Covariance-First Simulation of Gaussian Processes

Karhunen–Loève Expansions for Brownian Motion, Brownian Bridges, and Financial Monte Carlo

Moritz J. Schmid, PhD

moritz.schmid@wu.ac.at

Abstract

This technical note develops a covariance-first view of Gaussian-process simulation for quantitative finance. The central idea is that Monte Carlo path generation is not merely about drawing random numbers, but about constructing paths with the correct variance, covariance, and temporal dependence structure. The Karhunen–Loève expansion provides a natural framework for this task by decomposing a covariance kernel into orthogonal directions of variation and assigning independent Gaussian coefficients to those directions. The note derives the expansions of Brownian motion and the Brownian bridge on $[0, 1]$, connects them to finite-dimensional covariance-matrix decompositions, and validates the resulting simulations numerically. The aim is to show how a single covariance-based principle links stochastic-process theory, numerical linear algebra, and practical simulation in financial Monte Carlo.

July 2026

1 Introduction

Simulation is one of the central tools of quantitative finance. Pricing, hedging, and risk management often require expectations of path functionals rather than functions of a single random variable. A payoff may depend on an asset's terminal value, a running maximum or minimum, the time spent above a barrier, realized variance, or the joint evolution of several correlated risk factors. A generic valuation problem can be written as

$$V_0 = e^{-rT}\mathbb{E}^{\mathbb{Q}}\left[\Phi\big(X_{[0,T]}\big)\right], \tag{1}$$

where V_0 is the time-zero value, r is the continuously compounded risk-free rate, T is the maturity, $\mathbb{Q}$ is the risk-neutral pricing measure, and Φ is a payoff or loss functional of the path $X_{[0,T]}$. Closed-form solutions are available only in special cases, which is why Monte Carlo methods remain central in financial engineering (Glasserman, 2004; Jäckel, 2002; Asmussen and Glynn, 2007).

Brownian motion is the canonical noise source in continuous-time finance. It drives the Black–Scholes model (Black and Scholes, 1973; Merton, 1973), appears in short-rate and forward-rate models (Hull and White, 1990; Heath et al., 1992), and underlies multi-asset diffusion and stochastic-volatility models. More recently, fractional Brownian motion and Gaussian Volterra processes have become important in rough-volatility modelling (Bayer et al., 2016; Bennedsen et al., 2017; Gatheral et al., 2018). Across these settings, the numerical task is not simply to generate random numbers. It is to generate random paths with the correct variance, covariance, and temporal dependence structure.

Brownian motion itself is easy to simulate incrementally. Standard Brownian motion satisfies

$$W_0 = 0, \qquad W_t - W_s \sim \mathcal{N}(0, t-s), \qquad 0 \le s < t, \tag{2}$$

with independent increments and almost surely continuous paths. These properties suggest the usual simulation procedure: generate independent Gaussian increments and accumulate them. Many Gaussian processes, however, do not admit such a convenient incremental construction. In those cases, the covariance function becomes the natural starting point.

For a mean-zero Gaussian process, the covariance function

$$K(s,t) = \operatorname{Cov}(X_s, X_t) \tag{3}$$

determines all finite-dimensional distributions. Constructing the process in distribution is therefore equivalent to constructing random variables with the correct covariance structure. This covariance-first view is especially useful in financial simulation, where the dependence structure of risk factors, yield curves, volatility drivers, or conditioned paths is often more important than any single marginal distribution (Rasmussen and Williams, 2006; Lifshits, 2012).

The Karhunen–Loève theorem provides a systematic covariance-based construction. It decomposes a covariance kernel into deterministic orthogonal functions and non-negative eigenvalues, allowing a Gaussian process to be represented as a weighted sum of deterministic modes with independent Gaussian coefficients (Loève, 1946; Loève, 1977; Ash and Gardner, 1975). The guiding principle is simple:

To simulate a Gaussian process, decompose its covariance structure and inject randomness along the resulting orthogonal directions.

This is the infinite-dimensional analogue of simulating a multivariate normal random vector. If $Z \sim \mathcal{N}(0, I_n)$ and Γ is a positive semidefinite covariance matrix, then

$$X = \Gamma^{1/2} Z, \qquad \Gamma^{1/2}\left(\Gamma^{1/2}\right)^{\top} = \Gamma, \tag{4}$$

has covariance matrix Γ. The Karhunen–Loève expansion applies the same logic to a covariance operator: eigenvectors become eigenfunctions, and a finite covariance matrix becomes a covariance kernel.

This note develops that idea using Brownian motion and the Brownian bridge as canonical examples. Brownian motion is the fundamental building block of diffusion models. The Brownian bridge is equally important in simulation because it describes Brownian motion conditioned on an endpoint and appears naturally in barrier option pricing, variance reduction, quasi-Monte Carlo path construction, and conditional path generation (Glasserman, 2004; Jäckel, 2002). The note derives the continuous-time Karhunen–Loève expansions of both processes on $[0, 1]$, connects them to finite-dimensional covariance-matrix decompositions, and validates the resulting simulations numerically.

