6.1 Interest Rates Derivatives

In this section, we consider the pricing of *collateralized mortgage obligations* and the valuation of *zero coupon bonds*. Both applications lead to high-dimensional integrals where the large number of dimensions arises from small time steps in the time discretization. To reduce the effective dimension of these integrals we compare the different methods from Section 5.1 to generate sample paths of the stochastic processes which model the interest rate movements. Since the resulting integrands are smooth functions, no smoothing is required and sparse grid methods are particularly efficient as we will demonstrate by numerical experiments.

6.1.1 Zero Coupon Bonds

We here consider the problem to price zero coupon bonds by simulating the short-term interest rate $r(t)$ using the Vasicek model.[3] The same problem is also considered in [115, 110, 156] to analyse the behaviour of QMC methods.

6.1.1.1 Modeling

In the Vasicek model the movement of the short-term interest rate is given by

$$dr(t) = \kappa(\theta - r(t))dt + \sigma dW(t), \tag{6.1}$$

where $W(t)$ is the standard Brownian motion, $\theta > 0$ denotes the mean reversion level, $\kappa > 0$ denotes the reversion rate and $\sigma \geq 0$ denotes the volatility of the short rate dynamic.

For the solution of the above stochastic differential equation until time $t = T$, we use an Euler-Maruyama discretization with step size $\Delta t := T/d$ on the grid $t_k := \Delta k$, $k = 1, \ldots, d$, which yields the time-discrete version

$$r_k = r_{k-1} + \kappa(\theta - r_{k-1})\Delta t + \sigma z_k. \tag{6.2}$$

Here $z_k := W(t_k) - W(t_{k-1})$ is normally distributed with mean zero and variance Δt and denotes the increment of the Brownian motion in the k-th time interval.

Based on the short-term interest rates (6.2) the price $P(0,T)$ at time $t = 0$ of a zero coupon bond with maturity date T is given by

$$P(0,T) = \mathbb{E}\left[\exp\left\{-\Delta t \sum_{k=0}^{d} r_k\right\}\right]. \tag{6.3}$$

[3] For more information on the Vasicek model and other short rate models we refer to [12].

It can be written as a d-dimensional integral on $\mathbf{R}^d$ with Gaussian density since the expected value is via (6.2) taken over d many normally distributed random variables z_k. The value $P(0,T)$ can also be derived in closed-form, which we use to validate the numerical methods and results.

Lemma 6.1 (Closed-form solution). *The price of the bond (6.3) at time $t = 0$ with the short-term interest rate (6.2) is given by*

$$P(0,T) = \exp\left\{-\frac{(\gamma + \beta_d\, r(0))\, T}{d}\right\} \tag{6.4}$$

where

$$\beta_k := \sum_{j=1}^{k} (1-\kappa\Delta t)^{j-1},\ k = 1,\dots,d, \quad \textit{and}\ \gamma := \sum_{k=1}^{d-1} \beta_k \kappa\theta\Delta t - (\beta_k \sigma\Delta t)^2/2.$$

Proof. The proof uses that r_k is normally distributed. The derivation can be found in [151].

In our following numerical experiments, we always use the parameters

$$\kappa = 0.1817303, \theta = 0.0825398957, \sigma = 0.0125901, r(0) = 0.021673 \text{ and } T = 5$$

which are also used in [115]. We consider the high-dimensional case $d = 512$.

6.1.1.2 Effective Dimensions

We first compute the effective dimensions of the expected value (6.3) to investigate the impact of the different path constructions from Section 5.1. Here, we always compare random walk (RW), Brownian bridge (BB), principal component (PCA) and the LT-construction (LT).[4] The results are shown in Table 6.1.

We first observe that the Vasicek bond pricing problem (6.3) is of very low superposition dimension $d_s \leq 2$ in the anchored case and that d_s is almost independent of the path construction.

One can furthermore see that the truncation dimensions d_t in the classical case almost coincide with the truncation dimensions d_t in the anchored case. For instance, for $\alpha = 0.99$ we obtain $d_t = 420, 7, 1, 1$ using RW, BB, PCA and LT, respectively, for the anchored case and $d_t = 419, 7, 1, 1$ for the classical one. We conjecture that a more precise numerical computation would yield even exact equal results and we believe that this equality holds for a wider class of functions. Note however that this does not hold in general. We will give a counterexample in Section 6.1.2.

