



Efficient exposure computation by risk factor decomposition

C. S. L. de Graaf, D. Kandhai & C. Reisinger

To cite this article: C. S. L. de Graaf, D. Kandhai & C. Reisinger (2018) Efficient exposure computation by risk factor decomposition, Quantitative Finance, 18:10, 1657-1678, DOI: [10.1080/14697688.2018.1435902](https://doi.org/10.1080/14697688.2018.1435902)

To link to this article: <https://doi.org/10.1080/14697688.2018.1435902>



© 2018 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 24 Apr 2018.



Submit your article to this journal [↗](#)



Article views: 316



View Crossmark data [↗](#)

Efficient exposure computation by risk factor decomposition

C. S. L. DE GRAAF^{*†‡}, D. KANDHAI^{†§} and C. REISINGER[‡]

[†]Computational Science Lab, University of Amsterdam, 1098 XH, Amsterdam, The Netherlands

[‡]Mathematical Institute & Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, OX2 6GG/6ED, UK

[§]Quantitative Analytics, ING Bank, 1102 BD Amsterdam, The Netherlands

(Received 6 November 2016; accepted 25 January 2018; published online 24 April 2018)

The focus of this paper is the efficient computation of counterparty credit risk exposure on portfolio level. Here, the large number of risk factors rules out traditional PDE-based techniques and allows only a relatively small number of paths for nested Monte Carlo simulations, resulting in large variances of estimators in practice. We propose a novel approach based on Kolmogorov forward and backward PDEs, where we counter the high dimensionality by a generalization of anchored-ANOVA decompositions. By computing only the most significant terms in the decomposition, the dimensionality is reduced effectively, such that a significant computational speed-up arises from the high accuracy of PDE schemes in low dimensions compared to Monte Carlo estimation. Moreover, we show how this truncated decomposition can be used as control variate for the full high-dimensional model, such that any approximation errors can be corrected while a substantial variance reduction is achieved compared to the standard simulation approach. We investigate the accuracy for a realistic portfolio of exchange options, interest rate and cross-currency swaps under a fully calibrated 10-factor model.

Keywords: Counterparty credit risk estimation; Nested conditional expectations; Anchored-ANOVA decompositions for high-dimensional PDEs

1. Introduction

Managing Counterparty Credit Risk (CCR) has become one of the core activities of financial institutions since the credit crisis in 2008. The Basel III rules introduced Credit Value Adjustment (CVA) which increased counterparty risk capital requirements (Gregory 2015) and has significant consequences for the valuation of derivatives. The CVA is the difference between the value of a portfolio without default risk and the true portfolio value which includes default risk of the counterparty (Pykhtin and Zhu 2007). By including CVA, one can theoretically hedge away credit risk. More recently, Debt Value Adjustment (DVA), Funding Value Adjustment (FVA) and Capital Value Adjustment (KVA) (Green *et al.* 2014) were defined and this collection of value adjustments is now referred to as XVA. For all these adjustments, Expected Exposure (EE) and Expected Positive Exposure (EPE) are main ingredients that have to be computed from the future distribution of the underlying portfolio. In this paper, we develop an approximation method for these exposures on portfolio level.

Many traded portfolios consist of multiple underlying assets and derivatives on these assets; these are valued by models

with multiple risk drivers such as stochastic interest rates, FX rates and stochastic volatility. Computing the future value distributions, for which closed-form solutions are typically not available, is a challenging high-dimensional computational problem. A typical benchmark for methods geared towards high-dimensional problems are basket options, where the payoff depends on a portfolio of assets, such as stocks, stock indices or currencies. These options serve the purpose of diversification and are therefore popular among investors (Bouzoubaa and Osseiran 2010). For the valuation of these options, contributions have been made by, e.g. Jain and Oosterlee (2012) for Monte Carlo estimation under early exercise and Reisinger and Wissmann (2015) for PDE approximations.

In the context of CVA, netting needs to be taken into account and possible negative or positive parts of the exposure distribution are to be considered. This makes CVA estimation for simple plain vanilla instruments similar to the valuation of a series of basket options with different maturity, but is more complex if the instruments themselves are more complex, i.e. no closed-form or semi-closed-form pricing formulae are available. The computational problem in that case can be formulated as the estimation of nested conditional expectations, where the inner expectations are future derivative prices conditional on the

*Corresponding author. Email: csldegraaf@gmail.com

underlying risk factors, and the outer expectations, taken over the risk factors, are the expected exposures for the derivative portfolio.

For computing these risk measures, typically, Monte Carlo methods are used to sample the future states of the underlying risk factors by discretization and simulation. Portfolio values for all these states can then be computed in multiple ways, including: fully nested Monte Carlo simulation (Gordy and Juneja 2010, Broadie *et al.* 2011), where derivative values along those paths are estimated by an ‘inner’ Monte Carlo simulation; regression-based approaches, introduced for American style options in Longstaff and Schwartz (2001) and applied in the context of exposure in Karlsson *et al.* (2014) and Joshi and Kang (2016); or a grid or Fourier-based method for the inner expectations and Monte Carlo sampling of the outer expectation (de Graaf *et al.* 2014).

All these different methods have their particular strengths and weaknesses. The naïve nested Monte Carlo method is generically applicable but computationally extremely expensive with a complexity of $O(\epsilon^{-4})$ for root-mean-square error (RMSE) ϵ . This can be improved when the computational budget between inner and outer simulations is optimally balanced, e.g. $O(\epsilon^{-3})$ when trading off ‘inner’ bias and ‘outer’ variance (Gordy and Juneja 2010), and $O(\epsilon^{-5/2-\delta})$ for arbitrary $\delta > 0$ for certain non-uniform estimators (Broadie *et al.* 2011). Regression-based approaches have in a sense optimal complexity $O(\epsilon^{-2})$ as no extra ‘inner’ paths are required (Broadie *et al.* 2015), but an additional bias creeps in when leaving out sets of regression variables in higher dimensions and, in our experience, regression methods still exhibit relatively slow convergence in applications. In practical portfolio CVA computations, only about 10^3 to 10^4 paths are feasible, giving a high variance of estimators for both of these approaches, in particular for the sensitivities with respect to market factors. PDE schemes, on the other hand, are very efficient for low dimensions but suffer from the ‘curse of dimensionality’—a term used to describe the computational complexity which increases exponentially with the number of risk factors—which made Monte Carlo methods the industry standard technique for problems with dimensions larger than two or three.

Notwithstanding this, we propose a method based on PDEs not just for the inner, but also for the outer expectation. We address the curse of dimensionality by decomposition of the conditional expectations into groups of risk factors, hence breaking the high-dimensional problem up into a sequence of lower dimensional problems, whose solutions are then assembled into a truncated decomposition. For these low-dimensional problems numerical PDE methods are available which give sufficiently accurate approximations in acceptable time. We will find that for practically sufficient accuracy, the approximation requires solving only one one-dimensional base approximation and multiple two- and three-dimensional approximations which are used as corrections. Perhaps practically most relevant, the approximated portfolio dynamics underlying these expansions can be used to construct control variates for Monte Carlo estimators. As such, we do not see this methodology to compete with the above simulation approaches, but to be used in conjunction with them, as it provides a way to combine seamlessly the high accuracy of PDE schemes in low dimensions and the robustness of Monte Carlo schemes in high dimensions.

This includes improvements of Monte Carlo schemes using quasi-Monte Carlo sampling (both for regression and nested simulation).

Note that our approach is distinctly different from dimension and variance reduction methods based on conditional sampling as in Dang *et al.* (2015), Cozma and Reisinger (2017) and Dang *et al.* (2017).

The method proposed here draws on a variety of related previous work, but requires substantial extensions for the present context:

- in terms of application and problem structure, the method extends Reisinger and Wittum (2005) from basket options under Black–Scholes to portfolio risk management under a general class of models;
- specifically, for the exposure estimation problem formulated as nested conditional expectations, we solve a combination of forward and backward Kolmogorov equations and carry out numerical integration of the solutions;
- compared to the dimension-wise integration method of Griebel and Holtz (2010), the future distribution here is unknown, such that forward PDEs have to be approximated;
- compared to Reisinger and Wittum (2005), Reisinger and Wissmann (2015), we use original variables as factors instead of principal components; this makes the results more easily interpretable;
- compared to Reisinger and Wissmann (2017) which provides error bounds for the constant coefficient setting, we deal with the case of non-linear SDEs and variable-coefficient PDEs, which requires a generalized definition of the expansion;
- instead of theoretical error bounds, we assess the method on a complex real-world example, as outlined in the following.

For the numerical illustration, we consider typical FX portfolios consisting of multiple financial derivatives, namely Interest Rate Swaps (IRS), Cross Currency Basis Swaps (CCYS) and FX options. The risk factors are driven by a multi-dimensional Black–Scholes-2-Hull–White (BS2HW) model which captures stochasticity in the FX and the interest rates.

In actual practical applications, one may want to consider FX and interest rate smiles and a multi-curve framework in which a stochastic basis model should be used. In terms of the latter, however, little is known on how to calibrate the basis vol in an implied measure, and therefore in most pricing models basis is considered as deterministic. In the context of CCR, the problem is even more complex since one can argue that basis is a mixture of credit and liquidity and therefore, it is not clear if it should be modelled as a pure IR component or whether it is a separate risk factor strongly correlated with credit. As such, the question is linked with other modelling complexities such as wrong-way risk.

These modelling aspects are beyond the scope of the paper. We emphasize though that the computational framework presented here is rich enough to incorporate such extensions. We demonstrate the flexibility of the approach by giving results for an extension including stochastic volatility, namely the Heston-2-Hull–White model. For the three chosen currency pairs this

leads to a 10-dimensional PDE, where we demonstrate that a relative accuracy of about 1% can be achieved for expected positive exposure computations.

The models are calibrated to market data and for these different portfolios we compute EE and EPE and compare our results with a full-scale Monte Carlo benchmark, demonstrating vastly superior computational complexity.

The outline of this paper is as follows. In the following section 2, we describe the problem set-up. In section 3, we explain the main approximation technique used. (A brief description of the finite difference schemes used for the PDEs in the decomposition is given in appendix 2.) In section 4, we set up our case studies followed by section 5 where we present our results. In section 6, we provide a discussion and concluding remarks.

2. Problem formulation

We consider the general framework of a financial market described by a d -dimensional stochastic process $X = (X_t^i)_{i=1 \leq i \leq d}^{t \geq 0}$. In our applications, X will be given by a stochastic differential equation of the form:

$$dX_t^i = \mu_i(X_t, t) dt + \sigma_i(X_t, t) dW_t^i, \quad i = 1, \dots, d, \quad t > 0, \quad (1)$$

$$X_0^i = a_i, \quad i = 1, \dots, d, \quad (2)$$

where W is a d -dimensional standard Brownian motion with a given correlation matrix $(\rho_{i,j})_{1 \leq i,j \leq d}$, and $a \in \mathbb{R}^d$ a given initial state, μ_i is the drift, σ_i the volatility. We assume that the processes are written under a risk-neutral measure \mathbb{Q} , which is relevant for derivative valuation as well as regulatory CVA computations.

We now consider derivatives on X . For simplicity, we focus on the European case here. Let the pay-off function ϕ_k determine the amount $\phi_k(X_{T_k})$ to be received by the holder at time T_k for $1 \leq k \leq N_D$. The arbitrage-free value V_k of this claim for $0 \leq t \leq T$ is then,

$$V(x, t) = \mathbb{E}[\exp(-\int_t^T \rho(X_s) ds) \phi(X_T) | X_t = x],$$

where ρ is a discount factor (and we have dropped the subscripts k for simplicity), and satisfies the Kolmogorov backward PDE (see, e.g. Musiela and Rutkowski 2005),

$$\frac{\partial V}{\partial t} + \sum_{i=1}^d \mu_i \frac{\partial V}{\partial x_i} + \frac{1}{2} \sum_{i,j=1}^d \sigma_i \sigma_j \rho_{i,j} \frac{\partial^2 V}{\partial x_i \partial x_j} - \rho V = 0, \quad (3)$$

$$V(x, T) = \phi(x).$$

For Counterparty Credit Risk (CCR), not only the time 0 value is important, but also the evolution of this value after the deal has been entered. In this work, we use several measures to quantify CCR. Let

$$\sum_{k=1}^{N_D} w_k(t) V_k(X_t, t)$$

be the value of a portfolio of N_D derivatives at time t that a financial institution holds, where w_k , $1 \leq k \leq N_D$, are derivative portfolio weights. Here, we set $w_k(t) = 0$ for $t > T_k$, the expiry of the k -th derivative.

First, following Pykhtin and Zhu (2007), Expected Exposure (EE) is defined as the expected value of the portfolio at time $t \geq 0$,[†]

$$EE(t) = \mathbb{E} \left[\sum_{k=1}^{N_D} w_k(t) V_k(X_t, t) \middle| \mathcal{F}_0 \right], \quad (4)$$

where \mathcal{F}_0 is the filtration at time $t = 0$. Similarly, the Expected Positive Exposure (EPE) at a future time $t < T := \max_k T_k$ is given by,

$$EPE(t) = \mathbb{E} \left[\max \left(0, \sum_{k=1}^{N_D} w_k(t) V_k(X_t, t) \right) \middle| \mathcal{F}_0 \right]. \quad (5)$$

Note that the Expected Negative Exposure (ENE) can be computed by taking the minimum instead of the maximum in equation (5), or from $ENE = EE - EPE$. In short, EPE (ENE) gives the expected loss of the buyer (seller) of the portfolio in case he or she is in the money and the counterparty defaults, which makes them of prime importance when computing CVA and other XVAs. For example, CVA is defined as,

$$\begin{aligned} CVA &= (1 - R) \int_0^T \mathbb{E} \left[\exp(-\int_t^T \rho(X_s) ds) \max \right. \\ &\quad \times \left. \left(0, \sum_{k=1}^{N_D} w_k(t) V_k(X_t, t) \right) \middle| \mathcal{F}_0 \right] dPD(t), \quad (6) \\ &= (1 - R) \int_0^T D(t) EPE^*(t) dPD(t), \end{aligned}$$

where R is the recovery rate, $EPE^*(t)$ is the discounted EPE and $PD(t)$ denotes the default probability of the counterparty at time t (Gregory 2010). Note that (6) incorporates the debatable but common assumption that the portfolio value is independent of default events (no ‘wrong-way risk’).