2 The covariance-first view of simulation

The covariance-first construction is most transparent in finite dimensions. Let $X \in \mathbb{R}^n$ be a mean-zero Gaussian random vector with covariance matrix Γ, whose entries are

$$\Gamma_{ij} = \mathrm{Cov}(X_i, X_j), \qquad i, j = 1, \ldots, n. \tag{5}$$

The diagonal entries of Γ are variances; the off-diagonal entries are covariances. In finance, such a vector may represent asset returns, shocks to different points of a yield curve, innovations to several risk factors, or simulated process values on a time grid.

The objective is to transform independent Gaussian noise into correlated Gaussian noise with the desired covariance structure. Let

$$Z = (Z_1, \ldots, Z_n)^{\top}, \qquad Z \sim \mathcal{N}(0, I_n), \tag{6}$$

where I_n is the $n \times n$ identity matrix. If Γ is positive semidefinite, it admits an eigenvalue decomposition

$$\Gamma = Q \Lambda Q^{\top}. \tag{7}$$

Here, Q is an orthogonal matrix of eigenvectors and Λ is a diagonal matrix of non-negative eigenvalues. The eigenvectors identify orthogonal directions of variation, while the eigenvalues measure variance along those directions. This is the same geometric principle underlying principal component analysis.

A covariance square root of Γ is any matrix A satisfying $AA^\top = \Gamma$. The eigenvalue decomposition gives one such square root,

$$A = Q\Lambda^{1/2}, \tag{8}$$

where $\Lambda^{1/2}$ contains the square roots of the eigenvalues. Defining

$$X = Q\Lambda^{1/2}Z \tag{9}$$

therefore produces a Gaussian vector with covariance matrix Γ, since

$$\begin{aligned}
\mathrm{Cov}(X) &= \mathrm{Cov}\Big(Q\Lambda^{1/2}Z\Big) && (10)\\
&= Q\Lambda^{1/2}\,\mathrm{Cov}(Z)\Lambda^{1/2}Q^\top && (11)\\
&= Q\Lambda^{1/2}I_n\Lambda^{1/2}Q^\top && (12)\\
&= Q\Lambda Q^\top && (13)\\
&= \Gamma. && (14)
\end{aligned}$$

This covariance-matching principle appears whenever one simulates correlated asset returns, correlated Brownian increments, multi-factor interest-rate shocks, or joint market scenarios. It is also the basis of PCA-based risk-factor models, where eigenvectors represent orthogonal factors and eigenvalues measure explained variance (Glasserman, 2004; Jäckel, 2002; Golub and Van Loan, 2013; Trefethen and Bau, 1997).

The Karhunen–Loève theorem is the infinite-dimensional version of the same construction. Instead of a covariance matrix Γ, one works with a covariance kernel $K(s,t)$. Instead of diagonalizing a matrix, one diagonalizes an integral operator. Instead of eigenvectors, one obtains eigenfunctions. Instead of a finite weighted sum of Gaussian variables, one obtains an infinite series representation of a stochastic process.

This perspective is useful because continuous-time models are ultimately implemented on finite grids. The covariance matrix used in numerical simulation can be viewed as a discrete approximation to the covariance operator of the underlying stochastic process.

3 The Karhunen–Loève expansion

Let $(X_t)_{t\in[0,1]}$ be a mean-zero, square-integrable stochastic process with covariance kernel

$$K(s,t) = \mathrm{Cov}(X_s, X_t) = \mathbb{E}[X_sX_t]. \tag{15}$$

The Karhunen–Loève expansion is obtained from the integral eigenvalue problem

$$\int_0^1 K(s,t)e_k(s)\ \mathrm{d}s = \lambda_k e_k(t). \tag{16}$$

The eigenfunctions e_k play the role of eigenvectors in finite-dimensional linear algebra, while the eigenvalues λ_k measure the variance carried by each orthogonal direction.

Under suitable regularity conditions, a Gaussian process admits the representation

$$X_t = \sum_{k=1}^{\infty} \sqrt{\lambda_k}\, \xi_k e_k(t), \qquad \xi_k \sim \mathcal{N}(0,1), \tag{17}$$

where the coefficients ξ_k are independent standard normal random variables. The deterministic functions $e_k(t)$ determine the shape of the components over time; the factors $\sqrt{\lambda_k}$ scale their random amplitudes.

For simulation, the infinite expansion is truncated:

$$X_t^{(m)} = \sum_{k=1}^{m} \sqrt{\lambda_k}\, \xi_k e_k(t). \tag{18}$$

The truncated process has covariance

$$\mathrm{Cov}\left(X_s^{(m)}, X_t^{(m)}\right) = \sum_{k=1}^{m} \lambda_k e_k(s) e_k(t), \tag{19}$$

which approaches the full covariance kernel

$$K(s,t) = \sum_{k=1}^{\infty} \lambda_k e_k(s) e_k(t) \tag{20}$$

as more components are retained. Thus, truncation is not an arbitrary numerical shortcut; it is a controlled covariance approximation.

