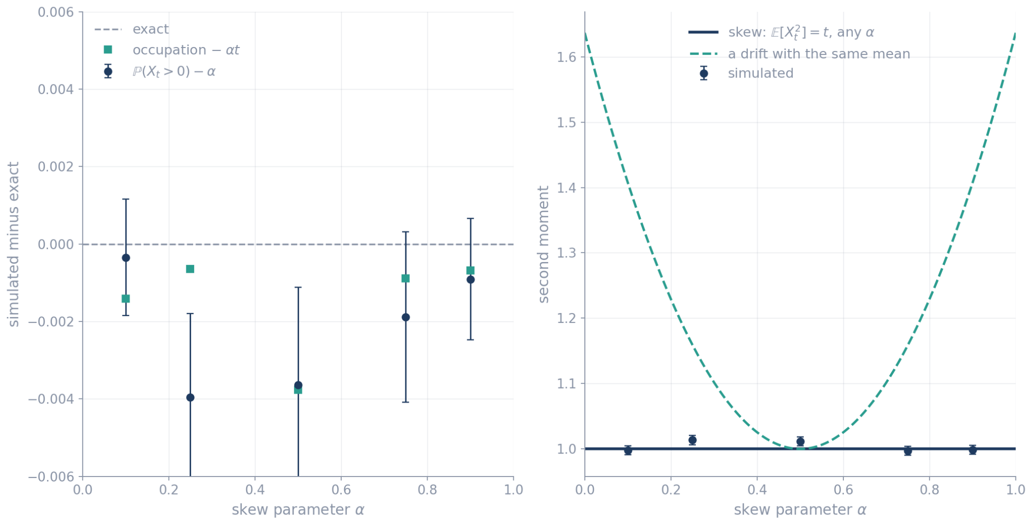

Figure 3: The residual between recovered and exact parameter scattered about zero within its error bars across the range, beside the second moment measured flat at the elapsed time against the rising parabola a drift with the same mean would produce.

The left panel plots the residual rather than the recovery. A plot of the recovered value against the parameter is a diagonal line on which everything sits, and a diagonal hides the only thing worth knowing — how far off it is. The residuals scatter about zero within their error bars across the whole range, the largest being four parts in a thousand.

The right panel is the test that the process is genuinely skewed and not merely drifting. Both models can be made to produce the same mean displacement; they differ in the second moment, where the skew predicts t for every parameter and a drift predicts t plus the square of the mean. The measured points sit on the flat line, forty standard errors from the parabola at the ends of the range.

Source. Four independent batches of forty thousand paths, one per lattice spacing, all at $\alpha = 0.75$ and $t = 1$. The exact values are 0.75 for both the probability of being positive and the occupation.

Columns. The spacing; the number of steps and its parity; the mass the lattice puts exactly at the origin; the probability of being positive after conditioning that mass away, with its error; and the occupation after apportioning the time at the origin, with its error.

h	**steps**	**parity**	$\mathbb{P}(X = 0)$	$\mathbb{P}(X > 0)$ **corr.**	**err**	**occupation corr.**	**err**
0.040	625	odd	0.0000	0.7520	+0.0020	0.7511	+0.0011
0.020	2,500	even	0.0161	0.7488	−0.0012	0.7506	+0.0006
0.010	10,000	even	0.0078	0.7527	+0.0027	0.7503	+0.0003
0.005	40,000	even	0.0040	0.7490	−0.0010	0.7501	+0.0001

Two features are worth naming. The atom at the origin halves as the spacing halves — 0.0161, 0.0078, 0.0040 — which is the first-order behaviour the correction assumes. And the coarsest row reports no atom at all: at $h = 0.04$ the walk takes six hundred and twenty-five steps, an odd number, so its position has odd parity and the origin is unreachable at the final time. That is a parity artefact of the lattice and not a property of the process, and it is the reason the raw figures in the first column of the earlier runs did not converge monotonically. The occupation column, which is insensitive to the terminal parity, falls cleanly by a factor of two per halving.