In this paper, we focus solely on the computation of EE and EPE. For exposure calculations at time 0, we then consider functionals of the form,

$$\begin{aligned} V(x, 0; t) &= \mathbb{E} \left[\Psi \left(\sum_{k=1}^{N_D} w_k(t) V_k(X_t, t) \right) \middle| X_0 = x \right] \\ &= \mathbb{E} \left[\Psi \left(\sum_{k=1}^{N_D} w_k(t) \mathbb{E} [\phi_k(X_{T_k}) | X_t] \right) \middle| X_0 = x \right], \end{aligned} \quad (7)$$

where Ψ is an exposure function, e.g. $\Psi(x) = x$ in the case of EE, or $\Psi(x) = \max(x, 0)$ for EPE. This can be adapted in the obvious way for American or path-dependent options.

Remark 1 To find $V(x, 0; t)$, the expected (positive) exposure at time t as seen from time 0, we could therefore solve backward PDEs in (t, T) for $V_k(\cdot, t)$, and then another backward PDE over $(0, t)$ with $\Psi[\sum_{k=1}^{N_D} w_k(t) V_k(\cdot, t)]$ as terminal condition. This would however require the solution of a different backward PDE in the second stage for different t . To avoid this, in the following we utilize a forward PDE for the transition densities to compute exposures.

The reason we use the backward PDE to approximate V_k instead of using the density function, is that we need the value of derivatives in the whole state-space reachable by X .

[†]Here, we do not include collateralization. This can be added by including a collateral model on top of the exposure dynamics.

Let $p(x, t; a, 0)$ be the transition density function of X_t at x given state a at $t = 0$, then

$$V(x, 0; t) = \int_{\mathbb{R}} p(y, t; x, 0) \Psi \left(\sum_{k=1}^{N_D} w_k(t) V_k(y, t) \right) dy. \quad (8)$$

Here, p is given through an adjoint relation by the Kolmogorov forward equation,

$$-\frac{\partial p}{\partial t} - \sum_{i=1}^d \frac{\partial}{\partial x_i} (\mu_i p) + \frac{1}{2} \sum_{i,j=1}^d \frac{\partial^2}{\partial x_i \partial x_j} (\sigma_i \sigma_j \rho_{i,j} p) = 0, \quad (9)$$

$$p(x, 0; a, 0) = \delta(x - a), \quad (10)$$

where δ is the Dirac distribution centred at 0.

Hence, we can obtain expected (positive) exposures by the solution of one forward PDE for the density and one backward PDE for each derivative in the portfolio, plus one integration for each t .

3. Approximation by risk factor decomposition

The principal difficulty in practice arises from the dimensionality d of X in (1). Although N_D in (7) is typically very large, the computational complexity is linear in N_D , while it is exponential in d . In this section, we introduce an approximation technique which makes large d computationally manageable.

To introduce the concepts, we focus in section 3.1 on the problem of approximating

$$V(x, t) = \mathbb{E}[\phi(X_T) | X_t = x] \quad (11)$$

$$= \int_{\mathbb{R}^d} \phi(y) p(y, T; x, t) dy, \quad (12)$$

i.e. for simplicity the discount factor is 1, where $p(\cdot, T; x, t)$ is the probability density function of X_T . We first discuss an extension of the anchored-ANOVA concept to PDEs with variable coefficients and the use of a ‘moving anchor’. In section 3.2, we explain how the approximations can be used as effective control variates for unbiased Monte Carlo estimators. The application to the complete expected exposure problem (8) will be discussed in section 3.3.

3.1. An anchored-ANOVA-type approximation

The proposed method extends anchored-ANOVA decompositions as considered by Griebel and Holtz (2010) in the context of integration problems. In the setting and notation of (12), the methodology in Griebel and Holtz (2010) chooses an anchor $a \in \mathbb{R}^d$ and then, for a given index set $u \subset \mathcal{D} = \{i : 1 \leq i \leq d\}$, defines projections of the integrand $f = \phi p$ by $f(a \setminus x_u)$, where $a \setminus x_u$ denotes a d -vector such that,

$$(a \setminus x_u)_i = \begin{cases} x_i & i \in u, \\ a_i & i \notin u. \end{cases} \quad (13)$$

This leads to a decomposition of f into lower dimensional functions, which can be exploited for successive quadrature approximations (see Griebel and Holtz 2010 and the references therein).

In contrast to there, we cannot assume here that the joint probability density function for X_T is analytically known and therefore, we have to consider the dynamics of X directly. To improve the accuracy, e.g. to account for the term-structure of interest rates, here we do not concentrate the measure at a fixed anchor a , but we allow the anchor to move with a certain conditional expectation of the underlying process, as detailed in the following.

We first define a deterministic approximation to X from (1), (2) by,

$$\xi^i(t) = \mathbb{E}[\bar{X}_t^i | \bar{X}_0^i = a_i], \quad 1 \leq i \leq d, \quad t \geq 0,$$

$$\begin{aligned} \bar{X}_t^i &= a_i + \int_0^t \mu_i(a_1, \dots, a_{i-1}, \bar{X}_s^i, a_{i+1}, \dots, a_d, s) ds \\ &\quad + \int_0^t \sigma_i(a_1, \dots, a_{i-1}, \bar{X}_s^i, a_{i+1}, \dots, a_d, s) dW_s^i, \quad t \geq 0. \end{aligned}$$

In Reisinger and Wissmann (2015), a more simplistic approach is used where the drift is first approximated by $\mu(a, 0)$, and then eliminated by a coordinate transformation $x \rightarrow x - \mu(a, 0)t$. The above construction takes into account more information about the term structure and also deals systematically with non-constant volatilities.

For given subset $u \subseteq \mathcal{D}$, we then define a process,

$$\begin{aligned} X_t^u &= \begin{cases} a_i + \int_0^t \mu_i(\xi(t) \setminus X_s^u, s) ds + \int_0^t \sigma_i(\xi(t) \setminus X_s^u, s) dW_s^i, & i \in u, \\ \xi^i(t), & i \notin u, \end{cases} \end{aligned} \quad (14)$$

where $\xi(t) \setminus X_s^u$ follows the notation from (13).

Here, we have replaced a subset of the processes by their conditional expectations. For instance, if X^1 is an exchange rate and X^2 and X^3 are the domestic and foreign short rate in a Hull–White model, then ξ^1 is the expectation of the exchange rate under a constant interest rate model, while ξ^2 is simply the expectation of the domestic short rate. The point is that the dynamics is effectively of dimension $|u|$, and the expectation (11) can be approximated by lower dimensional problems.

Accordingly, we define

$$V_u(a; x, t) = \mathbb{E}[\phi(X_T^u) | X_t^u = x], \quad (15)$$

where the right-hand side depends on a implicitly through the definition of X_t^u in (14).

Specifically, the backward PDE for (15) under (14) is,

$$\begin{aligned} \frac{\partial V_u}{\partial t} &+ \underbrace{\sum_{i \in u} \mu_i(\xi(t) \setminus x_u) \frac{\partial V_u}{\partial x_i} + \frac{1}{2} \sum_{i,j \in u} \sigma_i(\xi(t) \setminus x_u) \sigma_j(\xi(t) \setminus x_u) \rho_{i,j} \frac{\partial^2 V_u}{\partial x_i \partial x_j}}_{\equiv \mathcal{L}_u V_u} = 0, \end{aligned} \quad (16)$$

$$V_u(x, T) = \phi(x). \quad (17)$$

The significance of the arguments $\xi(t) \setminus x_u$ in the coefficients of (16) is that the full information on the coordinates x_j , $j \in u$ is used, while the solution $V(a; a, 0)$ at the anchor point can be computed solving only a $|u|$ -dimensional PDE in the variables x_j , $j \in u$.

Remark 2 The generalization from (15) to

$$V_u(a; x, t) = \mathbb{E}[\exp(-\int_t^T \rho(X_s^u) ds) \phi(X_T^u) | X_t^u = x]$$

is straightforward and the corresponding backward PDE instead of (16) is,

$$\frac{\partial V_u}{\partial t} + \mathcal{L}_u V_u + \rho(\xi(t) \setminus x_u) V_u = 0.$$

For future reference, we also have the forward PDE,

$$\begin{aligned} \frac{\partial p_u}{\partial t} + \sum_{i \in u} \frac{\partial}{\partial x_i} (\mu_i(\xi(t) \setminus x_u) p_u) \\ - \frac{1}{2} \sum_{i,j \in u} \frac{\partial^2}{\partial x_i \partial x_j} (\sigma_i(\xi(t) \setminus x_u) \sigma_j(\xi(t) \setminus x_u) \rho_{i,j} p_u) = 0, \end{aligned} \quad (18)$$

$$p_u(x, 0; a, 0) = \delta(x - a). \quad (19)$$

From here onwards, we can follow the path of [Griebel and Holtz \(2010\)](#) and the references therein, to define a decomposition recursively through a difference operator Δ , by $\Delta V_\emptyset = V_\emptyset$ and, for $u \neq \emptyset$,

$$\Delta V_u(a; x, t) = V_u - \sum_{w \subseteq u} \Delta V_w = \sum_{w \subseteq u} (-1)^{|w| - |u|} V_w, \quad (20)$$

where the inclusion in the first summation is strict.

This is indeed a decomposition because,

$$V(a; x, t) = \sum_{u \subseteq \{1, \dots, d\}} \Delta V_u = \sum_{k=0}^d \sum_{|u|=k} \Delta V_u. \quad (21)$$

We note that for $u = \{i_1, \dots, i_{|u|}\}$, $\Delta V_u(a; \xi(t) \setminus x_u, t)$ only depends non-trivially on the subset of coordinates $\{x_{i_1}, \dots, x_{i_{|u|}}\}$, and satisfies (16). Therefore, V_u and hence ΔV_u can be found by the solution of problems of dimension not higher than $|u|$.

As an example, consider $d = 3$ and $u = \{1, 2\}$, then

$$\begin{aligned} \Delta V_u(a; x, t) &= V_{\{1,2\}} - (\Delta V_{\{1\}} + \Delta V_{\{2\}} + \Delta V_\emptyset) \\ &= V_{\{1,2\}} - V_{\{1\}} - V_{\{2\}} + V_\emptyset. \end{aligned}$$

We can now define an approximation by,

$$V_{0,s}(a; x, t) = \sum_{k=0}^s \sum_{|u|=k} \Delta V_u = \sum_{k=0}^s c_k \sum_{|u|=k} V_u, \quad (22)$$

where c_k are integer constants, which also depend on s and d (we suppress this to keep the notation simple). The point is that the approximation $V_{0,s}(a; a, t)$ at the anchor can be found by solving PDEs of dimension at most s .

Remark 3 The approximations $V_{0,1}$ and $V_{0,2}$ have certain similarities with delta and delta-gamma approximations, respectively (see, e.g. [Alexander et al. 2006](#)). In the latter, the derivative value is approximated by,

$$\begin{aligned} V(X_t, t) &\approx V(a, 0) + \frac{\partial V}{\partial t} t + (X_t - a)^T \nabla V \\ &\quad + \frac{1}{2} (X_t - a)^T \nabla \nabla^T V (X_t - a), \end{aligned}$$

where $a = X_0$, ∇V is the gradient or delta and $\nabla \nabla^T V$ the Hessian or gamma. One notable difference is that $V_{0,2}$ does not make any approximations if V is two-dimensional or a sum of two-dimensional functions, while delta-gamma always involves linear and quadratic approximations.

An alternative approximation, extending (22), is given by,

$$V_{r,s}(a; x, t) = \sum_{k=0}^s c_k \sum_{|u|=k} V_{u \cup \{1, \dots, r\}}, \quad r + s \leq d. \quad (23)$$

Here, we always retain the first r coordinates, and apply the splitting only to the remaining $d - r$ coordinates. The approximation $V_{r,s}(a; a, t)$ at the anchor can be found by solving PDEs of dimension at most $r + s$. It is this approximation that we will use in the numerical computations later on in the paper, with $r = 1$ and $s = 1$ or $s = 2$. The coordinate x_1 is chosen to capture most of the dynamics, either through prior knowledge or small pilot runs with reduced accuracy.

For instance, for $s = 1, r = 1$, we have,

$$\begin{aligned} V_{1,1}(a; x, t) &= \sum_{1 < i \leq d} (V_{\{1,i\}} - V_{\{1\}}) + V_{\{1\}} \quad (24) \\ &= \sum_{1 < i \leq d} V_{\{1,i\}} - (d - 2) V_{\{1\}}, \end{aligned}$$

i.e. $c_0 = -(d - 2)$, $c_1 = 1$ in (23), and, for $s = 2, r = 1$,

$$\begin{aligned} V_{1,2}(a; x, t) &= \sum_{1 < i < j \leq d} (V_{\{1,i,j\}} - V_{\{1,i\}} - V_{\{1,j\}} + V_{\{1\}}) \\ &\quad + \sum_{1 < i \leq d} (V_{\{1,i\}} - V_{\{1\}}) + V_{\{1\}} \quad (25) \\ &= \sum_{1 < i < j \leq d} V_{\{1,i,j\}} - (d - 2) \sum_{1 < i \leq d} V_{\{1,i\}} \\ &\quad + \frac{(d - 2)(d - 1)}{2} V_{\{1\}}, \end{aligned}$$

i.e. $c_0 = \binom{d-1}{2}$, $c_1 = -(d - 2)$, $c_2 = 1$.

3.2. A control variate

In some situations, the approximation (23) with small r and s will be sufficiently accurate. There is an expectation—but no guarantee in general—that increasing either r or s will improve the accuracy. The computation of further terms may also be prohibitively expensive due to the high dimensionality and number of PDEs involved. In such cases, it will be valuable to simulate certain corrections by a MC method—instead of computing them by PDEs—to obtain more accurate approximations.

Therefore, we now discuss a way to turn the approximation (23) into a control variate for a Monte Carlo scheme for (11). Denote by $\{\omega_k : 1 \leq k \leq N_\omega\}$ a set of N_ω independent samples of the d -dimensional Brownian motion W , and $X_T(\omega_k)$ the (strong) solution to (1) for a given sample path ω_k of W .