[4] Here and in the following we apply the LT-construction only to the first column of the matrix $\mathbf{A}$. The remaining columns are arbitrary but orthonormal. This way the LT-construction is significantly less expansive than the PCA-construction. The optimization of additional columns increases the computing times but led to no or only minor improvements in our numerical experiments.

Observe in Table 6.1 that the path construction has a significant impact on the truncation dimensions. For α close to one the dimensions d_t are almost as large as the nominal dimension $d = 512$ if we employ the RW approach. The dimensions d_t are significantly smaller if BB, PCA or LT is used instead. The LT-construction even obtains the optimal result $d_t = 1$ for this problem. While it is not surprising that such an optimal transformation exists,[5] it is nevertheless interesting that it is correctly identified by the LT-construction, which takes only the gradient of the integrand at a certain anchor point into account.

Table 6.1 Effective dimensions of the Vasicek bond pricing problem (6.3). The nominal dimension is $d = 512$.

	anchored-ANOVA superposition dim.				anchored-ANOVA truncation dim.				classical ANOVA truncation dim.			
$1-\alpha$	RW	BB	PCA	LT	RW	BB	PCA	LT	RW	BB	PCA	LT
1e-1	1	1	1	1	302	2	1	1	305	2	1	1
1e-2	1	1	1	1	420	7	1	1	419	7	1	1
1e-3	1	1	1	1	471	16	2	1	471	16	2	1
1e-4	2	2	1	1	494	59	5	1	494	52	5	1

This situation is also illustrated in Figure 6.1. There, we show the values T_j, $j = 0, \ldots, d-1$, from (2.22) and (2.11) in the anchored and in the classical case, respectively. The results quantify the impact of the different path constructions on the modeling error. The smaller the values T_j the more important contributions are concentrated in the first j dimensions by the path construction. One can see that BB leads to a significantly faster decay of the importance of the dimensions as RW. We observe in Figure 6.1(a) that the level-wise construction of the Brownian bridge (the ℓ-th level corresponds to the first 2^ℓ dimensions) is mirrored in the decay of the values T_j. The PCA in turn leads to a significantly faster decay as BB. The LT even reduces the problem to only one effective dimension. In Figure 6.1(b) one can see that the values T_j in the classical case show a similar overall behaviour as in the anchored case but exhibit oscillations in the small T_j-values. These oscillations result can be reduced if we spend more sampling points in the employed approximation algorithm from [159] which unfortunately easily becomes to expensive, though. One can see that the oscillations are avoided in the anchored case although the results are significantly cheaper to obtain.

In summary, we see that the Vasicek bond pricing problem (6.3) is of very low superposition dimension d_s. It is also of low truncation dimension if BB, PCA or LT is used. Moreover, we observe that the effective dimensions in the anchored case provide the same or almost the same information on the importance of dimensions as in the classical case but have the advantage that they are significantly cheaper to compute. While the computation of the effective dimensions in the classical case require the computation of *many* high-dimensional integrals with up to $2d-1$ di-

[5] This way the closed-form pricing formula is derived.

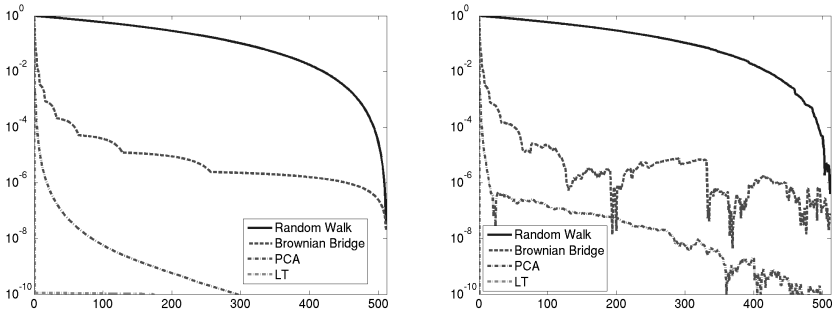

(a) Anchored case (2.22) (b) Classical case (2.11)

Fig. 6.1 Decay of the importance of the dimensions of the Vasicek bond pricing problem (6.3) with $d = 512$. Note that the results of the LT-construction almost coincide with the y-axis.

mensions,[6] the computation of the effective dimensions in the anchored case require only the computation of *one* d-dimensional integral with a sparse grid method or an other method of the class (3.18). Note, however, that in general the effective dimensions in the anchored case does not coincide with the effective dimension in the classical case. We give a counterexample in Section 6.1.2.