In quantitative finance, this matters whenever the dependence structure is more naturally specified through a covariance kernel or covariance matrix than through independent increments. Brownian bridges, Gaussian term-structure models, correlated risk-factor simulations, and Gaussian or Volterra-type volatility drivers all fit this logic. Low-dimensional Karhunen–Loève approximations can also improve Monte Carlo and quasi-Monte Carlo efficiency by concentrating the dominant variance in the leading random coordinates (Glasserman, 2004; Jäckel, 2002). The classical foundations of the expansion are due to Loève (1946) and related developments in second-order stochastic processes; modern Gaussian-process treatments include Rasmussen and Williams (2006) and Lifshits (2012).

4 Brownian motion

Brownian motion is the fundamental continuous-time noise process in mathematical finance. It drives the Black–Scholes model, many short-rate and forward-rate models, multi-asset diffusion models, and stochastic-volatility models (Black and Scholes, 1973; Merton, 1973; Shreve, 2004). Let $(W_t)_{t\in[0,1]}$ denote standard Brownian motion on the unit interval. It satisfies

$$W_0 = 0, \qquad W_t - W_s \sim \mathcal{N}(0, t-s), \qquad 0 \le s < t \le 1, \tag{21}$$

with independent increments and almost surely continuous paths. Its covariance function is

$$K(s,t) = \operatorname{Cov}(W_s, W_t) = \min(s,t). \tag{22}$$

The expression $\min(s,t)$ reflects the shared accumulated randomness up to the earlier of the two times.

The Karhunen–Loève eigenvalue problem for Brownian motion is

$$\int_0^1 \min(s,t) e_k(s)\ \mathrm{d}s = \lambda_k e_k(t). \tag{23}$$

Solving this problem yields

$$e_k(t) = \sqrt{2}\sin\left(\left(k-\frac{1}{2}\right)\pi t\right), \qquad \lambda_k = \frac{1}{\left(k-\frac{1}{2}\right)^2\pi^2}, \qquad k = 1,2,\ldots. \tag{24}$$

The half-integer frequencies reflect the boundary structure of Brownian motion on $[0,1]$: the process is fixed at the origin but free at the endpoint.

Substitution into the Karhunen–Loève representation gives

$$\begin{aligned} W_t &= \sum_{k=1}^{\infty} \sqrt{\lambda_k}\,\xi_k e_k(t) \\ &= \sqrt{2}\sum_{k=1}^{\infty} \xi_k \frac{\sin\left(\left(k-\frac{1}{2}\right)\pi t\right)}{\left(k-\frac{1}{2}\right)\pi}, \qquad \xi_k \sim \mathcal{N}(0,1). \end{aligned} \tag{25}$$

For simulation, the series is truncated after m terms:

$$W_t^{(m)} = \sqrt{2}\sum_{k=1}^{m} \xi_k \frac{\sin\left(\left(k-\frac{1}{2}\right)\pi t\right)}{\left(k-\frac{1}{2}\right)\pi}. \tag{26}$$

Small values of m produce overly smooth paths because high-frequency components are omitted; increasing m restores progressively finer fluctuations.

For financial simulation, this representation is useful because it constructs an entire path from global modes of variation rather than from local increments. That perspective is particularly relevant for variance reduction, quasi-Monte Carlo path generation, and low-dimensional approximations of path-dependent payoffs, where the ordering and concentration of variance across random coordinates can affect simulation efficiency (Glasserman, 2004; Jäckel, 2002).

5 Brownian bridge

The Brownian bridge is Brownian motion conditioned to end at a fixed value. On $[0,1]$, the standard Brownian bridge starts at zero and returns to zero at time one. It can be written as

$$B_t = W_t - tW_1, \qquad 0 \le t \le 1. \tag{27}$$

This construction immediately gives

$$B_0 = 0, \qquad B_1 = 0. \tag{28}$$

Brownian bridges are important in financial simulation because they describe conditional path evolution given endpoint information. This is central in barrier option pricing, where a path may cross a barrier between two simulated dates even if both endpoints remain on the same side. Brownian-bridge constructions are also widely used in variance reduction and quasi-Monte Carlo methods (Glasserman, 2004; Jäckel, 2002).

The covariance function of the Brownian bridge is

$$K(s,t) = \mathrm{Cov}(B_s, B_t) = \min(s,t) - st. \tag{29}$$

The term $\min(s,t)$ is the Brownian-motion covariance. The correction term st removes the covariance contribution associated with the random terminal value W_1, thereby enforcing the endpoint condition $B_1 = 0$.

The corresponding eigenvalue problem is

$$\int_0^1 \big(\min(s,t) - st\big) e_k(s)\ \mathrm{d}s = \lambda_k e_k(t). \tag{30}$$

Its solution is

$$e_k(t) = \sqrt{2}\sin(k\pi t), \qquad \lambda_k = \frac{1}{k^2\pi^2}, \qquad k = 1, 2, \ldots. \tag{31}$$

The integer sine frequencies reflect the fact that the eigenfunctions must vanish at both endpoints.