The standard estimator of $V(x, 0)$ in (11) is then,

$$\hat{V}^{N_\omega} = \frac{1}{N_\omega} \sum_{k=1}^{N_\omega} \phi(X_T(\omega_k)).$$

Denote by $\Phi = \phi(X_T)$ the random pay-off, i.e. for event ω_k define $\Phi(\omega_k) = \phi(X_T(\omega_k))$ and

$$\Phi_{r,s} = \sum_{k=0}^s c_k \sum_{|u|=k} \phi(X_T^{u \cup \{1, \dots, r\}})$$

(compare with (23)), then we define the estimator

$$\widehat{V}_{r,s}^{N_\omega} = \frac{1}{N_\omega} \sum_{k=1}^{N_\omega} (\Phi(\omega_k) - \alpha (\Phi_{r,s}(\omega_k) - V_{r,s})), \quad (26)$$

where α will be determined later to achieve the best variance reduction, and $V_{r,s}$ given by (23) will be constructed from numerical PDE solutions. From

$$V = \mathbb{E}[\phi(X_T)], \quad V_{u \cup \{1, \dots, r\}} = \mathbb{E}[\phi(X_T^{u \cup \{1, \dots, r\}})],$$

we get

$$\mathbb{E}[\widehat{V}_{r,s}^{N_\omega}] = V,$$

i.e. the standard estimator and the control variate estimator are both unbiased.

Typically, we will find that $\alpha \approx 1$ is optimal. Indeed, if $\alpha = 1$, then

$$\widehat{V}_{r,s}^{N_\omega} = V_{r,s} + \frac{1}{N_\omega} \sum_{k=1}^{N_\omega} (\Phi(\omega_k) - \Phi_{r,s}(\omega_k)), \quad (27)$$

i.e. we sample the approximation error of the truncated decomposition. The motivation for choosing $\widehat{V}_{r,s}^{N_\omega}$ over the standard estimator \widehat{V}^{N_ω} is that for given ω_k , $\Phi_{r,s}$ and Φ will be close and therefore the variance is much reduced compared to the standard estimator.

The value of α which minimizes the variance is determined by the co-variances of the control variate and the standard estimator, which can be approximated by the estimators (see Glasserman 2004)

$$\begin{aligned} \widehat{\sigma^2} &= \frac{1}{N_\omega} \sum_{k=1}^{N_\omega} (\Phi_{r,s}(\omega_k) - V_{r,s})^2, \\ \widehat{\rho} &= \frac{1}{N_\omega} \sum_{k=1}^{N_\omega} (\Phi(\omega_k) - \widehat{V}^{N_\omega}) (\Phi_{r,s}(\omega_k) - V_{r,s}), \end{aligned}$$

and the optimal value is estimated as $\widehat{\alpha} = \widehat{\rho} / \widehat{\sigma^2}$. The variance reduction is then approximately,

$$\frac{\text{Var}[\Phi - \alpha (\Phi_{r,s} - V_{r,s})]}{\text{Var}[\Phi]} \approx 1 - \frac{\text{cov}(\Phi, \Phi_{r,s})}{\sqrt{\text{var}(\Phi)} \sqrt{\text{var}(\Phi_{r,s})}}.$$

Summarizing, in situations where the approximation $V_{r,s}$ from the previous section is not accurate enough, it can be corrected with a relatively small number of Monte Carlo samples of the correction terms (the right-hand sum) in (27).

3.3. Application to derivative portfolios

We now discuss approximations to expected exposures, as per (7). To benefit from the dimension reduction of risk-factor decomposition, we compute the solutions of all PDEs involved in the spirit of section 3.1.

The general principle is to replace the process X in (7) by X^v as defined in (14). Thus, we approximate V in (8) by,

$$\begin{aligned} V_u(a; x, t) &= \int_{\mathbb{R}} p_u(y, t; x, 0) \Psi \left(\sum_{k=1}^{N_D} \omega_k(t) V_{k,u}(\xi(t) \setminus y_u, t) \right) dy, \end{aligned} \quad (28)$$

where p_u is the transition density function of X^u , which satisfies the forward PDE (18), (19) instead of (9). Similarly, $V_{k,u}$ is given as the solution to (16), with pay-off function ϕ_k in (17). The key point is that we compute p_u and $V_{k,u}$ by solving $|u|$ -dimensional PDEs, and (28) is a $|u|$ -dimensional integration problem, as p is a Dirac measure in dimensions $\mathcal{D} \setminus u$.

Then $V_{r,s}$ is defined by (28) and (23). The complexity of the whole computation is linear in N_D , and exponential in $r + s$.

For products where model derivative prices are available in closed-form (such as most swaps), we will use those in the computations for speed and accuracy. For others (such as options), we will use numerical PDE approximations as outlined in appendix 2.

4. Case studies

In this section, we present the market set-up for the detailed numerical studies in section 5, where we compute expected exposures (4) and expected positive exposures (5) of portfolios with increasing complexity.

The example consists of various exchange rate and interest rate products, and is chosen to be representative for the scenario where each product in a portfolio is exposed to one or more of a pool of risk factors. So for instance an exchange option of medium term maturity is exposed to interest rate risk in both the domestic and foreign markets, and conversely the domestic interest rate affects an exchange option as well as, say, a domestic zero coupon bond.

4.1. Driving risk factors

We consider a portfolio of FX and interest rate products. Each FX rate is assumed to be governed by a full BS2HW model (see, e.g. Clark 2011). Every FX rate F_t^i , thus, has a stochastic domestic and foreign interest rate. For clarity of exposition, we consider the case where all exchange rates are relative to a single currency, in the case below the EUR, so that we can write the joint dynamics of FX and interest rates as,

$$dF_t^i = (R_t^d - R_t^{f,i}) F_t^i dt + \sigma^i(t) F_t^i dW_t^{F,i}, \quad 1 \leq i \leq m, \quad (29a)$$

$$dR_t^d = \lambda_d(\Theta_d(t) - R_t^d) dt + \eta_d dW_t^d, \quad (29b)$$

$$\begin{aligned} dR_t^{f,i} &= \left[\lambda_f^i(\Theta_f^i(t) - R_t^{f,i}) - \eta_f^i \rho_{I(i), J(i)} \sigma^i(t) \right] dt \\ &\quad + \eta_f^i dW_t^{f,i}, \quad 1 \leq i \leq m, \end{aligned} \quad (29c)$$

with given Brownian motions $(W^i)_{1 \leq i \leq 2m+1} = (W^{F,1}, W^d, W^{f,1}, W^{f,2}, W^{f,3}, W^{f,3})$, with

$$d[W^i, W^j]_t = \rho_{i,j} dt, \quad \text{for } 1 \leq i, j \leq 2m+1, \quad (30)$$

and $I(i)$ is the index of the i -th exchange rate and $J(i)$ the index of the i -th foreign rate (e.g. $I(1) = 1, J(1) = 3$). Here, R^d is the domestic short rate and $R^{f,i}$ the foreign short rate for exchange rate $F^i, i = 1, \dots, m$, with $m+1$ the total number of markets.

The interest rates are modelled by a mean-reverting Hull–White process with $\Theta_d(t)$ and $\Theta_f^i(t), 1 \leq i \leq m$, designed to fit the forward rate curve in the respective markets. More details on the calibration will be given in Section 4.3.

We defer the discussion of stochastic volatility models to section 5.5, but remark here that it is generally not straightforward to extend these models from a single currency pair to multiple pairs while preserving important FX characteristics like symmetry and the triangle inequality (see Doust 2012, De Col *et al.* 2013).

In the numerical examples, we choose EURUSD, EURGBP and EURJPY, i.e. $m = 3$. Together with piecewise constant volatility, the interest rates in the EUR (domestic) and USD, GBP and JPY (all foreign) markets, this gives $d = 2m + 1 = 7$ risk factors. In the case of stochastic volatility, EURUSD, EURGBP and EURJPY are all modelled with a Heston model such that there are 10 risk factors. This example is complex enough in the sense that a 7- or 10-dimensional PDE solution is not feasible and that we can demonstrate the effect of different decompositions.

4.2. Derivatives

We consider portfolios of cross-currency and interest rate swaps and FX options.

4.2.1. Zero-coupon bonds. Although the portfolios below do not contain any bonds explicitly, the swaps studied here are all based on coupon legs which can be replicated by (and hence valued from the time 0 value of) a string of zero-coupon bonds that pay 1 at different times T . The time t value a zero-coupon bond, $P(t, T)$, under Hull–White for the short rate R is,

$$P(t, T) = \frac{P(0, T)}{P(0, t)} e^{A(t, T) - B(t, T)R_t}, \quad (31)$$

where

$$A(t, T) = B(t, T)f(t) - \frac{\eta^2}{4\lambda} B^2(t, T)(1 - e^{-2\lambda t}),$$

$$B(t, T) = \frac{1 - e^{-\lambda(T-t)}}{\lambda},$$

and

$$f(t) = \frac{-\log(P(0, t))}{t}, \quad t \geq 0, \quad (32)$$

see Filipovic (2009).

For simplicity, we assume here a single curve framework and do not make a distinction between discounting and forwarding curves (see the discussion at the end of the introduction).

4.2.2. Cross-currency swap. We consider a series of FX forward swaps settled in arrears, with floating payment and receiving leg. Note that both legs need to be valued in the domestic currency, for which the future exchange rate is used. Following Filipovic (2009), the value $V_{\text{ccy}}(t, T_0, T)$ of the swap at time $t \leq T_0$ with payment dates $T_1, \dots, T_{N_C} = T$ is,

$$V_{\text{ccy}}(F_t, t; T_0, T) = C^d(t, T_0, T) - M F_t C^f(t, T_0, T),$$

where F_t is the FX rate at time t , $C^d(t, T_0, T)$ and $C^f(t, T_0, T)$ are the domestic and foreign floating rate notes, and M is the moneyiness as a percentage of the (future) FX rate, e.g. 100% is referred to as At-The-Money (ATM), 105% as In-The-Money

Table 1. Parameters for the CCYS and IRS.

		EURUSD	EURGBP	EURJPY		IR swap
Moneyiness	M	100%	95%	105%	M	100%
Maturity	T	5Y	3Y	2Y	T	5Y
Notional	Λ^d	100	100	50	Λ	150
Number of coupons	N_C	100	100	100	N_C	100

(ITM) and 95% as Out of-The-Money (OTM). The values of the floating rate notes are,

$$C^{\text{d/f}}(F_t, t; T_0, T) = P^{\text{d/f}}(t, T) \Lambda^{\text{d/f}} + \sum_{i=1}^{N_C} \left(P^{\text{d/f}}(t, T_{i-1}) - P^{\text{d/f}}(t, T_i) \right) \Lambda^{\text{d/f}} = P^{\text{d/f}}(t, T_0) \Lambda^{\text{d/f}},$$

where $\Lambda^{\text{d/f}}$ is the notional in domestic or foreign currency, $\Lambda^f = \Lambda^d / F_0$.

4.2.3. Interest rate swap. Here, we consider a floating vs. fixed interest rate swap, where the fixed leg is defined as a fixed coupon bond, such that the value with notional Λ equals

$$V_{\text{irs}}(t) = \left(K \Delta T \sum_{i=1}^{N_C} P(t, T_i) + P(t, T) - P(t, T_0) \right) \Lambda,$$

where K is the fixed rate and $\Delta T = T_i - T_{i-1}$ the interval between payment dates (see Filipovic (2009)).

4.2.4. FX option. Finally, we consider an FX call option on the EURUSD FX rate with strike K , notional Λ_{Opt} and maturity T with pay-off

$$\phi(F_T) = \max(0, F_T - K) \Lambda_{\text{Opt}}.$$

The value of the option at time $t < T$, under the BS2HW model from section 4.1, is equal to

$$V_{\text{opt}}(F_t, R_t^d, R_t^f, t) = \mathbb{E}[\exp(-\int_t^T R_s^d ds) \phi(F_T) | X_t],$$

where $X_t = (F_t, R_t^d, R_t^f)$ follow (29). For more details on these options we refer to Filipovic (2009) and Brigo and Mercurio (2013). We only note that the option value function is the solution of the corresponding three-dimensional backward PDE; see also remark 2.

4.2.5. Contract parameters. The CCYS are driven by the FX rate and the foreign and domestic interest rate processes. The contract specific parameters are shown in table 1. Note that the trades have different maturity, notional and moneyiness.

The interest rate swap is traded on the EUR interest rate, ATM, where the swap rate is such that the initial value of the trade at inception ($t = 0$) is equal to zero. The specific parameters can be found in table 1.

The option strike is set at $K = 0.95 F_0^1$ and the maturity to $T = 4$. The notional of this option is set to $\Lambda_{\text{Opt}} = 100$.

4.3. Market and calibration

As the behaviour and accuracy of the approximation method from section 3 may depend significantly on the model parameters, we perform a careful calibration to market data before undertaking numerical tests.

We use a data-set from 2 December 2014. At this time, the interest rates were low, which resulted in low yields for all markets. Shown in figure 1(b) are the yields up to five years, computed by polynomial interpolation of the market forward curve $f(t)$ as per (32). They are used for bond pricing via (31) and to calibrate $\Theta \in \{\Theta_d(t), \Theta_f^i(t), 1 \leq i \leq m\}$ in (29) by

$$\Theta(t) = \frac{1}{\lambda} \frac{df(t)}{dt} + f(t) + \frac{\eta^2}{2\lambda^2} (1 - e^{-2\lambda t}) \quad (33)$$

with $f(t)$ from (32), which gives an exact fit to bond prices (see Filipovic (2009)).

For the FX volatilities in (29a), we assume a piecewise constant volatility function which follows the ATM volatilities over time,

$$\sigma(t) = \sum_{i=1}^{N_M} \sigma_i 1_{\{T_{i-1} < t \leq T_i\}},$$

where $0 = T_0 < \dots < T_{N_M} = T_{\max}$ is a partition of the interval $[0, T_{\max}]$ for some maximum maturity T_{\max} such that σ_i is the volatility corresponding to the time period $(T_{i-1}, T_i]$. For simplicity, this partition corresponds to the maturities of available quoted options, hence N_M is the number of different maturities.