6.1.1.3 Quadrature Error versus Costs

We next compute the integral value (6.3) using different numerical quadrature and path generating methods. By a comparison of the results with the analytical solution from Lemma 6.1, we determine the relative errors of the different numerical methods. The relative errors are shown in Figure 6.2 for different numbers n of function evaluations.

One can see that the convergence rate of the MC method is always about 0.5 as predicted by the law of large numbers. The rate is not affected by the path construction since the total variance stays unchanged. The convergence rate of the QMC method increases if BB, PCA or LT is used since these path constructions concentrate the total variance in the first few dimensions. This way QMC outperforms MC and achieves higher convergence rates of almost one, smaller relative errors and a less oscillatory convergence behaviour. We furthermore observe in Figure 6.2 that the convergence of SGP and SGH is significantly accelerated by the path constructions BB, PCA and PCA which lead to a low effective dimension.

Note here that two different regimes have to be distinguished to describe the convergence behaviour of these methods, compare, e.g., Figure 6.2(d). In the preasymptotic regime, SGP and SGH first search for the important dimensions and interac-

[6] The computation of the superposition dimension in the classical case in addition suffers from cancellation problems and costs which are exponential in the superposition dimension.

tions, whereas, in the asymptotic regime, the important dimensions are identified and the grid is then refined only in these directions. Since the LT-construction reduces the problem to only one dimension, its combination with dimension-adaptive methods is particularly efficient. We see from Figure 6.2(d) that SGP and SGH correctly recognize the solely important dimension and then only refine in this respect, which leads to an extremely rapid convergence in the asymptotic regime. This way the dependence on the dimension is completely avoided in the asymptotic regime. There, the convergence of the methods is as fast as it is known for univariate problems despite the high nominal dimension $d = 512$. Overall, SGH is the most efficient method. It is more efficient than SGP since it avoids the singular transformation to the unit cube. It outperforms (Q)MC by several orders of magnitude independent of the employed path construction. By exploiting the low effective dimension and smoothness of the integrand, SGH achieves in combination with PCA or LT almost machine accuracy with only about 1,000 function evaluations.

While these results can not be easily transferred to general situations the nevertheless illustrate the improvements with respect to accuracy and computing time which are possible in comparison to a standard approach (such as MC) if the integrand satisfies certain properties (such as smoothness) and if these properties are exploited by suitable numerical techniques.

To better understand the fast performance of the SGH method we look at the index set $\mathscr{I}$ that is build up by Algorithm 4.3 in a dimension-adaptive way. For visualization we consider the two-dimensional slices through the index set $\mathscr{I}$ that correspond to the set of 512-dimensional indices $\mathbf{k}$ of the form $(k_1, k_2, 1, \ldots, 1)$ and $(k_1, 1, \ldots, 1, k_2)$ for $k_1, k_2 \geq 1$, respectively. The resulting index sets are shown in Figure 6.3 for the example that the Brownian bridge construction is used. In Figure 6.3 all indices $\mathbf{k}$ are marked with a dot that are included in the index set $\mathscr{I}$ which is build up by Algorithm 4.3 if the threshold $\varepsilon = 10^{-3}$ is used, compare also Figure 4.2. One can see for instance that the index $(3, 1, \ldots, 1)$ is included in $\mathscr{I}$ but not the index $(1, \ldots, 1, 3)$. Moreover, the values $|\Delta_{\mathbf{k}} f|$ are shown color-coded from 10^0 (black) to 10^{-16} (white) for $1 \leq k_1, k_2 \leq 7$. One can see that these values are decaying rapidly for increasing k_1, k_2. They are already below 10^{-10} if $k_1 > 4$ or $k_2 > 4$. With respect to the dimension 512 we see in Figure 6.3(b) that $|\Delta_{\mathbf{k}} f| < 10^{-10}$ already if $k_2 > 1$, which reflects the low truncation dimension of the Vasicek problem with the Brownian bridge construction. The results shown in Figure 6.3 indicate for this particular example that already a rather small index set $\mathscr{I}$ suffices to capture all indices $\mathbf{k}$ that correspond to significant contributions to the integral value, which explains why high precisions can be achieved with only little costs by the dimension-adaptive SGH method.[7]

[7] Note that the results shown in Figure 6.3 also indicate that Algorithm 4.3 indeed correctly identifies the most important contributions to the integral value.