The Karhunen–Loève expansion is therefore

$$B_t = \sqrt{2}\sum_{k=1}^{\infty} \xi_k \frac{\sin(k\pi t)}{k\pi}, \qquad \xi_k \sim \mathcal{N}(0,1). \tag{32}$$

The truncated simulation is

$$B_t^{(m)} = \sqrt{2}\sum_{k=1}^{m} \xi_k \frac{\sin(k\pi t)}{k\pi}. \tag{33}$$

Because each sine function vanishes at $t = 0$ and $t = 1$, the truncated process also satisfies

$$B_0^{(m)} = 0, \qquad B_1^{(m)} = 0. \tag{34}$$

The contrast with Brownian motion is instructive. Brownian motion is fixed at the origin but free at the endpoint, leading to half-integer frequencies. The Brownian bridge is pinned at both endpoints, leading to integer frequencies. In applications, this difference is not merely cosmetic: conditioning on endpoint information changes the covariance structure between observation dates, which can materially affect crossing probabilities and path-dependent payoffs.

6 Finite-dimensional covariance decompositions

In practice, stochastic processes are simulated on finite time grids. Even when a model is formulated in continuous time, a Monte Carlo implementation evaluates it at finitely many dates, such as option monitoring dates, cash-flow dates, rebalancing dates, or discretization points for a path-dependent payoff.

Let

$$0 = t_1 < t_2 < \cdots < t_n = 1 \tag{35}$$

be a grid on the unit interval. Instead of simulating the full continuous-time process $(X_t)_{t\in[0,1]}$, one simulates the vector

$$X = \left(X_{t_1}, X_{t_2}, \ldots, X_{t_n}\right)^\top. \tag{36}$$

For a mean-zero Gaussian process, this vector is fully determined by its covariance matrix

$$\Gamma_{ij} = K(t_i, t_j), \qquad i, j = 1, \ldots, n. \tag{37}$$

For Brownian motion,

$$\Gamma_{ij} = \min(t_i, t_j), \tag{38}$$

whereas for the Brownian bridge on $[0, 1]$,

$$\Gamma_{ij} = \min(t_i, t_j) - t_i t_j. \tag{39}$$

The continuous covariance kernel is thus converted into a finite covariance matrix that can be decomposed numerically.

Let $Z \sim \mathcal{N}(0, I_n)$ be a vector of independent standard normal random variables. If

$$\Gamma = Q\Lambda Q^\top, \tag{40}$$

then

$$X = Q\Lambda^{1/2} Z \tag{41}$$

has covariance matrix Γ. Alternatively, if $\Gamma = LL^\top$ is a Cholesky decomposition, then $X = LZ$ has the same covariance. Cholesky methods are often computationally efficient; eigenvalue methods are more closely aligned with the Karhunen–Loève perspective because they explicitly decompose covariance into orthogonal variance components.

This finite-dimensional view is central to quantitative finance. Correlated multi-asset simulations use covariance or correlation matrices across assets. Term-structure simulations use covariance across maturities or factors. Path-dependent derivative simulations use covariance across monitoring dates. In each case, independent Gaussian random variables are transformed into correlated shocks with the desired dependence structure.

The link to the Karhunen–Loève expansion is direct. The continuous-time expansion decomposes a covariance operator into eigenfunctions and eigenvalues; the finite-dimensional simulation decomposes a covariance matrix into eigenvectors and eigenvalues. The mathematical object changes, but

the principle remains the same: decompose covariance, simulate independent Gaussian coefficients, and reconstruct the vector or process with the desired dependence structure (Glasserman, 2004; Jäckel, 2002; Golub and Van Loan, 2013; Trefethen and Bau, 1997).

7 Simulation and numerical validation

The preceding sections describe the Karhunen–Loève expansion as a covariance-based construction of Gaussian processes. This section illustrates the construction numerically. The simulations are implemented in R by truncating the infinite series and evaluating the resulting process on a finite time grid.

For Brownian motion, the truncated simulation uses

$$W_t^{(m)} = \sqrt{2} \sum_{k=1}^{m} \xi_k \frac{\sin\left(\left(k - \frac{1}{2}\right) \pi t\right)}{\left(k - \frac{1}{2}\right) \pi}. \tag{42}$$

The core simulation step is straightforward. The following R function evaluates the truncated Brownian-motion expansion on a regular grid. The Brownian-bridge implementation is obtained by replacing the half-integer sine basis with the integer sine basis.

```r
simulate_bm_kl <- function(steps = 600, n_terms = 350,
                           paths = 1000, seed = NULL) {
  if (!is.null(seed)) set.seed(seed)
  t <- seq(0, 1, length.out = steps)
  k <- seq_len(n_terms)
  basis <- outer(
    t, k,
    function(t, k) sin((k - 0.5) * pi * t) / ((k - 0.5) * pi)
  )
  z <- matrix(rnorm(n_terms * paths), n_terms, paths)
  t(sqrt(2) * basis %*% z)
}
```

Listing 1: Core R implementation of the Brownian-motion KL simulation.

Figure 1 shows simulated Brownian motion paths generated from this representation. The paths start at zero and diffuse outward as variance accumulates over time.

