Understanding Brownian Motion

Álvaro Martín Ruiz

February 25, 2026

1 What is Brownian Motion?

Brownian motion is a kind of stochastic process, that is, a way to model how a variable moves over time. Specifically, Brownian motion is the mathematical model of pure randomness evolving in time. It is the mathematical idealization of a path that moves continuously but in a completely unpredictable way, with no memory of where it has been.

Intuitive idea: Imagine a tiny particle suspended in water. It gets constantly hit by water molecules from all directions. Each hit is tiny, random, and unpredictable. The particle jitters randomly, changing direction all the time. Brownian motion models this limit behavior when hits happen extremely frequently and each hit is extremely small.

2 Formal definition

A Brownian motion $(W_t)_{t \geq 0}$ is a stochastic process such that:

- **Starts at zero:** $W_0 = 0$.
- **Independent increments:** Increments over non-overlapping time intervals are independent; that is, each increment is independent of what happened before.
- **Normally distributed increments:** For two time instants t and s ($0 \leq s < t$), the increment is Gaussian with mean 0 and variance $t - s$:

  $$W_t - W_s \sim \mathcal{N}(0, t - s) \iff \Delta W \sim \mathcal{N}(0, \Delta t)$$

  Properties: mean 0 (no drift), variance Δt, typical size $\sqrt{\Delta t}$.

  Concrete example: Let $s = 1, t = 1.01$. Then $t - s = 0.01$ and:

  $$W_{1.01} - W_1 \sim \mathcal{N}(0, 0.01) \quad \text{(typical size } \sqrt{0.01} = 0.1\text{)}$$

- **Continuous paths:** No jumps. [1]

[1] More precisely, Brownian motion has almost surely continuous paths. This means that with probability 1, a realized trajectory of Brownian motion is continuous and has no jumps. Properties are said to hold 'almost surely' if the set of exceptions has probability zero. In other words, while one can theoretically construct pathological outcomes where continuity fails, those outcomes are so negligible that they never occur in practice and have no impact on modeling or applications.

3 The Interesting Properties of Brownian Motion

3.1 Continuous but nowhere differentiable

Intuitively, we can say that a Brownian path:

- has no jumps
- but is so jagged that you cannot draw a tangent anywhere

Over a tiny time interval Δt:

$$\Delta W \sim \sqrt{\Delta t} \quad \text{more precisely, } \mathbb{E}[|\Delta W|] = \sqrt{\frac{2}{\pi}}\sqrt{\Delta t}.$$

So, the magnitude of the rate of change is:

$$\left|\frac{\Delta W}{\Delta t}\right| \sim \frac{\sqrt{\Delta t}}{\Delta t} = \frac{1}{\sqrt{\Delta t}} \to \infty \quad \text{as } \Delta t \to 0.$$

As $\Delta t \to 0$, the slope blows up. This is why $\frac{dW}{dt}$ does not exist.

Conceptually, this lack of a derivative makes perfect sense. If you could calculate a derivative, you would know the exact 'speed' and 'direction' of the path at that moment, allowing you to predict its next move. Because Brownian motion is entirely unpredictable, that short-term predictability mathematically cannot exist.

3.2 Infinite variation but finite quadratic variation

3.2.1 Infinite Variation

Total variation measures how much the path "wiggles" in total along the vertical axis:

$$TV = \lim_{N\to\infty} \sum_{i=1}^{N} |\Delta W_i| = \infty$$

This implies the Brownian path has infinite length.

Demonstration We sum the absolute increments. Since increments are normally distributed, $\Delta W \sim \mathcal{N}(0, \Delta t)$, the expected absolute step size is:

$$\mathbb{E}[|\Delta W|] = \sqrt{\frac{2}{\pi}}\sqrt{\Delta t}$$

Summing over the total time interval T:

$$\begin{aligned}\mathbb{E}[TV] &\approx \sum_{i=1}^{N} \mathbb{E}[|\Delta W_i|] \\ &= N \times \left(\sqrt{\frac{2}{\pi}}\sqrt{\Delta t}\right)\end{aligned}$$

Since the number of steps is $N = \frac{T}{\Delta t}$, we substitute to get:

$$= \frac{T}{\Delta t} \times \sqrt{\frac{2}{\pi}}\sqrt{\Delta t}$$
$$= T\sqrt{\frac{2}{\pi}} \times \frac{1}{\sqrt{\Delta t}}$$

Taking the limit as step size vanishes ($\Delta t \to 0$):

$$\lim_{\Delta t \to 0} \left(C \cdot \frac{T}{\sqrt{\Delta t}} \right) = \infty$$

This computation shows that the expected total variation diverges as we refine the partition. In fact, one can prove more rigorously that the total variation of a Brownian path on any time interval is infinite almost surely (with probability 1, but not literally for every imaginable outcome)

Intuition In standard calculus (smooth paths), a small displacement is proportional to time ($\Delta x \propto \Delta t$). However, for Brownian motion, displacement scales with the **square root** of time ($\Delta W \propto \sqrt{\Delta t}$). Because $\sqrt{\Delta t} \gg \Delta t$ for small Δt, the path wiggles infinitely more than a smooth curve.