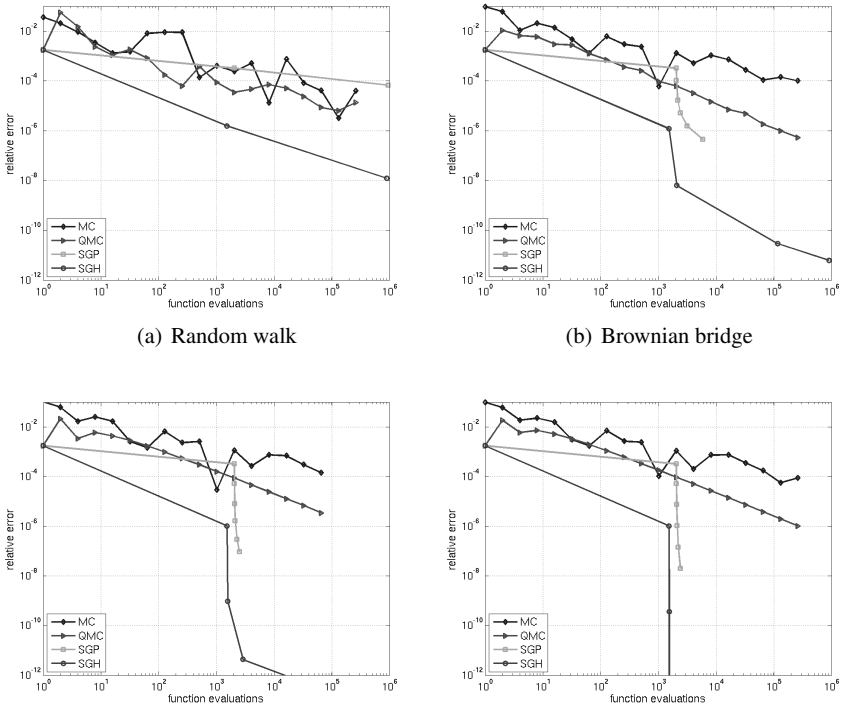

(a) Random walk (b) Brownian bridge

(c) Principal components (d) Linear transformation

Fig. 6.2 Convergence behaviour of the different numerical methods for the Vasicek bond pricing problem (6.3) with $d = 512$.

6.1.2 Collateralized Mortgage Obligations

The problem to price mortgage-backed securities is a commonly used test problem from finance to study the performance of numerical quadrature algorithms. It was first presented in [129]. There, a collateralized mortgage obligation (CMO) is calculated as 360-dimensional integral which was given to the authors Paskov and Traub by the investment bank Goldman Sachs. They observed that QMC methods outperform MC for this problem, which initiated extensive research to generalise their results to further problems in finance and to explain the success of the QMC despite the high dimension. The CMO problem is also considered in [115, 17] to study the performance of QMC methods and in [45, 46, 130] to demonstrate the efficiency of SG methods.

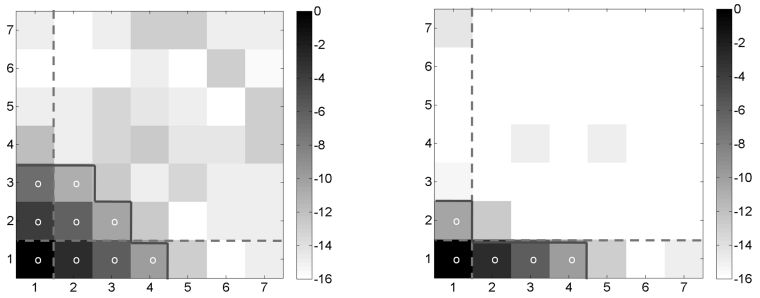

(a) Dimensions 1 and 2 (b) Dimensions 1 and 512

Fig. 6.3 Two-dimensional slices through the index set $\mathscr{I}$ corresponding to the threshold $\varepsilon = 10^{-3}$ for the Vasicek problem with the Brownian bridge construction ($d = 512$).