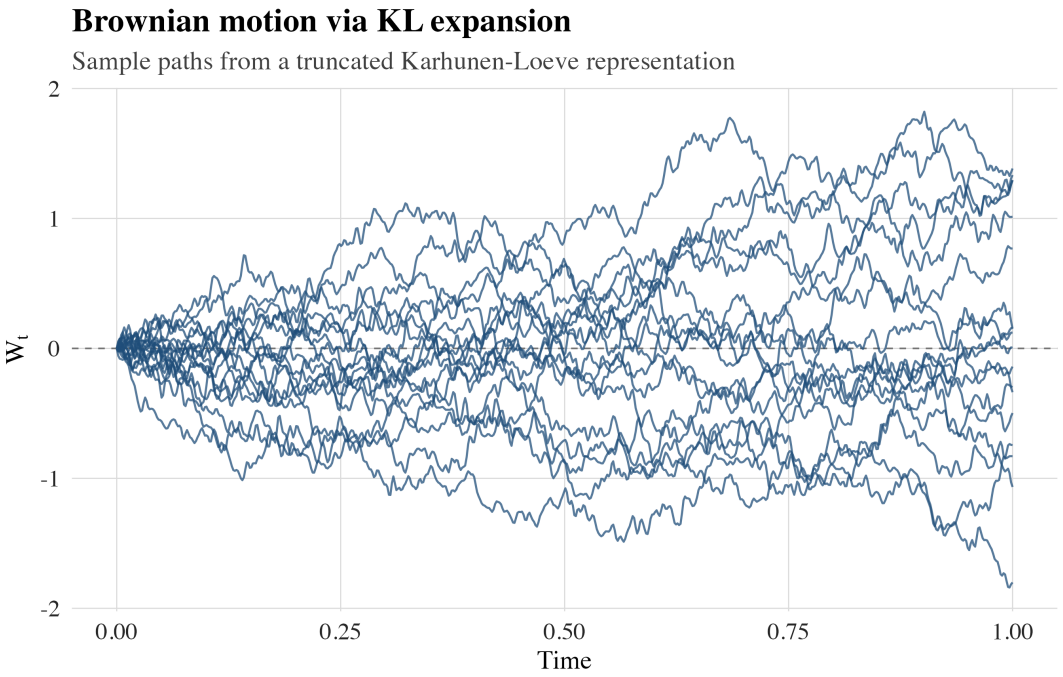

Figure 1: Sample paths of Brownian motion simulated from a truncated Karhunen–Loève expansion.

For the Brownian bridge, the truncated simulation uses

$$B_t^{(m)} = \sqrt{2} \sum_{k=1}^{m} \xi_k \frac{\sin(k\pi t)}{k\pi}. \tag{43}$$

Because each basis function vanishes at both endpoints, all simulated paths satisfy the bridge constraints at $t = 0$ and $t = 1$, as shown in Figure 2.

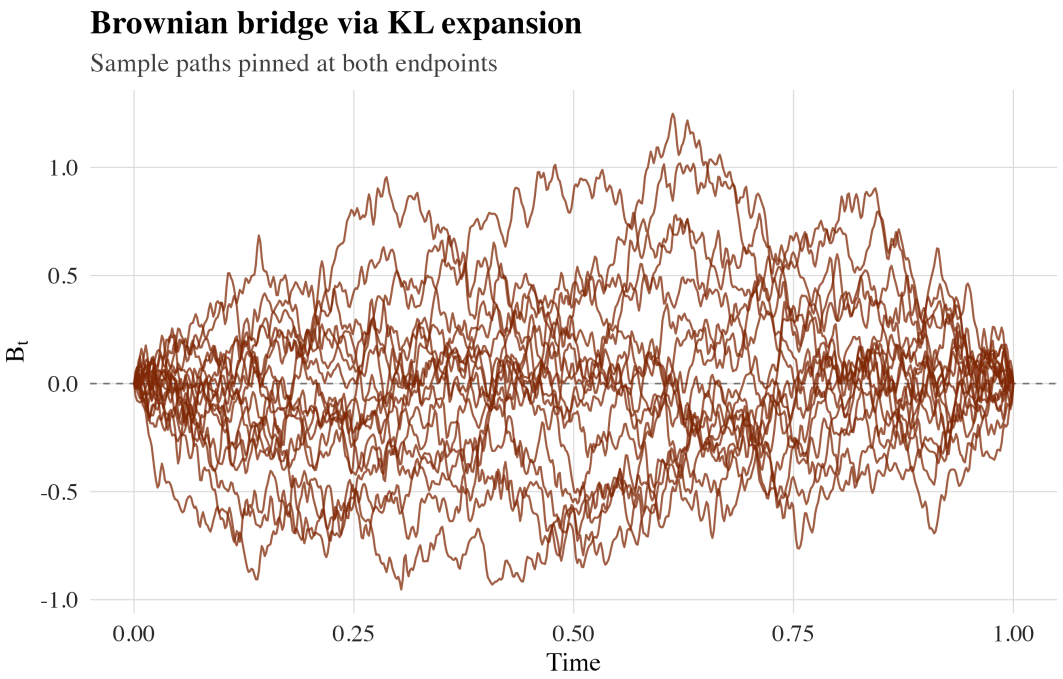

Figure 2: Sample paths of the Brownian bridge simulated from a truncated Karhunen–Loève expansion.