The calibration matches σ_i to European At-The-Money (ATM) options' volatilities assuming constant rates.[†] Denote $\tilde{\sigma}^{T_i}$ the implied volatility of an ATM option with maturity T_i , then the σ_i can be determined inductively from the system of equations,

$$(\tilde{\sigma}^{T_i})^2 T_i = \sum_{j=1}^i \sigma_j^2 (T_j - T_{j-1}), \quad i = 1, \dots, N_M.$$

The resulting different volatility levels are shown in figure 1(a).

The Hull–White parameters are calibrated to fit the prevailing yield curve and swaption data using analytic expressions, i.e. Θ is calibrated by (33) and η and λ by a least-squares fit to co-terminal swaptions that terminate in 10 years (1×9 , 2×8 , 3×7 , 4×6 , 5×5 , 6×4 , 7×3 , 8×2 and 9×1), as these swaptions can also be used to replicate the CVA of a swap that matures in 10 years, as is shown in Sorensen and Bollier (1994).

The correlations between the observable factors (exchange and interest rates) are estimated using weekly historical time-series data from the preceding three-year period.

The estimated correlation matrix is regularized to be positive semi-definite by setting any negative eigenvalues in a singular value decomposition to zero (Rebonato 1999). The resulting full correlation matrix is shown in appendix 1 together with the other parameters. The yield curve data are provided in a spread sheet at de Graaf et al. (2016).

[†]Such a step-by-step calibration is industry practice in CCR modelling.

5. Results

In this section, we analyse the accuracy of the numerical approximations from section 3, applied to the market described in section 4. We compare our results to an accurate approximation computed by a brute-force Monte Carlo method. In particular, we study the relative differences expressed in percentages,

$$e_{L_2} = 100 \frac{\sqrt{\sum_{i=1}^{N_T} (V_{r,s}(t_i) - V_{MC}(t_i))^2}}{\sqrt{\sum_{i=1}^{N_T} V_{MC}(t_i)^2}},$$

$$e_{L_\infty} = 100 \frac{\max_{i=1}^{N_T} (|V_{r,s}(t_i) - V_{MC}(t_i)|)}{\max_{i=1}^{N_T} (|V_{MC}(t_i)|)}, \quad (34)$$

where $r = 1$ and $s = 1$ or $s = 2$ to denote two- or three-dimensional corrections (see (23)). For the computations, t_i are chosen to coincide with the swap payment dates, so that $N_T = N_C$ is the total number of coupon payments. In addition, we study the mean difference (MD) over time relative to the sum of all notionals (N_{total}) in the portfolio expressed in basis points,

$$\text{MD} = \frac{10000}{N_T N_{\text{total}}} \sum_{i=1}^{N_T} |V_{r,s}(t_i) - V_{MC}(t_i)|.$$

Similarly, the normalized standard error for EE is defined as,

$$SE = 100 \sqrt{\sum_{i=1}^{N_T} SE(t_i)^2} / \sqrt{\sum_{i=1}^{N_T} EE(t_i)^2},$$

$$SE(t_i) = \frac{\text{Std}(EE(t_i))}{\sqrt{N_\omega}}, \quad (35)$$

where $\text{Std}(EE(t_i))$ is the standard deviation of the estimator for $EE(t_i)$, and an analogous definition for EPE.

The settings used for the numerical methods (domain and mesh sizes, time steps, etc) are reported in appendix B.3.

In the following, assuming piecewise constant volatilities, we use up to seven risk factors, $d = 7$, where (see (1) and (29))

$$(X_t^1, \dots, X_t^7) = (F_t^1, R^d, R_t^{f,1}, F_t^2, R_t^{f,2}, F_t^3, R_t^{f,3}) \quad (36)$$

are the EURUSD, EURGBP, EURJPY exchange rates (F^i), and EUR, USD, GBP, JPY short rates (R^d and $R^{f,i}$, respectively).

To simplify the notation (28), we will use, e.g. the shorthand

$$V(R^d, F^2, R^{f,3}) \equiv V_{\{2,4,7\}}(R_0^d, F_0^2, R_0^{f,3}; R_0^d, F_0^2, R_0^{f,3}, t)$$

$$= V_{\{2,4,7\}}(X_0^2, X_0^4, X_0^7; X_0^2, X_0^4, X_0^7, t),$$

where we suppress the dependence on the anchor $a = X_0$ and time t , and it is understood implicitly that the arguments identify the function being used. Also note that the function is evaluated at X_0 , since we are interested in exposures conditional on the current state \mathcal{F}_0 .

5.1. Case A: single EURUSD CCYS

As a first test case, we focus on the BS2HW model (section 4.1) for a single CCY swap trade on EURUSD. In this case, only the three relevant risk factors F^1 , R^d and $R^{f,1}$ (the EURUSD exchange rate, the EUR and USD short rates) are used, so effectively $d = 3$.

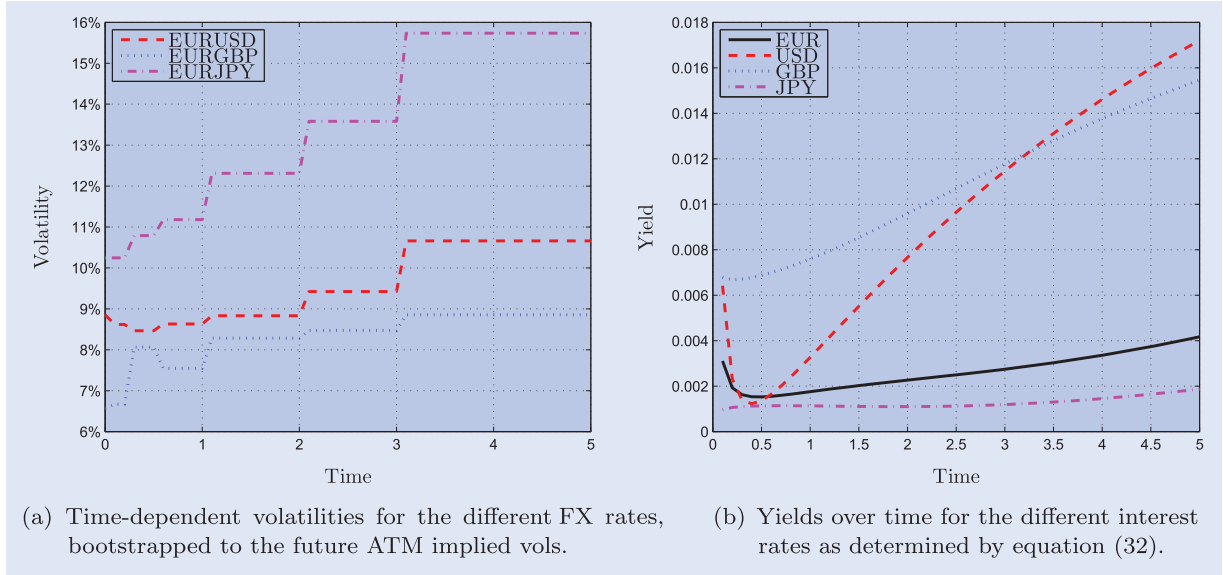


Figure 1. Volatilities and yields over time on 2 December 2014.

In that case

$$V_{1,1} = V(F^1) + \left(V(F^1, R^d) - V(F^1) \right) + \left(V(F^1, R^{f,1}) - V(F^1) \right) \quad (37)$$

$$V_{1,2} = V_{1,1} + \left(V(F^1, R^d, R^{f,1}) - V(F^1, R^d) - V(F^1, R^{f,1}) + V(F^1) \right) \quad (38)$$

$$= V(F^1, R^d, R^{f,1}). \quad (39)$$

The first term in (37) is a one-dimensional approximation with only the EURUSD rate stochastic, and then the next two brackets correct separately for stochastic domestic (EUR) and stochastic foreign (USD) interest rate, respectively. Note how in this particular case equation (38) with the three-dimensional corrections simplifies to the exact solution (39).

In figure 2(a) and (b), the one-, two- and three-dimensional decomposed approximations $V_{1,0}$, $V_{1,1}$ and $V_{1,2}$ are plotted together with the full-scale Monte Carlo approximation V_{MC} . The exposure decreases over time, which means that the FX rate drives the expected value of this swap down. The EPE is positive by definition and increases over time.

Both the decomposed two- and three-dimensional approximations are reasonably close to the full Monte Carlo estimator, but while the results for Case A in table 2 show that the maximum error is over 5% in the two-dimensional case, it is around 0.5% for the three-dimensional approximation. Note that the latter result should be exact, so that any discrepancy are finite difference errors. These may seem large, but we note that the standard error of Monte Carlo with the industry standard of 10^4 paths would be approximately $0.1\% \sqrt{4 \cdot 10^6 / 10^4} = 2\%$ (see table 2). While we could certainly improve the finite difference accuracy by finer meshes, we do not investigate this further at this point.[†]

[†]See appendix 4 for an analysis of the accuracy of the finite difference schemes used.

5.2. Case B: three CCYS (EURUSD, EURGBP and EURJPY)

In the case of three CCYS, there are $d = 7$ risk factors, such that adding three-dimensional corrections will no longer yield an exact solution. In this case, we make an *a priori* choice of the base risk factor, namely the EURUSD FX rate.[‡] The two-dimensional decomposed approximation for the portfolio of three FX swaps now has two extra terms to correct for the other FX rates and two extra terms to correct for the foreign interest rates:

$$\begin{aligned} V_{1,1} &= V(F^1) + \left(V(F^1, F^2) - V(F^1) \right) \\ &+ \left(V(F^1, F^3) - V(F^1) \right) + \left(V(F^1, R^d) - V(F^1) \right) \\ &+ \left(V(F^1, R^{f,1}) - V(F^1) \right) \\ &+ \left(V(F^1, R^{f,2}) - V(F^1) \right) \\ &+ \left(V(F^1, R^{f,3}) - V(F^1) \right) \\ &= V(F^1) + \sum_{j=2,3} \Delta V(F^1, F^j) + \Delta V(F^1, R^d) \\ &+ \sum_{j=1,2,3} \Delta V(F^1, R^{f,j}), \end{aligned} \quad (40)$$

where the short-hand $\Delta V(X^1, X^j) \equiv V(X^1, X^j) - V(X^1)$ was used in the last line.

For the three-dimensional corrections, there are in principle $\binom{6}{2} = 15$ extra terms, because in addition to X_1 , we choose two out of the $d - 1 = 6$ remaining factors. However, only eight of them are non-zero, as given here:

$$\begin{aligned} V_{1,2} &= V_{1,1} + \Delta V(F^1, F^2, F^3) + \Delta V(F^1, F^2, R^d) \\ &+ \Delta V(F^1, F^3, R^d) \\ &+ \sum_{j=1,2} \Delta V(F^1, F^2, R^{f,j}) \end{aligned}$$

[‡]See section 5.7 for an assessment of alternative choices.