6.1.2.1 Modeling

We next describe the CMO problem following [17]. We consider a mortgage-backed security with a maturity of d months. Its holder receives d payments m_k of an underlying pool of mortgages at the times $t_k = \Delta t\, k$, $k = 1, \ldots, d$, where Δt denotes the period length of one month. The present value of the sum of all payments is then given by

$$g(\mathbf{z}) := \sum_{k=1}^{d} u_k m_k$$

where

$$u_k := \prod_{j=0}^{k-1} (1 + i_j)^{-1}$$

is the discount factor for the month k, which depends on the interest rates i_j. The interest rate i_k for month k is modelled by

$$i_k := K_0^k e^{\sigma(z_1 + \ldots + z_k)} i_0$$

using k many standard normally distributed random numbers z_j, $j = 1, \ldots, k$. Here, i_0 is the interest rate at the beginning of the mortgage, σ is a positive constant and $K_0 := e^{-\sigma^2/2}$ such that $\mathbb{E}[i_k] = i_0$. The payments m_k at time t_k are given by

$$m_k := c\, r_k((1 - w_k) + w_k c_k),$$

where c denotes the monthly payment. Moreover,

$$c_k := \sum_{j=0}^{d-k} (1+i_0)^{-j},$$
$$r_k := \prod_{j=1}^{k-1} (1-w_j),$$
$$w_k := K_1 + K_2 \arctan(K_3 i_k + K_4).$$

Here, r_k denotes the fraction of remaining mortgages at month k, which in turn depends on the fraction w_k of mortgages which prepay in month k. The values K_1, K_2, K_3, K_4 are constants of the model for the prepayment rate w_k.

The expected present value of the sum of all payments can then be written as d-dimensional integral

$$PV := \int_{\mathbf{R}^d} g(\mathbf{z}) \varphi_d(\mathbf{z}) d\mathbf{z} \tag{6.5}$$

on $\mathbf{R}^d$ with Gaussian weight φ_d or using the usual transformation with the inverse of the cumulative normal distribution function Φ as d-dimensional integral

$$PV := \int_{[0,1]^d} f(\mathbf{x}) d\mathbf{x}$$

on $[0,1]^d$ with the integrand $f(\mathbf{x}) = g(\Phi^{-1}(x_1), \ldots, \Phi^{-1}(x_d))$.

We next focus on the efficient numerical computation of this integral. In our numerical experiments we thereby use the parameters

$$i_0 = 0.007, c = 1, K_1 = 0.01, K_2 = -0.005, K_3 = 10, K_4 = 0.5 \text{ and } \sigma = 0.0004$$

which are also used in [17, 45, 130]. We consider the case $d = 256$ for which we obtain the reference solution $PV = 119.21588257$.

6.1.2.2 Effective Dimensions

To illustrate the impact of the different path constructions we first consider the effective dimensions of the integral (6.5). From Table 6.2, we see that the CMO prob-

	anchored-ANOVA superposition dim.				anchored-ANOVA truncation dim.				classical ANOVA truncation dim.			
$1-\alpha$	RW	BB	PCA	LT	RW	BB	PCA	LT	RW	BB	PCA	LT
1e-1	1	1	1	1	123	18	13	123	60	5	2	1
1e-2	2	2	1	2	191	134	108	191	110	10	5	1
1e-3	2	2	2	2	225	192	235	225	158	36	11	1
1e-4	2	2	2	2	242	229	254	242	181	80	23	1

Table 6.2 Effective dimensions of the CMO pricing problem (6.5). The nominal dimension is $d = 256$.