The eigenvalues determine how much variance is carried by each component. For Brownian motion and the Brownian bridge, respectively,

$$\lambda_k^W = \frac{1}{\left(k - \frac{1}{2}\right)^2 \pi^2}, \qquad \lambda_k^B = \frac{1}{k^2 \pi^2}. \tag{44}$$

Both sequences decay at rate k^{-2}. Figure 3 plots this decay on a logarithmic scale. The decrease is fast enough for the leading components to capture the dominant large-scale variation, but it is still relatively slow compared with the eigenvalue decay generated by smoother covariance kernels. This matters in practice: Brownian paths are rough, and higher-order components remain relevant when the payoff depends on fine path features such as barrier crossings, extrema, or realized variation.

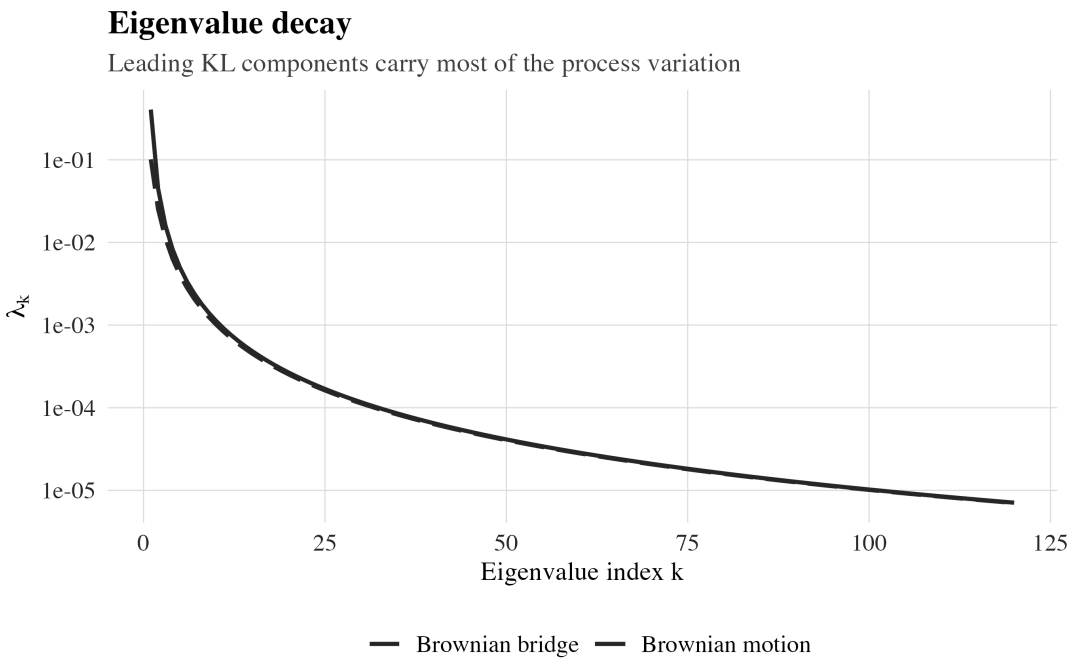

Figure 3: Eigenvalue decay for the Karhunen–Loève expansions of Brownian motion and the Brownian bridge.

A useful diagnostic is to compare the empirical covariance matrix of the simulated paths with the theoretical covariance matrix. For N simulated paths, let

$$X^{(i)} = \left(X_{t_1}^{(i)}, \ldots, X_{t_n}^{(i)}\right)^\top \tag{45}$$

denote the i-th path evaluated on the grid. The empirical covariance matrix computed from these vectors should approximate the theoretical matrix implied by the process.

For Brownian motion, the theoretical covariance matrix has entries

$$\Gamma_{ij} = \min(t_i, t_j). \tag{46}$$

Figure 4 plots the empirical covariance matrix minus the theoretical covariance matrix. The errors are small relative to the covariance scale and reflect ordinary Monte Carlo sampling error.

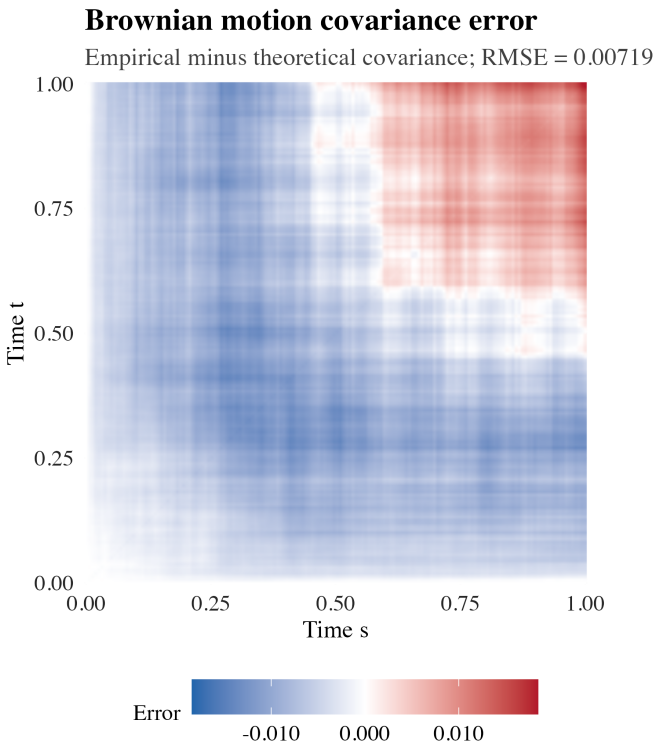

Figure 4: Empirical covariance error for Brownian motion. The figure shows the empirical covariance matrix minus the theoretical covariance matrix.