3.2.2 Finite Quadratic Variation

While the absolute length explodes, something interesting happens if we square the increments. The Quadratic Variation (QV) is defined as:

$$QV = \lim_{\|\Pi\| \to 0} \sum (\Delta W_i)^2$$

As the partition gets finer, this sum converges to the total time elapsed:

$$\sum (\Delta W_i)^2 \to T$$

More formally, this is defined as:

$$[W, W]_t = \lim_{\|\Pi\| \to 0} \sum_{j=1}^{n} (W_{t_j} - W_{t_{j-1}})^2 = t \quad \iff \quad [W]_t = \lim_{n \to \infty} \sum_{j=1}^{n} (\Delta W_j)^2 = t$$

- Π represents a partition of the time interval $[0, T]$.
- $\|\Pi\| \to 0$ implies the mesh size (largest time step) goes to zero.

Intuitively, although the expected change between two instants is 0, the expected *squared* increment is exactly Δt. When we sum these small Δt pieces over the entire interval (of length T), we recover the total time T.

Demonstration We look at the expected value of the sum of squared increments. Recall that for $\Delta W \sim \mathcal{N}(0, \Delta t)$, the variance is $\mathbb{E}[(\Delta W)^2] = \Delta t$.

$$\mathbb{E}\left[\sum_{i=1}^{N}(\Delta W_i)^2\right] = \sum_{i=1}^{N}\mathbb{E}[(\Delta W_i)^2] = \sum_{i=1}^{N}\Delta t$$

Since the sum of time steps covers the total duration:

$$\sum_{i=1}^{N}\Delta t = T$$

Thus, the sum of squared increments averages out to exactly T. (A more complex variance analysis would show that the random deviation from T vanishes as $N \to \infty$).

Conclusion: The Itô Rule If we write the previous limit in differential notation, we obtain the fundamental rule of Stochastic Calculus (Itô calculus):

$$(dW_t)^2 = dt$$

This is the key difference from standard calculus, where second-order differentials are negligible.

4 Arithmetic Brownian Motion

Now that we know what Standard Brownian Motion is, we can understand what Arithmetic Brownian Motion is:

1. W_t **(Standard Brownian Motion):** This is the 'raw ingredient'. It is the pure, standardized "random noise" and it is defined by the rules (axioms) we saw earlier.

2. X_t **(Arithmetic Brownian Motion):** This is the 'final product'. It is a scaled and shifted version of the standard motion.

   $$dX_t = \mu dt + \sigma dW_t$$

   This formula says: "My custom process (X_t) is made of a predictable trend (μ) plus some Standard Brownian Motion (W_t) scaled by σ."

   If you integrate this differential equation, you get the explicit solution:

   $$X_t = X_0 + \mu t + \sigma W_t$$

   This shows that the final position (X_t) is simply the starting point (X_0), plus the predictable trend over time (μt), plus the scaled randomness (σW_t)

   Notice that a standard Brownian motion is a **martingale**; at any point in time, the expected future position is simply the current position. Mathematically:

$$\mathbb{E}[W_t \mid \mathcal{F}_s] = W_s, \quad \text{where } \mathcal{F}_s \text{ denotes a filtration (what you know has happened so far).}$$

Analogy: A Car

Think of X_t as a car driving on a road.

- μ **(The Drift):** The driver creates a steady path forward.
- σ **(The Suspension):** How loose or tight the suspension is.
- W_t **(The Standard Brownian Motion):** The **bumps in the road**. The road itself is the source of randomness.

The formula $dX_t = \mu dt + \sigma dW_t$ essentially says:

> "The movement of the car (dX) = The driver's steady driving (μdt) + The suspension's reaction (σ) to the bumps in the road (dW)."

The "bumps" (W_t) exist independently of the car. The car just experiences them.