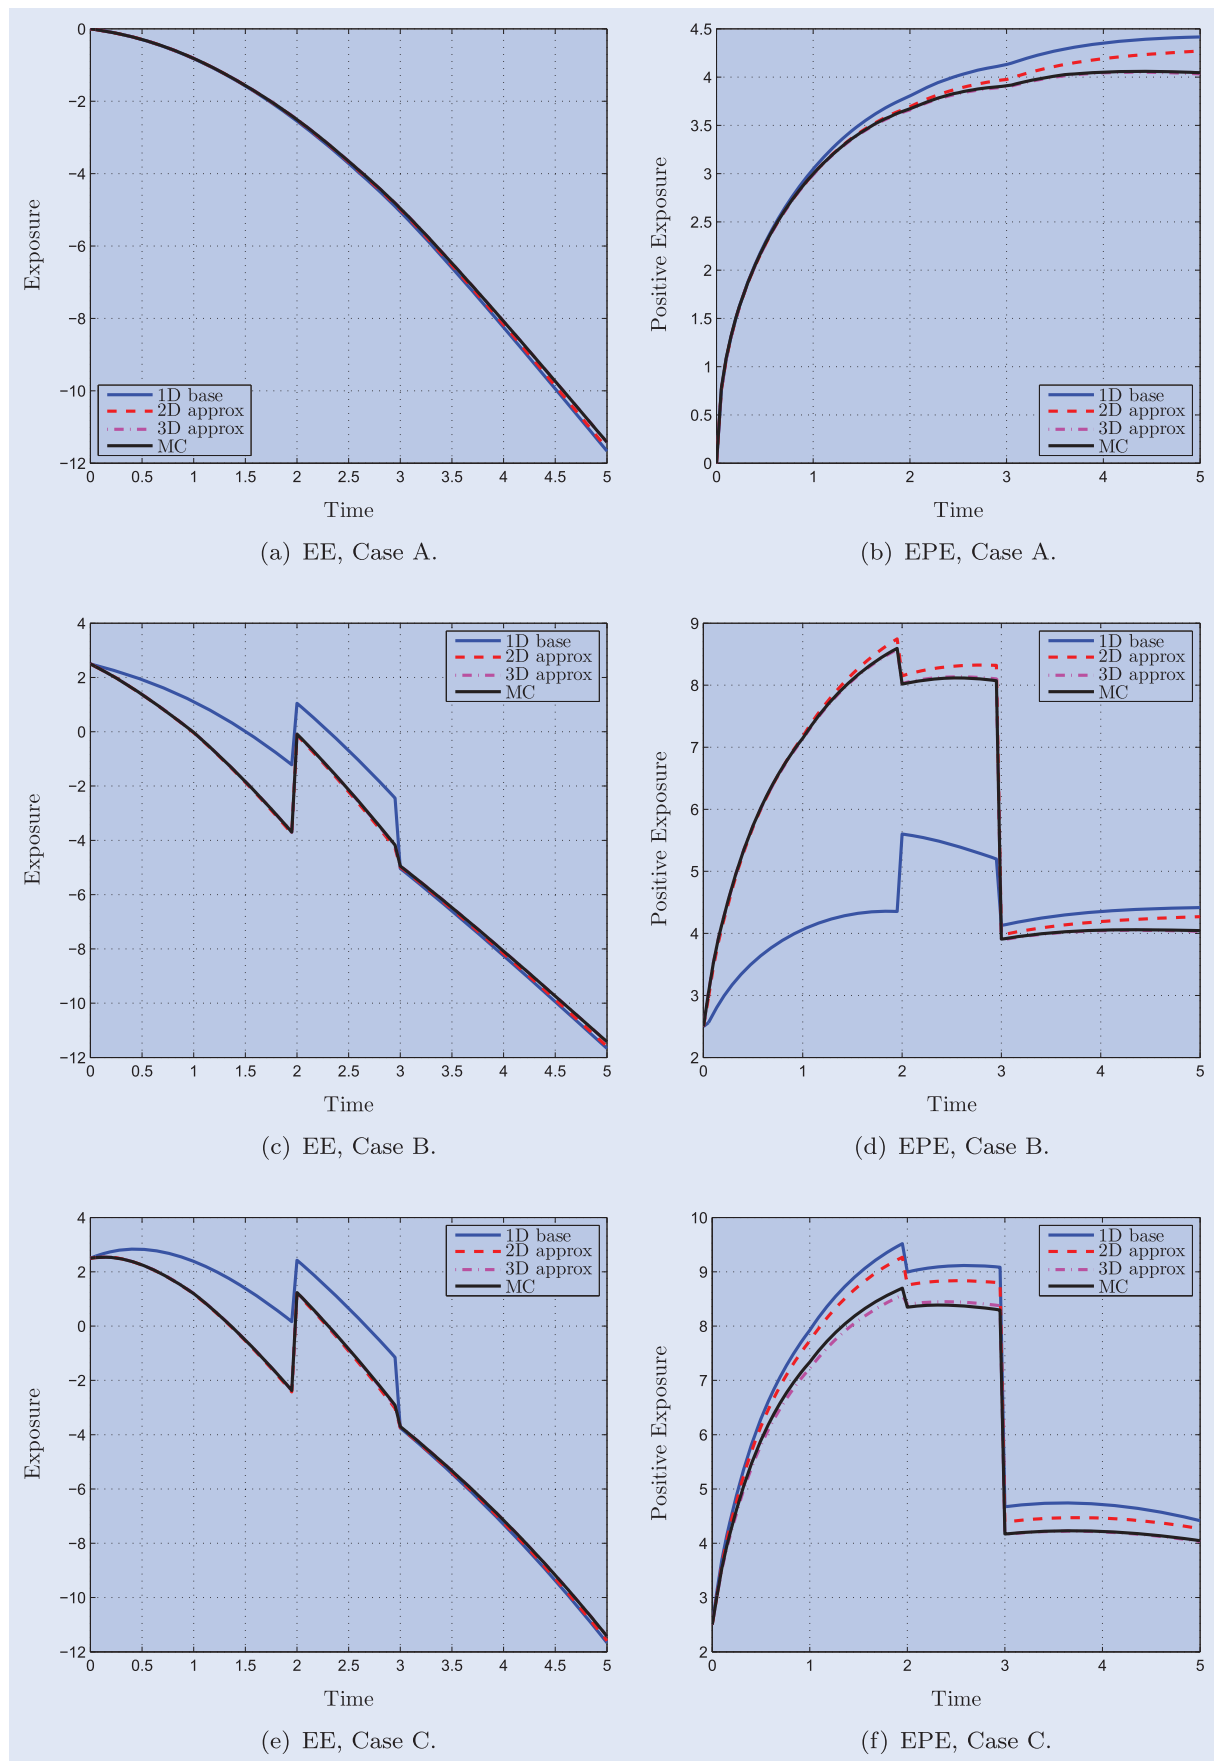


Figure 2. EE (left) and EPE (right) for a single ATM CCYS (2(a) and (b)), for three CCYS (2(c) and (d)), and for three CCYS and an IR swap (2(e) and (f)). The risk factors are driven by BS2HW and the EURUSD FX rate is taken as a base.

Table 2. Errors of exposure approximations for Cases A, B and C. The finite difference approximations are computed with $m_1 = 60$ grid points and 500 time steps. The errors are expressed in percentages together with the standard error of the Monte Carlo benchmark with $4 \cdot 10^6$ paths and 1000 time steps. The SE is defined in (35) as the root sum squared of the standard errors over time relative to the root sum squared of the sampled EE or EPE.

		$e_{L_2} (\%)$			$e_{L_\infty} (\%)$			MD (bp)			SE (%)
		1D	2D	3D	1D	2D	3D	1D	2D	3D	
Case A	EE	2.07	1.20	0.038	2.16	1.50	0.049	8.58	4.43	0.17	0.13
	EPE	6.11	2.83	0.14	9.27	5.66	0.55	18.29	7.35	0.40	0.085
Case B	EE	19.54	1.41	0.20	21.67	1.50	0.34	32.46	2.52	0.25	0.17
	EPE	36.51	2.29	0.26	49.44	3.09	0.54	359.1	4.93	0.46	0.075
Case C	EE	22.23	1.48	0.18	22.19	1.44	0.30	33.26	2.36	0.22	0.19
	EPE	29.83	5.63	1.17	41.87	6.85	1.57	324.7	13.24	2.32	0.073

$$+ \sum_{j=1,3} \Delta V(F^1, F^3, R^{f,j}) + \Delta V(F^1, R^d, R^{f,1}), \quad (41)$$

where the short-hand $\Delta V(X^1, X^i, X^j) \equiv V(X^1, X^i, X^j) - V(X^1, X^i) - V(X^1, X^j) + V(X^1)$ was used.

The corrections can be interpreted in the following way. For instance, in the first term in (41), $V(F^1, F^2, F^3)$ treats all exchange rates stochastic but with deterministic interest rates (simple Black–Scholes models); and in the last term, $V(F^1, R^d, R^{f,1})$ comes from the full BS2HW model for the EURUSD rate where all other processes are approximated by their expectations (see section 3.1). Similar interpretations can be found for the other correction terms.

In contrast, we have for instance

$$\begin{aligned} \Delta V(F^1, R^d, R^{f,3}) &= \underbrace{V(F^1, R^d, R^{f,3}) - V(F^1, R^d)}_{=0} \\ &\quad - \underbrace{V(F^1, R^{f,3}) + V(F^1)}_{=0} = 0, \end{aligned}$$

because in this approximation where (the EURJPY exchange rate) F^3 is deterministic, the (JPY) short rate $R^{f,3}$ has no impact on the exposure of the swaps. Similar arguments hold for the other six corrections that vanish.

Figure 2(c) and (d) show again the one- to three-dimensional decomposed approximations of the exposures together with the full scale Monte Carlo approximation V_{MC} . The different expiries for the different swaps are reflected clearly in the profiles. At $T = 2$, the EURJPY swap expires, which results in an upwards jump in total EE and EPE whereas at $T = 3$ the EURGBP swap expires and the exposures drop.

Both the two- and especially the three-dimensional approximation are close to the full Monte Carlo estimator; see especially the rows relating to Case B in table 2. The mean differences are now relative to the sum of all notionals in the portfolio. As none of the correction terms accounts simultaneously for, say, the EURJPY rate, EUR and JPY short rates, even the decomposition with three-dimensional approximations is not exact in this case, such that all errors arise from a combination of truncation error for the expansion and discretization error for the PDE solution (as well as a slightly smaller Monte Carlo error for the benchmark).

We give a table of all the correction terms separately in appendix 4. Table D4 there shows while some terms are dominant, like the one involving the EURUSD and EURGBP rates among the two-dimensional corrections and the one with EURUSD

and the associated short rates among the three-dimensional ones, it is certainly not the case that the other terms are negligible. This demonstrates that the problem is genuinely high-dimensional and does not have a lower superposition dimension (see Wang and Sloan 2005), by which we mean loosely speaking that the solution cannot be expressed exactly as a linear combination of solutions to low-dimensional problems. Indeed, while the approximations $V_{0,3}$ and hence $V_{1,3}$ would be exact for EE here because of its linearity, $V_{1,2}$ is not. For EPE, generally only decompositions involving all risk factors, such as $V_{0,7}$ or any $V_{r,s}$ with $r+s = 7$ will be exact. The results show, however, that approximations with much smaller r and s and therefore, much lower computational cost can be sufficiently accurate.

In section 5.6 we will show the variance reduction achieved in Case B using the above approximation as control variate.

5.3. Case C: Case B with an additional IRS in EUR

By adding an IR swap to the portfolio, the risk factors do not change, as the EUR interest rate was already modelled as a factor that drives the FX rate in Section 5.2. In figure 2(e) and (f), one observes the exposure levels are increased because of the interest rate swap. Again, the different profiles match closely over time. In table 2 one sees more clearly that, for all error measures, both the EE and EPE errors improve significantly when we include three-dimensional corrections, as expected.

One thing to note is that the errors for EE are almost exactly the same as in Case B, because the IR swap exposure is valued exactly by all models that have R^d in them, e.g. $V(F^1, R^d)$.

5.4. Case D: Case C with an additional EURUSD FX call option

We now add an FX option on the EURUSD rate as described in section 4.2. The exposure level increases first, but drops at $t = 4$, as seen in figure 3(a) and (b).

In the previous examples, sections 5.1–5.3, analytical pricing formulas were available for the swaps in the portfolio, which made Monte Carlo estimation of exposures straightforward. With the FX option added, we use the regression-based MC method, similar to the popular Longstaff–Schwartz algorithm (Longstaff and Schwartz 2001) for American options. An outline of the algorithm used is given in appendix 3.

Because of this regression, all paths need to be stored and less paths can be used, which results in a higher standard

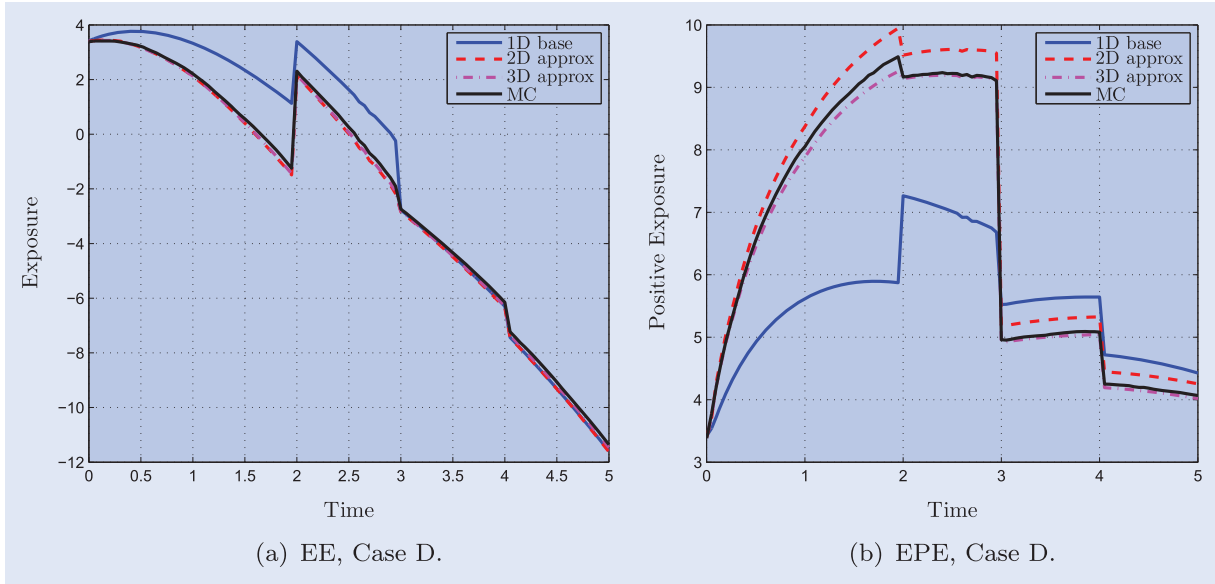


Figure 3. EE (left) and EPE (right) for a portfolio of CCYS, a call option on the EURUSD FX rate and a floating vs. fixed interest rate swap. The factors are driven by BS2HW and the EURUSD FX rate is taken as the base factor.

error. We also note that the benchmark in this case is not perfect due to the limited span of the basis functions. The resulting bias is not taken into account in the SE given in table 3. Nonetheless, the discrepancy between PDE results and Monte Carlo behaves as expected—see table 3—with a sharp improvement as corrections are added in.

5.5. Stochastic volatility

To investigate the applicability of this method to stochastic volatility, we model all three FX rates with the Heston model (Heston 1993), where the stochastic volatility is modelled by a CIR process:

$$dF_t^i = (R_t^d - R_t^{f,i})F_t^i dt + \sqrt{Y_t^i} F_t^i dW_t^{F,i}, \quad 1 \leq i \leq m, \quad (42a)$$

$$dY_t^i = \kappa_i (\bar{v}_i - Y_t^i) dt + \gamma_i \sqrt{Y_t^i} dW_t^{Y,i}, \quad 1 \leq i \leq m, \quad (42b)$$

$$dR_t^d = \lambda_d (\Theta_d(t) - R_t^d) dt + \eta_d dW_t^d, \quad (42c)$$

$$dR_t^{f,i} = \left[\lambda_f^i (\Theta_f^i(t) - R_t^{f,i}) - \eta_f^i \rho_{I(i),J(i)} \sqrt{Y_t^i} \right] dt + \eta_f^i dW_t^{f,i}, \quad 1 \leq i \leq m, \quad (42d)$$

with $Y_0^i = v_{0,i}$ given, and other notation as earlier. The three sets of Heston parameters are calibrated to volatility smiles from 2 December 2014 and are given with the calibration fit in appendix A.1. It is unclear how to estimate the correlations between the stochastic variance and the other risk factors (apart from its own FX rate). In this test, they are assumed to be zero.† The extension to local-stochastic volatility models would not cause any computational difficulty and the accuracy of the

approximation is expected to be similar. In this case, there are $d = 10$ risk factors and we again choose the EURUSD FX rate *a priori* as base risk factor. We consider Case B with three CCYS with different moneyness and maturities.

The two-dimensional decomposed approximation for the portfolio of three FX swaps now has one extra term to correct for the stochastic volatility of the EURUSD base risk factor:

$$V_{1,1} = V(F^1) + \sum_{j=2,3} \Delta V(F^1, F^j) + \Delta V(F^1, R^d) + \sum_{j=1,2,3} \Delta V(F^1, R^{f,j}) + \Delta V(F^1, Y^1),$$

where the short-hand notation from equation (40) was used.