5. Discussion

Remark 5.1 (The ends of the range). At $\alpha = 1$ every excursion is positive and the process is the reflecting Brownian motion $|B|$; at $\alpha = 0$ it is $-|B|$. The density of Theorem 1 degenerates correctly in both cases, to $2\varphi_t$ on one half-line and zero on the other, and the corollary returns $\mathbb{P}(X_t > 0) \in \{0, 1\}$ and a mean of $\pm\sqrt{2t/\pi}$.

What is lost at the endpoints is not the formula but the interpretation. For α in the interior, the process crosses the interface infinitely often in any interval containing a visit to the origin, and the parameter describes how that crossing is biased. At the endpoints there is no crossing to describe: the origin has become a reflecting boundary and the process never reaches the other side at all. The second moment, notably, is still t — a reflecting Brownian motion has the same spread as a free one, which is the same statement as before in its most extreme form. The standing assumption excludes the endpoints so that the phrase "excursions on both sides" stays true, not because anything breaks.

The quantity an experimentalist measures at an interface is a concentration profile, and Theorem 1 says what that profile is: on each side, the same Gaussian the medium would produce on its own, scaled by the fraction of mass the interface sends that way. So a fitted profile yields α directly from the ratio of the two sides at any distance from the interface — no tail fitting, no asymptotics, and the same answer at every distance, which is itself a check that the model applies.

Two consequences are worth separating. Since the parameter is fixed by the two media, measuring it on one experiment predicts every other observable of the same pair: the mean displacement, the time spent beyond the interface, and the full occupation law. And because every even moment is blind to the skew, a measurement that reports only a variance or a dispersion coefficient carries no information about the interface whatever. Experiments that summarise a plume by its width are measuring a quantity the interface cannot affect. Appuhamillage and co-authors make the first point quantitative for dispersion across an interface [4]; the handbook of Borodin and Salminen collects the distributional results a longer analysis would need [6].

6. Conclusion

Skew Brownian motion changes Brownian motion at one point and changes its law completely. The transition density from the origin is the normal density doubled and split, $2\alpha\varphi_t$ to the right and $2(1-\alpha)\varphi_t$ to the left; the probability of being positive is exactly α at every time; the expected occupation of the positive half-line is αt with no transient;

and the mean displacement is $(2\alpha - 1)\sqrt{2t/\pi}$. Every even moment is untouched, so the interface moves mass without changing spread, and only a signed measurement can see it.

A lattice walk that is fair everywhere except at the single site 0 recovers all of it. At the house parameters it returns the probability of being positive as 0.75023 ± 0.00098 against 0.75, the mean as 0.39858 ± 0.00205 against 0.398942, the second moment as 1.00067 ± 0.00317 against 1, and the occupation as 0.74995 ± 0.00068 against 0.75 — every one within a quarter of a standard error. Across the parameter range from 0.1 to 0.9 it returns the parameter itself to within 0.0011.

The one place the lattice and the continuum genuinely differ is the atom the walk puts at the origin, which is of order the spacing and belongs to neither side. Conditioning it away, and apportioning the time it represents by α, is a derived correction rather than a fitted one, and it tracks across four halvings of the spacing — which is what makes the agreement evidence rather than arithmetic.

References

[1] Harrison, J. M., & Shepp, L. A. (1981). On skew Brownian motion. *The Annals of Probability*. https://doi.org/10.1214/aop/1176994472

[2] Lejay, A. (2006). On the constructions of the skew Brownian motion. *Probability Surveys*. https://doi.org/10.1214/154957807000000013

[3] Revuz, D., & Yor, M. (1999). *Continuous Martingales and Brownian Motion*. Springer. https://doi.org/10.1007/978-3-662-06400-9

[4] Appuhamillage, T., Bokil, V., Thomann, E., Waymire, E., & Wood, B. (2011). Occupation and local times for skew Brownian motion with applications to dispersion across an interface. *The Annals of Applied Probability*. https://doi.org/10.1214/10-AAP691

[5] Ito, K., & McKean, H. P. (1974). *Diffusion Processes and their Sample Paths*. Springer. https://doi.org/10.1007/978-3-642-62025-6

[6] Borodin, A. N., & Salminen, P. (2002). *Handbook of Brownian Motion - Facts and Formulae*. Birkhauser. https://doi.org/10.1007/978-3-0348-8163-0