For the Brownian bridge, the theoretical covariance matrix has entries

$$\Gamma_{ij} = \min(t_i, t_j) - t_i t_j. \tag{47}$$

Figure 5 shows the corresponding covariance error. The error is again small and centered around zero, confirming that the simulated paths reproduce the target covariance structure.

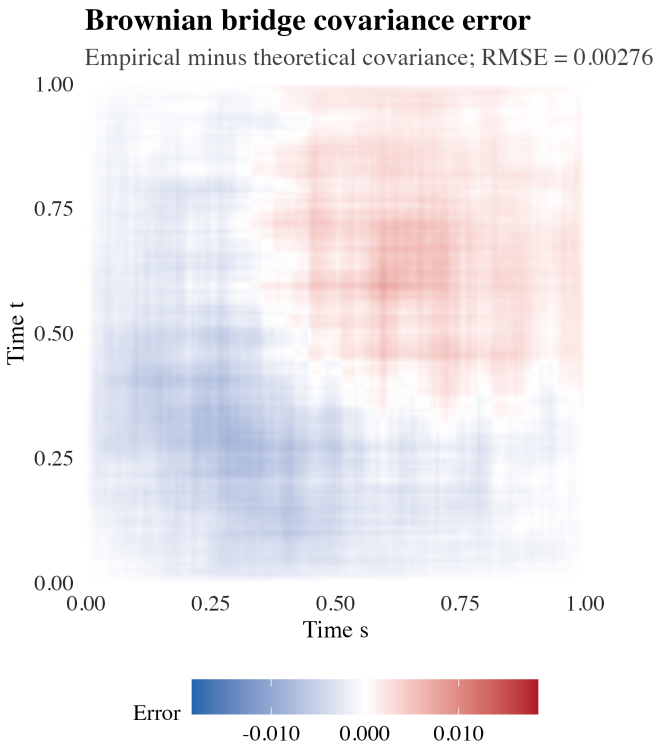

Figure 5: Empirical covariance error for the Brownian bridge. The figure shows the empirical covariance matrix minus the theoretical covariance matrix.

These diagnostics confirm the covariance-matching logic. The truncated Karhunen–Loève expansion does not merely generate visually plausible paths; it generates paths whose empirical covariance structure approximates the theoretical covariance kernel. In quantitative finance, this validation step is important because pricing and risk estimates can be sensitive to dependence, especially for path-dependent payoffs, barrier events, volatility-linked claims, and multi-factor models.

8 Practical considerations

The Karhunen–Loève expansion gives a clean covariance-based representation of Gaussian processes, but practical simulation requires choices about truncation, time discretization, sample size, and numerical stability.

The truncation level m controls how many covariance modes are retained:

$$X_t^{(m)} = \sum_{k=1}^{m} \sqrt{\lambda_k}\, \xi_k e_k(t). \tag{48}$$

Increasing m improves the covariance approximation but increases computational cost. A practical way to choose m is to retain enough components to explain a target fraction of total variance. For

example, one may choose the smallest m such that

$$\frac{\sum_{k=1}^{m} \lambda_k}{\sum_{k=1}^{M} \lambda_k} \geq 0.99, \tag{49}$$

where M is a sufficiently large reference truncation level. This explained-variance criterion is not a substitute for payoff-level validation, but it gives a useful first diagnostic for the numerical dimension of the simulation problem.

The eigenvalue decay provides further guidance. For Brownian motion and the Brownian bridge, the eigenvalues decay at rate k^{-2}, so low-order components capture the dominant large-scale movements. At the same time, the decay is not extremely fast, reflecting the roughness of Brownian paths. For smooth or moderately path-dependent quantities, relatively few components may be sufficient; for payoffs that depend on fine path features, such as barrier crossings or realized variation, more components may be required.

This ordering of components is also useful for quasi-Monte Carlo. When the leading Karhunen–Loève coordinates explain most of the variation, low-discrepancy points can be assigned to the most important directions first, with randomized shifts or scrambling used to retain error estimation. This is one reason Brownian-bridge and principal-component constructions are common in financial Monte Carlo implementations.

The time grid is equally important. A continuous-time model must be evaluated at finitely many time points. If the grid is too coarse, relevant path behavior between grid points may be missed. This matters for barrier options, lookback options, discretely monitored payoffs, and risk measures based on intraperiod extremes. Finer grids improve resolution but increase computational cost.