For the three-dimensional corrections, in this case there are in principle $\binom{9}{2} = 36$ extra terms, because in addition to X^1 , we choose two out of the $d - 1 = 9$ remaining factors. In this preliminary test, however, we only add a stochastic volatility factor to the four 2D corrections that contribute the most in the Black–Scholes case.‡ These include corrections for the volatility in the two additional FX rates (EURGBP and EURJPY):

$$V_{1,2} = V_{1,1} + \Delta V(F^1, F^2, F^3) + \Delta V(F^1, F^2, R^d) + \Delta V(F^1, F^3, R^d) + \sum_{j=1,2} \Delta V(F^1, F^2, R^{f,j}) + \sum_{j=1,3} \Delta V(F^1, F^3, R^{f,j}) + \Delta V(F^1, R^d, R^{f,1}) + \sum_{j=1,2} \Delta V(F^1, F^j, Y^j) + \Delta V(F^1, Y^1, R^d) + \Delta V(F^1, Y^1, R^f), \quad (43)$$

where the short-hand notation from equation (41) was used.

Note that all the corrections without stochastic volatility treat all exchange rates stochastic but with a time dependent volatility function equal to the expectation of the future volatility,

†After correcting the estimated full correlation matrix for positive-definiteness, these correlations can be slightly non-zero. The full correlation matrix used in this test can be found in equation (A1) in appendix A.1.

‡See table D4 in appendix 4.

Table 3. Errors of exposures of a portfolio with three CCYS, a IRS and a FX call option. The errors are expressed in percentages together with the standard error of the Monte Carlo benchmark with $2 \cdot 10^5$ paths. The SE is defined in (35) as the root sum squared of the standard errors over time relative to the root sum squared of the sampled EE or EPE.

	$e_{L_2} (\%)$			$e_{L_\infty} (\%)$			MD (bp)			SE (%)
	1D	2D	3D	1D	2D	3D	1D	2D	3D	MC
EE	23.20	1.43	0.32	22.57	1.56	0.36	33.72	2.08	0.52	0.86
EPE	26.36	4.83	1.29	37.87	5.44	1.87	319.00	12.51	2.62	0.30

Table 4. Errors of exposure approximations for Case B modelled with the Heston model. The finite difference approximations are computed with $m_1 = 80$ grid points and 500 time steps. The errors are expressed in percentages together with the standard error of the Monte Carlo benchmark with $4 \cdot 10^6$ paths and 1000 time steps. The SE is defined in (35) as the root sum squared of the standard errors over time relative to the root sum squared of the sampled EE or EPE.

		$e_{L_2} (\%)$			$e_{L_\infty} (\%)$			MD (bp)			SE (%)
		1D	2D	3D	1D	2D	3D	1D	2D	3D	
Case B Heston	EE	23.40	0.62	0.68	28.03	0.87	0.75	42.79	1.09	1.35	0.18
	EPE	31.84	8.55	0.96	41.70	12.00	1.08	368.01	15.59	2.13	0.076

which can be calculated analytically as,

$$\mathbb{E} \left[Y_t^i | Y_0^i \right] = v_{0,i} e^{-\kappa_i t} + \bar{v}_i (1 - e^{-\kappa_i t}). \quad (44)$$

The one- to three-dimensional decomposed approximations of the exposures including stochastic volatility are shown in figure 4(a) and (b), together with the full-scale Monte Carlo approximation V_{MC} . Again, the profiles reflect the different expiries, and we see that the two-dimensional approximation significantly outperforms the one-dimensional approximation, and for EPE the three-dimensional approximation significantly outperforms the two-dimensional approximation.

This is also clearly visible in table 4, where we show various error measures as earlier for the different approximations.

In the case of expected exposure, it appears that the two-dimensional decomposed approximation is already close to the full-scale Monte Carlo solution. This is in line with Simaitis *et al.* (2016), where it is found that for expected exposure, a time-dependent volatility function can be used, whereas only for non-linear exposures (such as expected positive or negative exposure), a full stochastic volatility model should be used.

5.6. Variance reduction

We now correct the bias in the results in the preceding sections by applying the methodology of section 3.2, at much reduced variance compared to the standard estimator. For simplicity, we restrict ourselves to the case where derivatives are priced analytically, so that no regression is required. The obtained variance reduction, calculated as the ratio of the two variances with and without control variate, is shown in figure 5 for the different decomposed approximations and test cases from sections 5.1–5.3. To be clear, a value of, e.g. 0.5% means that the variance is reduced by a factor of 200, while a value of 100% implies that no variance reduction at all was achieved.

In figure 5(a) and (b), the variance reduction factor for EE and EPE in the case of a single CCYS is in the range of 100–1000. For the two-dimensional control variate, the reduction is greater for shorter maturities and increases with time. For

Table 5. Computational times of the individual solvers in seconds. For the finite difference (FD) solvers, we use $m_1 = 60$ grid points and 100 or 500 (in brackets) time steps (see also Appendix B.3). The full Monte Carlo (MC) result is obtained with $4 \cdot 10^6$ paths and 100 or 1000 (in brackets) time steps. For the control variate (CV MC) results, we reduce the number of paths commensurate with the variance reduction. From Figure 5, Case B, the variance reduction factor is seen to be around 50 in the case of 2D and 200 for the 3D corrections. Therefore, we choose $8 \cdot 10^4$ paths in the case of 2D corrections and $2 \cdot 10^4$ for the 3D corrections.

Solver	time (s)
1D FD	1.53 (3.65)
2D FD	2.35 (10.67)
3D FD	344 (601.34)
Full MC	4251 (7738)
2D CV MC	81.5 (258.4)
3D CV MC	30.1 (107.6)

the three-dimensional control variate, the variance reduction is constant, which is due to the exactness of the control variate. The only discrepancy between the control variate and the Monte Carlo estimator is the (time) discretization error, which also explains why the variance is not equal to zero.

Figure 5(a) and (b) show the variance reduction in the case of three CCYS. Here, the three-dimensional decomposed approximation is no longer exact and thus, also in the three-dimensional case, the variance increases over time. The sharp drop at $T = 2$ and $T = 3$ is due to the expiry of the EUR-JPY and EURGBP CCYS. The variance reduction is largest between time 3 and 5 because during that time period only the EURUSD CCYS is not terminated. When we use two-dimensional corrections, the reduction is around 200 for EE and 50 for EPE. When three-dimensional corrections are taken into account, we obtain a reduction by a factor of 10^5 for EE and 200 for EPE.†

†The extreme variance reduction for EE and 3D is more clearly seen quantitatively from figure 6 later on, where the data are plotted on a log scale.

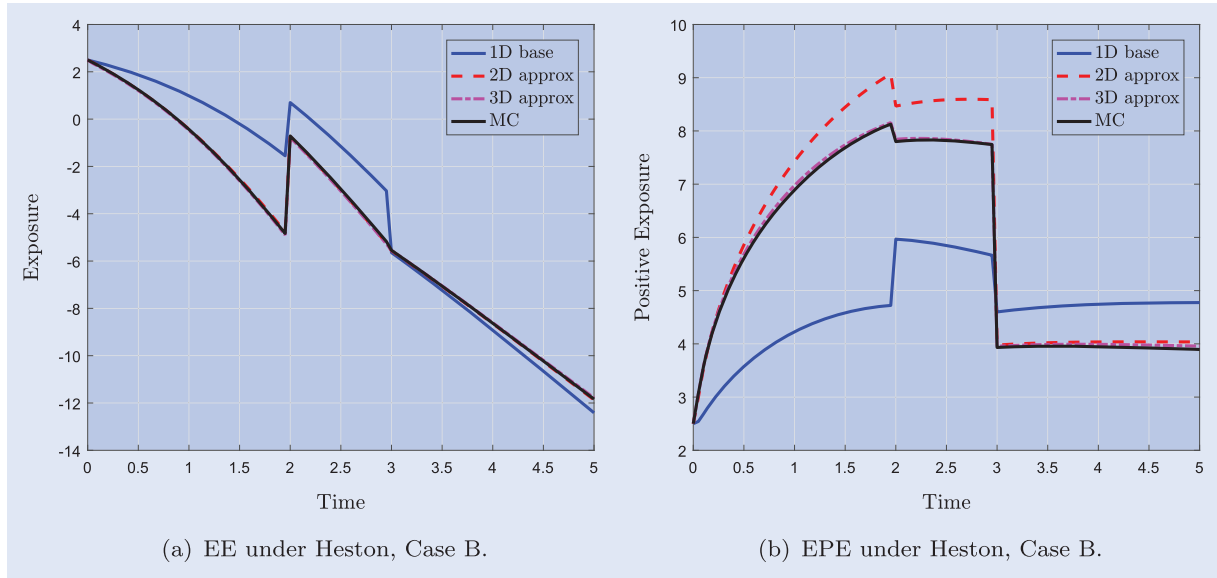


Figure 4. EE (left) and EPE (right) for three CCYS under Heston. The EURUSD FX rate is taken as a base risk factor and all FX rates are simulated with stochastic volatility.

Table 6. Errors of exposures for Case B for different choices of the base risk factor. The errors are expressed in percentages together with the standard error of the Monte Carlo benchmark with 10^6 paths. The SE is defined in (35) as the root sum squared of the standard errors over time relative to the root sum squared of the sampled EE or EPE.

		e_{L_2}		e_{L_∞}		MD (bp)		SE
		2D	3D	2D	3D	2D	3D	
EURUSD	EE	2.02%	0.69%	1.83%	0.59%	2.43	0.86	0.23%
	EPE	1.97%	0.36%	2.74%	0.51%	2.64	0.51	0.15%
EURGBP	EE	2.32%	1.39%	2.26%	1.60%	2.47	1.26	0.23%
	EPE	4.01%	1.89%	4.77%	3.05%	5.45	2.02	0.15%
EURJPY	EE	2.44%	1.43%	2.26%	1.60%	2.72	1.39	0.23%
	EPE	7.91%	2.03%	19.35%	3.05%	8.95	2.46	0.15%

Figure 5(c) and (d) show the variance reduction in the case of three CCYS and an IRS. The EE results resemble those for three CCYS alone, because the stochasticity of the IRS due to R^d is modelled exactly by the corrections, as explained at the end of section 5.3. For EPE, the variance is reduced by a factor around 30 for two-dimensional corrections and a factor of 50 for three-dimensional corrections.

The set-up is chosen so that the finite difference accuracy is comparable to the Monte Carlo accuracy. As the finite difference approximation converges at a higher order (for low dimensions, typically up to about 3 or 4) than the Monte Carlo sampling, the computational effort of the Monte Carlo component will become dominant if higher accuracy is required. As a side note, we remark here that the different PDEs and integration problems can be solved by entirely different methods, e.g. Fourier methods could be used for some of the lower dimensional problems or quasi-Monte Carlo methods for some of the medium-dimensional problems.[†] The different solutions in the decomposition can also be computed fully in parallel,

[†]Similarly, quasi-Monte Carlo could be used in the benchmark method and offers a potential reduction in simulation paths from

and parallel to the Monte Carlo runs. We report here sequential run times.

From these run times for Case B, presented in table 5, we can see that, using the 2D approximation as a control variate, we gain a computational speed-up by a factor 50, as the computation time for the 2D corrections by finite differences is negligible. Using 3D corrections gives us a speed-up factor 12; in spite of the greater variance reduction, the increased time for the PDE solutions eats up any benefit over the 2D corrections.

5.7. Other base risk factors

In the previous sections, the EURUSD FX rate was chosen as the base risk factor *a priori*, based on the fact that the derivatives driven by this FX rate have the highest maturity. However, there is no guarantee that this choice as a base is optimal in any sense. Often, one may have prior knowledge what the main driving factors are. Failing that, it would be

$O(n^4)$ to $O(n^2)$ for simulation-in-simulation (which is rare), or from $O(n^2)$ to $O(n)$ for regression methods.