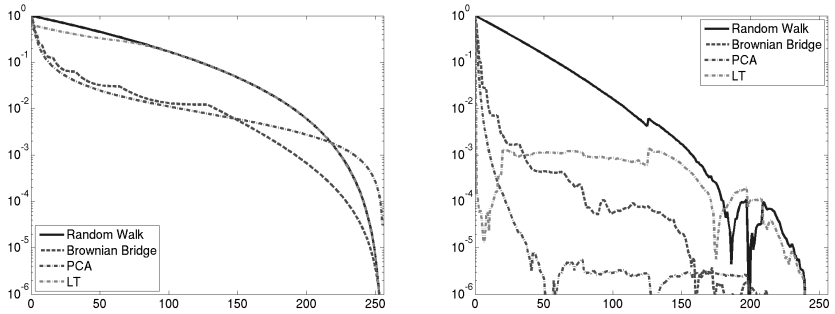

(a) Anchored case (2.22) (b) Classical case (2.11)

Fig. 6.4 Decay of the importance of the dimensions of the CMO pricing problem (6.5) with $d = 256$.

lem is of very small superposition dimension $d_s \leq 2$ in the anchored case for all $\alpha \in [0.9, 0.9999]$ and all path constructions.

In Table 6.2 also the truncation dimensions d_t of this problem are shown. It is striking that the path construction has only a small impact on the truncation dimension in the anchored case, i.e., the advantage of BB, PCA and LT compared to RW is not so clear for the CMO problem. For $\alpha = 0.9$ we have $d_t = 123$ in case of RW and LT. This truncation dimension is reduced to $d_t = 18$ and $d_t = 13$ if BB and PCA is used, respectively. For higher accuracy requirements, however, i.e. for $\alpha \geq 0.99$, significantly less or even no reduction at all is achieved with these constructions. Note that for the CMO problem the truncation dimensions d_t in the classical case clearly differ from the truncation dimensions in the anchored case. In the classical case, BB, PCA and LT lead to significant dimension reductions. LT even reduces the problem to the truncation dimension one.

This situation is also illustrated by the values T_j from (2.22) and (2.11) which we show in Figure 6.4. The differences between the anchored and the classical case are clearly visible. While LT leads to very similar results as RW in the anchored case, LT leads to a large dimension reduction compared to RW in the classical one. In the anchored case, BB turns out to be the most efficient path construction, whereas LT is most efficient in the classical case.[8] The results show that different path constructions can be different effective depending on the measure (classical or anchored) which is used to describe the importance of dimensions.

[8] Since LT aims by construction to minimize the effective dimension in the classical case it is to expect that LT gives poorer results in the anchored case as in the classical one.

6.1.2.3 Quadrature Error versus Costs

We next study the convergence behaviour of the different numerical methods for the CMO pricing problem. To measure their accuracy we used the reference solution $PV = 119.21588257$. The respective numerical results are illustrated in Figure 6.5. One can see that the QMC method converges faster, less oscillatory and superior to MC if we switch from RW to BB, PCA or LT. SGP performs similar to QMC in case of BB and PCA and slightly worse in case of RW and LT. SGH combined with BB or PCA is most efficient for the CMO problem. It achieves the highest convergence rate and the most precise results. With 10^4 function evaluations SGH obtains a relative error which is about 100 times smaller than the relative error of the QMC method.