Monte Carlo sampling error is another constraint. If N denotes the number of simulated paths, standard Monte Carlo error typically decreases at rate

$$N^{-1/2}. \tag{50}$$

Increasing the number of paths by a factor of four therefore reduces the standard error by roughly one half. This slow rate motivates variance-reduction methods. Brownian-bridge constructions and low-dimensional covariance decompositions are useful precisely because they can improve the efficiency of Monte Carlo or quasi-Monte Carlo simulation (Glasserman, 2004; Jäckel, 2002).

Numerical stability also matters. In finite-dimensional simulation, covariance matrices may be nearly singular, especially for finely spaced grids or constrained processes such as the Brownian bridge. Eigenvalue decompositions can produce tiny negative eigenvalues because of floating-point error. These values do not represent meaningful negative variances and are usually set to zero before taking square roots.

In applied finance, these details are not cosmetic. A covariance approximation that is adequate for a vanilla option may be inadequate for a barrier option, volatility-linked payoff, or multi-factor risk model. The advantage of the Karhunen–Loève perspective is that it makes the approximation explicit: it shows which components of variation are retained, which are discarded, and how the simulated process inherits its dependence structure from the underlying covariance kernel.

9 Conclusion

The Karhunen–Loève expansion provides a natural way to simulate Gaussian processes from their covariance structure. In finite dimensions, this is the familiar construction of a multivariate normal vector from a covariance matrix. In continuous time, the same logic becomes a spectral decomposition of a covariance operator.

Brownian motion and the Brownian bridge illustrate the construction clearly. Brownian motion has covariance kernel

$$K(s,t) = \min(s,t), \tag{51}$$

which leads to

$$W_t = \sqrt{2}\sum_{k=1}^{\infty} \xi_k \frac{\sin\left(\left(k-\frac{1}{2}\right)\pi t\right)}{\left(k-\frac{1}{2}\right)\pi}. \tag{52}$$

The Brownian bridge has covariance kernel

$$K(s,t) = \min(s,t) - st, \tag{53}$$

which leads to

$$B_t = \sqrt{2}\sum_{k=1}^{\infty} \xi_k \frac{\sin(k\pi t)}{k\pi}. \tag{54}$$

The difference between half-integer and integer sine frequencies reflects the difference in boundary conditions: Brownian motion is fixed only at the origin, whereas the Brownian bridge is pinned at both endpoints.

The numerical simulations confirm the covariance-matching logic. Truncated Karhunen–Loève expansions generate visually plausible paths, but more importantly, their empirical covariance matrices approximate the theoretical covariance structures of the target processes. The eigenvalue decay also explains why the first components capture much of the large-scale variation, while higher-order terms add finer path fluctuations.

For quantitative finance, the central message is that simulation is often a problem of dependence construction. Pricing, hedging, and risk management require paths with the correct variance, covariance, and temporal dependence structure. This is especially important for path-dependent derivatives, barrier events, volatility-linked payoffs, interest-rate models, multi-factor risk systems, and rough-volatility models. The Karhunen–Loève expansion offers a transparent framework for such simulations: decompose covariance, simulate independent Gaussian coefficients, and reconstruct paths with the desired dependence structure.

References

Ash, R. B. and M. F. Gardner (1975). *Topics in Stochastic Processes*. Academic Press.

Asmussen, S. and P. W. Glynn (2007). *Stochastic Simulation: Algorithms and Analysis*. Springer.

Bayer, C., P. Friz, and J. Gatheral (2016). "Pricing under rough volatility". In: *Quantitative Finance* 16.6, pp. 887–904.

Bennedsen, M., A. Lunde, and M. S. Pakkanen (2017). "Hybrid scheme for Brownian semistationary processes". In: *Finance and Stochastics* 21.4, pp. 931–965.

Black, F. and M. Scholes (1973). "The Pricing of Options and Corporate Liabilities". In: *Journal of Political Economy* 81.3, pp. 637–654.

Gatheral, J., T. Jaisson, and M. Rosenbaum (2018). "Volatility is rough". In: *Quantitative Finance* 18.6, pp. 933–949.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer.

Golub, G. H. and C. F. Van Loan (2013). *Matrix Computations*. 4th ed. Johns Hopkins University Press.

Heath, D., R. Jarrow, and A. Morton (1992). "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation". In: *Econometrica* 60.1, pp. 77–105.

Hull, J. and A. White (1990). "Pricing Interest-Rate-Derivative Securities". In: *Review of Financial Studies* 3.4, pp. 573–592.

Jäckel, P. (2002). *Monte Carlo Methods in Finance*. John Wiley & Sons.

Lifshits, M. (2012). *Lectures on Gaussian Processes*. Springer.

Loève, M. (1946). "Fonctions aléatoires de second ordre". In: *Comptes Rendus de l'Académie des Sciences* 222, pp. 401–403.

— (1977). *Probability Theory I*. 4th ed. Springer.

Merton, R. C. (1973). "Theory of Rational Option Pricing". In: *Bell Journal of Economics and Management Science* 4.1, pp. 141–183.

Rasmussen, C. E. and C. K. I. Williams (2006). *Gaussian Processes for Machine Learning*. MIT Press.

Shreve, S. E. (2004). *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer.

Trefethen, L. N. and D. Bau (1997). *Numerical Linear Algebra*. SIAM.