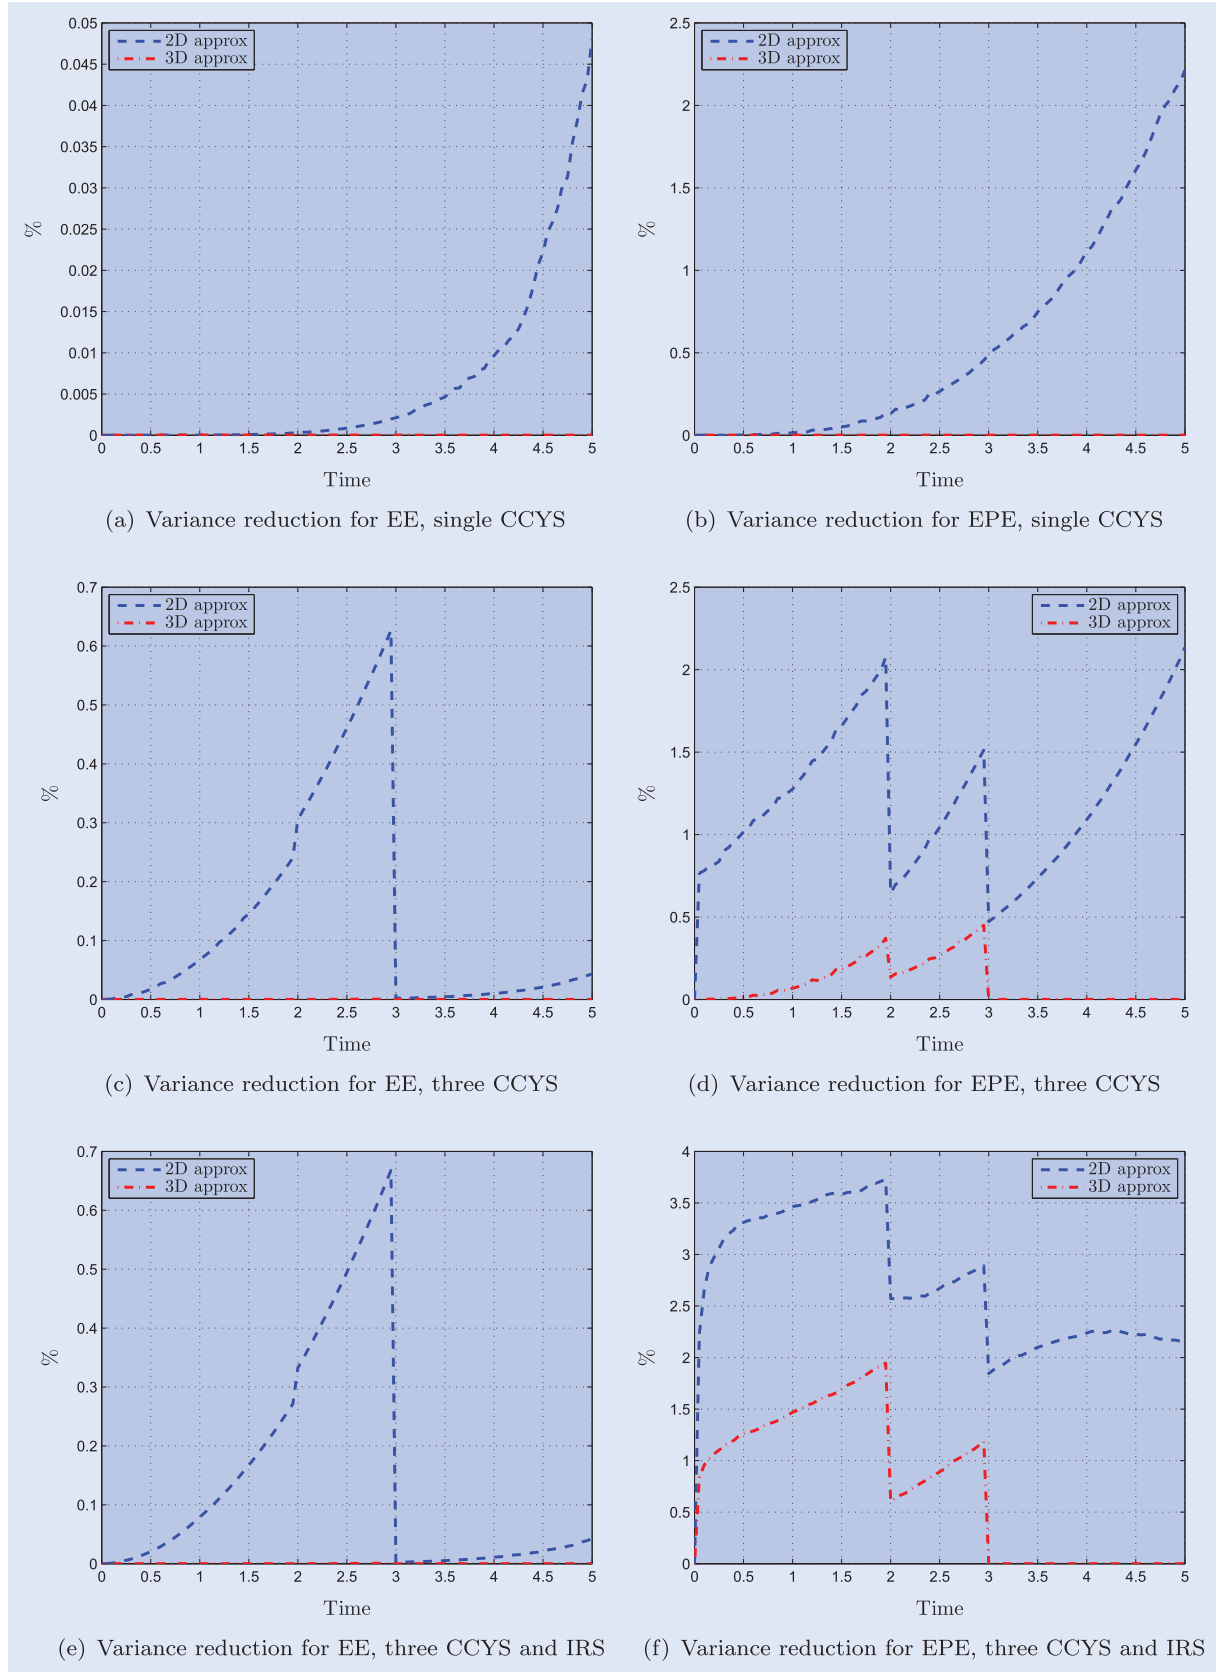


Figure 5. Variance reduction for different test cases with 2D and 3D corrections. The variance with and without control variate is compared. For this computation 10^6 paths were used.

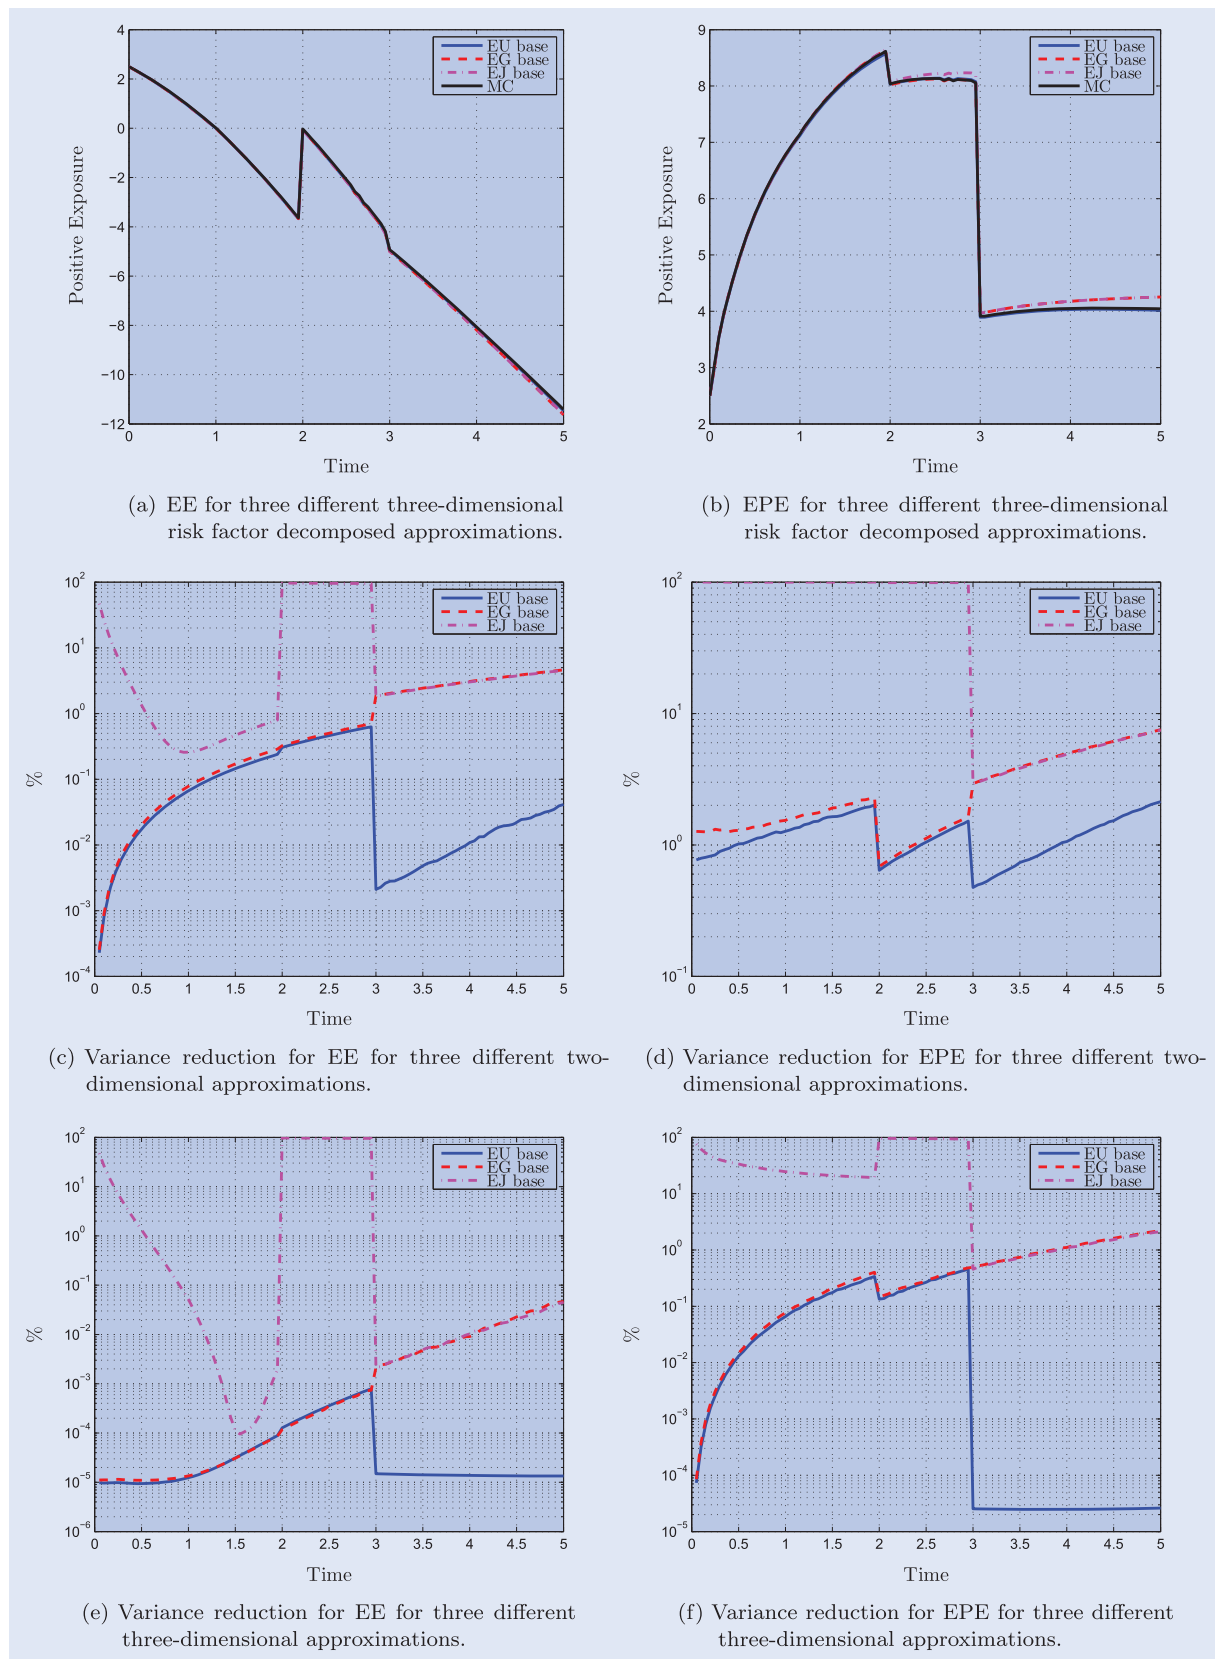


Figure 6. Exposure profiles (EE and EPE) for Case B and variance reduction with different base factors and 2D and 3D corrections. For the variance computations 10^6 paths were used.

practically feasible to estimate the variance reduction achieved by different factors by a relatively small number of samples in a trial run, and then do the actual large-scale estimation with the best performing base factor. In figure 6(a) and (b), the different EE and EPE profiles for Case B are shown when we choose the EURGBP or EURJPY FX rate as a base compared to choosing the EURUSD rate. Table 6 shows that the error is indeed smallest for the EURUSD rate.

In figures 5(e) and (f), we show the resulting variance reductions of the corresponding control variates. Clearly, the EURJPY FX rate as a base does not reduce the variance after $T = 2$, as that is when the base risk factor does not affect any non-terminated derivative in the portfolio. The EURUSD FX rate performs best especially after time $T > 3$ when the EURUSD rate affects the only non-terminated derivative in the portfolio. Note that we use a log scale for better visualization.

6. Conclusion

This paper is motivated by the computation of exposure profiles for portfolios depending on a moderate to large number of risk factors. For this problem, the industry standard technique is to employ forward Monte Carlo sampling to compute future scenarios. This scales linearly in both the number of dimensions (risk factors) and products. However, valuing the whole book across all scenarios is still a big computational challenge and relatively large standard errors have to be tolerated given the relatively low feasible number of sample paths.

We therefore, propose another approach which exploits the accuracy of PDE approximation schemes for low-dimensional estimation problems, through an anchored-ANOVA-style splitting of the high-dimensional problem into a sequence of lower dimensional ones. As the problem is stated in the form of nested conditional expectations, we use a combination of forward and backward PDEs to generate exposure profiles for all future times with a single backward and forward sweep.

This paper provides a proof of concept rather than a fully worked out black box algorithm. The detailed analysis of a moderately sized, realistic and fully calibrated test case showed the scale of the benefits achievable by risk factor decomposition coupled with numerical PDE solutions, used as standalone approximation or—if necessary—employed as a control variate for a Monte Carlo estimator. Some of the computational savings were dramatic, with a speed-up factor of 50 compared to standard Monte Carlo estimation, using a sum of two-dimensional estimators as control variate.

Further tests, not reported here, indicate that for longer maturities (10 years) some three-dimensional terms in the decomposition become significant for an accurate standalone approximation as well as for effective variance reduction.

It remains to investigate in practice the scalability with respect to the number of risk factors and number of products in the derivative portfolio. The PDE-based method, like standard Monte Carlo sampling, scales linearly in the number of derivatives. As for the number of risk factors, the number of 2D correction terms is linear, and the number of 3D ones quadratic in the dimension.

For larger portfolios, one can consider using approximations of the form $V_{0,2}$ or $V_{0,3}$ instead of $V_{1,2}$, or indeed a combination

of the two in the sense that only a subset of V_u with $|u| = 2$ and $|u| = 3$ are computed and included in the approximation, chosen adaptively by estimation of the individual terms. We find evidence for this in table D4, which shows that only a subset of the correction terms are significant. This also suggests that not all terms have to be computed with the same relative accuracy, and one can follow the principle of Griebel and Holtz (2010) to divide the total computational budget optimally between all correction terms. In this process, it is not necessary to use the same numerical method for all terms. Indeed, we have used a combination of closed-form and numerical solutions in the tests, and the framework is rich enough to use the best available method (e.g. Fourier-based methods for affine models, PDEs for early exercise options, or conceivably even a Monte Carlo method for strongly path-dependent derivatives) for a given sub-problem. Moreover, for the computation of a specific correction term only a subset of derivatives has to be considered, namely those affected by the risk factors considered in that correction. So it is conceivable that even for a large derivative portfolio each correction term only requires consideration of a small fraction of the derivatives. For instance, if the portfolio contains derivatives on equities, commodities and FX, the risk factor decomposition may provide a way to decompose the exposure computations into smaller sub-computations.