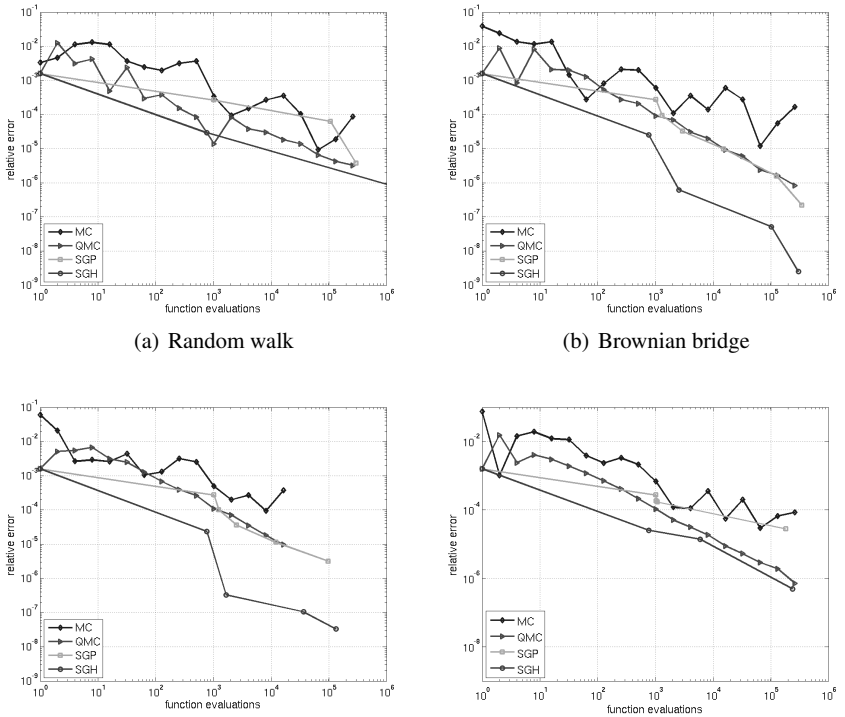

(a) Random walk (b) Brownian bridge

(c) Principal components (d) Linear transformation

Fig. 6.5 Convergence behaviour of the different numerical approaches for the CMO pricing problem (6.5) with $d = 256$.

We next discuss the relation of the convergence behaviour of the numerical methods to the effective dimension of the CMO problem. We already showed that the path construction affects both the performance of the numerical methods (except for MC) and the truncation dimension of the integral. Since the truncation dimen-

sion in the classical case differs from the truncation dimension in the anchored case for this problem, it is interesting to see which of these two notions better predicts the convergence behaviour of the numerical methods. Remember that LT does not lead to an improved convergence of the dimension-adaptive sparse grid methods SGP and SGH compared to RW. This observation can not be explained by the effective dimension in the classical case since LT obtains the optimal result $d_t = 1$ for the CMO problem. The observation is, however, in clear correspondence with the fact that LT provides no reduction of the truncation dimension in the *anchored* case. This indicates that the performance of SGP and SGH depends on the truncation dimension in the anchored case, but not on the truncation dimension in the classical case. Note that the convergence behaviour of the QMC method is rather related to the effective dimension d_t in the classical case than to the anchored one, since this method converges faster and less oscillatory with LT than with RW.

The different effective dimensions in the anchored and in the classical case are related to the fact that in the anchored-ANOVA decomposition the contributions $If_{\mathbf{u}}$ are of varying sign in the CMO problem. Summing the contributions thus leads to cancellation effects which are not seen in the anchored case since the absolute values $|If_{\mathbf{u}}|$ are taken into account there. Nevertheless, also the error indicator[9] of dimension-adaptive sparse grid methods is based on the absolute values $|\Delta_{\mathbf{k}}|$. These values are closely related to $|If_{\mathbf{u}}|$ as we showed in Section 4.3. The methods SGP and SGH can thus also not profit from such cancellation effects and their convergence behaviour therefore rather depends on the effective dimensions in the anchored case than on the effective dimension in the classical case.

Note finally that the truncation dimension d_t in the anchored case explains the impact of the path construction but not the high performance of the SGH method since d_t is high for this problem. The fast convergence is explained by the low superposition dimension $d_s \leq 2$ and by the smoothness of the integrand.

6.2 Path-dependent Options

In this section, we consider the pricing of European-style path-dependent options in the Black-Scholes model.[10] We denote the maturity date of the option by T and assume that the payoff of the option at time T depends on the prices $S(t)$ of an underlying asset at the equidistant points in time $t_k := \Delta t\, k$, $k = 1, \ldots, d$, with $\Delta t := T/d$. We collect these prices in the vector $\mathbf{S} := (S(t_1), \ldots, S(t_d))$.

We assume that the underlying asset follows a geometric Brownian motion

$$dS(t) = \mu S(t)dt + \sigma S(t)dW(t). \tag{6.6}$$

Here, μ denotes the drift, σ denotes the volatility and $W(t)$ is the standard Brownian motion. Using the usual Black-Scholes assumptions, see, e.g., [75, 86], the price of

[9] See Algorithm 4.3.

[10] For more information on the Black-Scholes model we refer to [75].

http://www.springer.com/978-3-642-16003-5

Sparse Grid Quadrature in High Dimensions with Applications in Finance and Insurance

Holtz, M.

2011, VIII, 192 p. 32 illus., Hardcover
ISBN: 978-3-642-16003-5