The Explicit Formula

If you integrate the differential equation $dX_t = \mu dt + \sigma dW_t$, you get the explicit solution, which makes the relationship much clearer:

$$\underbrace{X_t}_{\text{Final Position}} = \underbrace{X_0}_{\text{Start}} + \underbrace{\mu t}_{\text{Trend}} + \underbrace{\sigma W_t}_{\text{Randomness}}$$

5 Brownian Motion in Finance

5.1 Geometric Brownian Motion

While Arithmetic Brownian Motion (ABM) is good at capturing the randomness of stocks, it cannot be used directly to model its movement, as it lacks of several basic properties of stock changes:

- **Negative values:** In ABM, the price change is independent of the current price level. If a stock is at \$10 and the volatility is \$5, the model thinks there is a significant chance the stock will hit $-\$5$ tomorrow.
- **Scale invariance:** If a stock moves from \$10 to \$11, that's a 10% move. If that same stock grows to $\$1,000$, ABM would still predict it moves by roughly \$1 per day. At a $\$1,000$ price point, a \$1 move is negligible (0.1%). In stocks, volatility is proportional to the price.
- **Compounding interest:** We need to model that as the price (S_t) gets bigger, the absolute dollar increase (dS_t) also gets bigger.

To overcome this limitations, the change (move) in stock prices is modeled as:

$$dS = \mu S \, \mathrm{d}t + \sigma S \, \mathrm{d}W$$

This is the formula of **Geometric Brownian Motion (GBM)**, which is very similar to the Arithmetic Brownian Motion, but because it multiplies the current value by a random factor, it stays strictly positive (it can get very close to zero, but it will never cross it).

Here:

- $\mathrm{d}W$ = pure randomness. W is the Wiener process (Standard Brownian Motion).
- σ = scale of uncertainty
- μ = drift (expected return)

This means that over a very small time interval $\mathrm{d}t$, the random change in the stock price has a deterministic part and a random part.

$\mu S \, \mathrm{d}t \rightarrow$ **expected growth**: Proportional to price, proportional to time, and directional.

$\sigma S \, \mathrm{d}W \rightarrow$ **uncertainty**: Random, with mean 0 and standard deviation $\sigma S \sqrt{\mathrm{d}t}$

5.2 Why GBM is Not Perfect for Modeling Stock Prices

- **Memoryless Assumption:** Geometric Brownian Motion (GBM) is a Markov process, meaning the next price move depends only on the current price. This only happens under the Weak Form of Market Efficiency (past prices don't predict future prices). However, real-world markets often exhibit **autocorrelation**. Some examples are:
  - **Momentum:** Assets that have been rising tend to rise in the short term.
  - **Mean Reversion:** Over longer horizons, prices frequently revert toward a historical average.
  - **Regimes:** Markets shift between high-volatility "panic" states and low-volatility "calm" states. GBM assumes a constant drift (μ) and volatility (σ), failing to account for these structural shifts.
- **Fat Tails (Kurtosis):** GBM assumes that returns are normally distributed. In reality, stock markets experience "black swan" events—extreme price swings—much more frequently than a normal distribution predicts. This results in **leptokurtic** distributions with "fat tails."
- **No jumps:** GBM assumes price paths are continuous and move in infinitesimal increments. In reality, markets often experience discrete "jumps" or gaps in price due to overnight news, earnings announcements, or sudden macroeconomic shocks that the model cannot capture.

6 Derivation of Brownian Motion as a Scaling Limit

Standard Brownian Motion W_t can be formally derived as the continuous-time limit of a symmetric discrete-time random walk. This transition is justified by **Donsker's Invariance Principle**. [2]

6.1 The Discrete Random Walk

Let $X_1, X_2, \ldots$ be a sequence of independent and identically distributed (i.i.d.) random variables representing steps in a walk, defined by:

$$P(X_i = 1) = \frac{1}{2}, \quad P(X_i = -1) = \frac{1}{2}$$

The moments for each step are $E[X_i] = 0$ and $Var(X_i) = E[X_i^2] = 1$. We define the position of the random walk after k steps as the sum of steps:

$$S_k = \sum_{i=1}^{k} X_i$$

6.2 Scaling and the Continuous Limit

To arrive at a continuous process W_t on the interval $[0, T]$, we divide the time into n small increments of size $\Delta t = \frac{T}{n}$. We define the scaled random walk $W_n(t)$ as:

$$W_n(t) = \frac{1}{\sqrt{n}} S_{\lfloor nt \rfloor} = \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} X_i$$

The scaling factor $\frac{1}{\sqrt{n}}$ is essential to ensure the variance does not vanish or explode as $n \to \infty$. [3]

6.3 Application of the Central Limit Theorem (CLT)

For a fixed time t, we observe the properties of $W_n(t)$ as n increases:

- **Mean:** $E[W_n(t)] = \frac{1}{\sqrt{n}} \sum_{i=1}^{nt} E[X_i] = 0$.
- **Variance:** Due to independence, $Var(W_n(t)) = \frac{1}{n} \sum_{i=1}^{nt} Var(X_i) = \frac{1}{n}(nt) = t$.

By the **Central Limit Theorem**, as $n \to \infty$, the distribution of $W_n(t)$ converges in distribution to a normal distribution:

$$W_n(t) \xrightarrow{d} \mathcal{N}(0, t)$$

[2] Donsker's principle is not the only approach; Paul Lévy's construction also provides a rigorous proof of Brownian motion's existence and path continuity (See Peter Mörters and Yuval Peres, 2010, Section 1.2)

[3] To see why this specific scaling is necessary, consider taking n steps in one unit of time. If the step size remains ± 1, the total variance for n steps is n, which explodes to infinity as $n \to \infty$. Alternatively, if we scale the steps by $\frac{1}{n}$, the single-step variance is $\frac{1}{n^2}$, making the total variance $n \times \frac{1}{n^2} = \frac{1}{n}$, which vanishes to zero. By scaling the steps by $\frac{1}{\sqrt{n}}$, the variance of a single step becomes $\frac{1}{n}$. Thus, the total variance for n steps remains perfectly stable at $n \times \frac{1}{n} = 1$, ensuring the path stays contained but appropriately volatile in the continuous limit.

6.4 Donsker's Invariance Principle

While the CLT handles the distribution at a specific point in time, **Donsker's Invariance Principle** extends this to the convergence of the *entire path*. It states that the sequence of random functions $W_n(t)$ converges in distribution to a Wiener process W_t:

$$W_n \Rightarrow W$$

- W_n represents the entire jagged, discrete random walk drawn on a graph (a random function).
- W represents the perfectly continuous, fractal-like path of true Brownian Motion (the Wiener process).

Intuition: It means that calculating the probability of the path doing anything over time, the answer you get from the discrete random walk (W_n) will perfectly match the answer you get from continuous Brownian Motion (W) as $n \to \infty$.

This principle is "invariant" because the limit remains a Brownian Motion regardless of the underlying distribution of X_i. Donsker's Principle guarantees that the continuous limit will be exactly the same Brownian Motion, regardless of what the underlying steps actually are, as long as they follow two rules:

1. They are independent and identically distributed (i.i.d.).
2. They have a mean of 0 and a finite variance.

Intuition: This means you could build your discrete random walk using, for example, standard coin flips ($+1$ or -1) or a computer generating random decimals between -1.5 and $+1.5$.

If you speed them up and shrink them down using the $\frac{1}{\sqrt{n}}$ scaling factor we discussed, all of them will morph into the exact same continuous Wiener process (W_t). Technically this is a functional central limit theorem: many different step distributions (with mean 0 and finite variance) all share the same Brownian motion limit under the $\frac{1}{\sqrt{n}}$ scaling.

References

[1] Andy Dahl. A Rigorous Introduction to Brownian Motion. https://www.math.uchicago.edu/~may/VIGRE/VIGRE2010/REUPapers/Dahl.pdf, 2010. Accessed: February 24, 2026.

[2] Dummy R. Brownian Motion for Dummies. https://www.youtube.com/watch?v=5A2iNvpAv1w, 2024. YouTube Video. Accessed: February 9, 2026.

[3] FinanceAndEconomics. Brownian Motion / Wiener Process Explained. https://www.youtube.com/watch?v=YZL2xdhh7qE, 2023. YouTube Video. Accessed: February 11, 2026.

[4] Davar Khoshnevisan. Donsker's Invariance Principle. https://www.math.utah.edu/~davar/ps-pdf-files/donsker.pdf, 2006. Lecture Notes. Accessed: February 16, 2026.

[5] Peter Mörters and Yuval Peres. *Brownian Motion*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2010.

[6] QuestDB. Geometric Brownian Motion for Asset Prices. https://questdb.com/glossary/geometric-brownian-motion-for-asset-prices/, 2024. Accessed: February 16, 2026.

[7] Jayati Walia. Brownian Motion in Finance. https://www.simtrade.fr/blog_simtrade/brownian-motion-finance/, 2021. Accessed: February 16, 2026.

[8] Wikipedia contributors. Donsker's theorem — Wikipedia, The Free Encyclopedia. https://en.wikipedia.org/w/index.php?title=Donsker%27s_theorem&oldid=1255562719, 2024. Accessed: February 16, 2026.