The practical challenge will be to develop a framework which allows this to happen in a generic way, and which is easily adaptable as different derivatives are entered or the modelling framework changes.

Acknowledgements

The authors would like to thank Dr. Sumit Sourabh for helpful advice and Dr. Shashi Jain for providing the data.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was support by the Dutch Technology Foundation STW [project 12214] is gratefully acknowledged.

References

- Alexander, S., Coleman, T.F. and Li, Y., Minimizing CVaR and VaR for a portfolio of derivatives. *J. Banking Finance*, 2006, **30**, 583–605.
- Andersen, L.B. and Piterbarg, V.V., *Interest Rate Modeling*, 2010 (Atlantic Financial Press: London).
- Andreasen, J. and Huge, B., Finite difference based calibration and simulation. 2010. Available online at ssrn.com/abstract=1697545
- Bouzeubaa, M. and Osseiran, A., *Exotic Options and Hybrids: A Guide to Structuring, Pricing and Trading*, The Wiley Finance Series, 2010 (Wiley: Chichester).
- Brigo, D. and Mercurio, F., *Interest Rate Models Theory and Practice*, Springer Finance, 2013 (Springer: Berlin Heidelberg).
- Broadie, M., Du, Y. and Moallemi, C.C., Efficient risk estimation via nested sequential simulation. *Manage. Sci.*, 2011, **57**, 1172–1194.
- Broadie, M., Du, Y. and Moallemi, C.C., Risk estimation via regression. *Oper. Res.*, 2015, **63**, 1077–1097.

- Capriotti, L., Jiang, Y. and Macrina, A., Real-time risk management: An AAD-PDE approach. *Int. J. Financ. Eng.*, 2015, **2**, 1550039.
- Clark, I.J., *Foreign Exchange Option Pricing: A Practitioner's Guide*, The Wiley Finance Series, 2011 (Wiley: Chichester).
- Cozma, A. and Reisinger, C., A mixed Monte Carlo and PDE variance reduction method for foreign exchange options under the Heston-CIR model. *J. Comput. Finance*, 2017, **30**, 109–149.
- Dang, D.M., Jackson, K.R. and Mohammadi, M., Dimension and variance reduction for Monte Carlo methods for high-dimensional models in finance. *Appl. Math. Finance*, 2015, **22**, 522–552.
- Dang, D.M., Jackson, K.R. and Sues, S., A dimension and variance reduction Monte-Carlo method for option pricing under jump-diffusion models. *Appl. Math. Finance*, 2017, **24**, 175–215.
- De Col, A., Gnoatto, A. and Grasselli, M., Smiles all around: FX joint calibration in a multi-Heston model. *J. Banking Finance*, 2013, **37**, 3799–3818.
- de Graaf, C.S.L., Feng, Q., Kandhai, B.D. and Oosterlee, C.W., Efficient computation of exposure profiles for counterparty credit risk. *Int. J. Theor. Appl. Finance*, 2014, **17**, 1450024.
- de Graaf, C.S.L., Kandhai, B.D. and Reisinger, C., Efficient exposure computation by risk factor decomposition. 2016. arXiv preprint arXiv:1608.01197.
- Doust, P., The stochastic intrinsic currency volatility model: A consistent framework for multiple FX rates and their volatilities. *Appl. Math. Finance*, 2012, **19**, 381–445.
- Filipovic, D., *Term-Structure Models: A Graduate Course*, Springer Finance, 2009 (Springer: Berlin Heidelberg).
- Glasserman, P., *Monte Carlo Methods in Financial Engineering. Applications of Mathematics: Stochastic Modelling and Applied Probability*, 2004 (Springer: New York).
- Gordy, M.B. and Juneja, S., Nested simulation in portfolio risk measurement. *Manage. Sci.*, 2010, **56**, 1833–1848.
- Green, A.D., Kenyon, C. and Dennis, C.R., KVA: Capital valuation adjustment by replication. *Risk*, 2014, **27**, 82–87.
- Gregory, J., *Counterparty Credit Risk: The New Challenge for Global Financial Markets*, 2010 (Wiley: Chichester).
- Gregory, J., *The xVA Challenge: Counterparty Credit Risk, Funding, Collateral, and Capital*, 2015 (Wiley: Chichester).
- Griebel, M. and Holtz, M., Dimension-wise integration of high-dimensional functions with applications to finance. *J. Complexity*, 2010, **26**, 455–489.
- Haentjens, T. and in 't Hout, K.J., ADI finite difference schemes for the Heston-Hull-White PDE. *J. Comput. Finance*, 2012, **16**, 83–110.
- Haentjens, T. and in 't Hout, K.J., ADI schemes for pricing American options under the Heston model. *Appl. Math. Finance*, 2015, **20**, 207–235.
- Heston, S.L., A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.*, 1993, **6**, 327–343.
- Hundsdoerfer, W. and Verwer, J.G., *Numerical Solution of Time-dependent Advection-diffusion-reaction Equations*, 2003 (Springer-Verlag: Berlin).
- in 't Hout, K.J. and Foulon, S., ADI finite difference schemes for option pricing. *Int. J. Numer. Anal. Model.*, 2010, **7**, 303–320.
- in 't Hout, K.J. and Welfert, B., Unconditional stability of second-order ADI schemes applied to multi-dimensional diffusion equations with mixed derivative terms. *Appl. Numer. Math.*, 2009, **59**, 677–692.
- Hout, In't, K.J. and Wyns, M., Convergence of the Modified Craig-Sneyd scheme for two-dimensional convection-diffusion equations with mixed derivative term. *J. Comput. Appl. Math.*, 2016, **296**, 170–180.
- Itkin, A., High-order splitting methods for forward PDEs and PIDEs. *Inter. J. Theor. Appl. Finance*, 2015, **18**, 1550031.
- Jain, S. and Oosterlee, C.W., Pricing high-dimensional Bermudan options using the stochastic grid method. *Int. J. Comput. Math.*, 2012, **89**, 1186–1211.
- Joshi, M. and Kang, O., Least squares Monte Carlo credit value adjustment with small and unidirectional bias. Available at SSRN, 2016, **2717250**.
- Karlsson, P., Jain, S. and Oosterlee, C.W., Counterparty credit exposures for interest rate derivatives using the stochastic grid bundling method. Available at SSRN, 2014, **2538173**, 1–25.
- Longstaff, F. and Schwartz, E., Valuing American options by simulation: A simple least-squares approach. *Rev. Financ. Stud.*, 2001, **1**, 113–147.
- Musiela, M. and Rutkowski, M., *Martingale Methods in Financial Modelling*, 2nd ed., 2005 (Springer-Verlag: Berlin).
- Pykhtin, M. and Zhu, S., A guide to modelling counterparty credit risk. *Global Assoc. Risk Professionals*, 2007, **37**, 16–22.
- Rebonato, R., On the simultaneous calibration of multifactor lognormal interest rate models to Black volatilities and to the correlation matrix. *J. Comput. Finance*, 1999, **2**, 5–27.
- Reisinger, C. and Wissmann, R., Numerical valuation of derivatives in high-dimensional settings via PDE expansions. *J. Comput. Finance*, 2015, **18**(4), 95–127.
- Reisinger, C. and Wissmann, R., Error analysis of truncated expansion solutions to high-dimensional parabolic PDEs. *ESAIM: Math. Model. Numer. Anal.*, 2017, **51**(6), 2435–2463.
- Reisinger, C. and Wittum, G., Efficient hierarchical approximation of high-dimensional option pricing problems. *SIAM J. Sci. Comput.*, 2005, **29**, 440–458.
- Simaitis, S., de Graaf, C.S., Hari, N. and Kandhai, D., Smile and default: The role of stochastic volatility and interest rates in counterparty credit risk. *Quant. Finance*, 2016, **16**, 1725–1740.
- Sorensen, E. and Bollier, T., Pricing swap default risk. *Financ. Analysts J.*, 1994, **50**, 23–33.
- Tavella, D. and Randall, C., *Pricing Financial Instruments: The Finite Difference Method*, 2000 (Wiley: New York).
- Wang, X. and Sloan, I.H., Why are high-dimensional problems in finance often of low effective dimension? *Soc. Ind. Appl. Math.*, 2005, **27**, 159–183.

Appendix 1. Model parameters

The full correlation matrix for the process $(X_t^1, \dots, X_t^7) = (F_t^1, R_t^d, R_t^{f,1}, F_t^2, R_t^{f,2}, F_t^3, R_t^{f,3})$ from (36), used in the tests of section 5, after regularization for positive definiteness (Rebonato 1999):

$$\begin{pmatrix} 1 & -0.3024 & 0.1226 & 0.5815 & -0.0142 & 0.5510 & 0.5351 \\ -0.3024 & 1 & 0.6293 & -0.2577 & 0.6895 & -0.4554 & 0.3188 \\ 0.1226 & 0.6293 & 1 & 0.0459 & 0.7453 & -0.3049 & 0.4181 \\ 0.5815 & -0.2577 & 0.0459 & 1 & 0.1230 & 0.5490 & -0.0848 \\ -0.0142 & 0.6895 & 0.7453 & 0.1230 & 1 & -0.3015 & 0.3587 \\ 0.5510 & -0.4554 & -0.3049 & 0.5490 & -0.3015 & 1 & -0.3260 \\ 0.5351 & 0.3188 & 0.4181 & -0.0848 & 0.3587 & -0.3260 & 1 \end{pmatrix}$$

The FX and interest rate SDEs are driven by the parameters in table A1.

As discussed in Section 4.3, the functions $\Theta_d(t)$, and $\Theta_f^i(t)$, $i = 1, 2, 3$, are calibrated to fit the forward rate curve of the respective markets as seen on 2 December 2014. The ATM volatilities of the FX rates at this date used for $\sigma(t)$ can be found in table A2.

Table A1. Parameters for the BS2HW model (29).

EUR		$i = 1$ (USD)	$i = 2$ (GBP)	$i = 3$ (JPY)
R_0^d	F_0^i	1.2470	0.7926	147.53
	$R_0^{f,i}$	-0.0036	0.0065	0.0011
λ_d	λ_f^i	0.010	0.0523	0.010
η_d	η_f^i	0.0092	0.0104	0.0057

The full correlation matrix, including stochastic volatility, for the process $(X_t^1, \dots, X_t^{10}) = (F_t^1, Y_t^1, R_t^d, R_t^{f,1}, F_t^2, Y_t^2, R_t^{f,2}, F_t^3, Y_t^3, R_t^{f,3})$ used in the tests of section 5.5, after regularization for positive definiteness (Rebonato 1999):

$$\begin{pmatrix} 1 & -0.2644 & -0.2910 & 0.1142 & 0.5673 & -0.0004 & -0.0107 & 0.5243 & -0.0055 & 0.5043 \\ -0.2644 & 1 & 0.0019 & -0.0018 & -0.0016 & -0.0001 & 0.0009 & -0.0043 & -0.0014 & -0.0052 \\ -0.2910 & 0.0019 & 1 & 0.6322 & -0.2558 & -0.0001 & 0.6892 & -0.4535 & -0.0015 & 0.3284 \\ 0.1142 & -0.0018 & 0.6322 & 1 & 0.0438 & 0.0001 & 0.7473 & -0.3091 & 0.0014 & 0.4171 \\ 0.5673 & -0.0016 & -0.2558 & 0.0438 & 1 & -0.0740 & 0.1248 & 0.5468 & 0.0013 & -0.0796 \\ -0.0004 & -0.0001 & -0.0001 & 0.0001 & -0.0740 & 1 & -0.0001 & 0.0003 & 0.0001 & 0.0003 \\ -0.0107 & 0.0009 & 0.6892 & 0.7473 & 0.1248 & -0.0001 & 1 & -0.3001 & -0.0007 & 0.3659 \\ 0.5243 & -0.0043 & -0.4535 & -0.3091 & 0.5468 & 0.0003 & -0.3001 & 1 & -0.3452 & -0.3103 \\ -0.0055 & -0.0014 & -0.0015 & 0.0014 & 0.0013 & 0.0001 & -0.0007 & -0.3452 & 1 & 0.0041 \\ 0.5043 & -0.0052 & 0.3284 & 0.4171 & -0.0796 & 0.0003 & 0.3659 & -0.3103 & 0.0041 & 1 \end{pmatrix} \quad (A1)$$

Note, that due to the regularization the correlation matrix is altered, and correlations are slightly different from the Black–Scholes case.

A.1. Heston model parameters

In table A3 we give the calibrated Heston parameters together with the mean implied volatility error. As shown in figures A1(a), (b) and (c), the skew in the EURJPY market is most pronounced (note the different scales on the y-axis). Due to this skew, the bootstrapped ATM volatility is different from the expectation of the future volatility in (44). This mean volatility and the bootstrapped volatility over time are presented in figure A1(d).

Table A2. ATM volatilities as seen on 2 December 2014 that are used for bootstrapping the piecewise constant volatility function as explained in section 4.3.

	EURUSD (%)	EURGBP (%)	EURJPY (%)
$T = 1M$	8.852	6.570	10.247
$T = 3M$	8.695	6.635	10.245
$T = 6M$	8.580	7.350	10.517
$T = 1Y$	8.605	7.447	10.848
$T = 2Y$	8.717	7.865	11.580
$T = 3Y$	8.952	8.068	12.247
$T = 5Y$	9.635	8.383	13.642

Table A3. Calibrated Heston parameters to 2014 market data. Including mean implied vol errors quoted in percentages. The models are calibrated to 10%-, 25%- and 50%-Δ FX put and call options with maturities 1Y, 2Y, 3Y and 5Y.

	Parameters	Error
EURUSD	κ	0.5449
	v_0	0.0072
	\bar{v}	0.0126
	ρ	-0.2752
	γ	0.1560
	κ	0.6740
EURGBP	v_0	0.0054
	\bar{v}	0.0098
	ρ	-0.0762
	γ	0.1771
	κ	0.0476
	v_0	0.0116
EURJPY	\bar{v}	0.1428
	ρ	0.1890
	γ	-0.3507
	κ	2.23%

Appendix 2. Numerical approximation of the Kolmogorov PDEs

Here, we outline the finite difference method used for the PDE solutions, as it is somewhat non-standard through the use of forward and backward equations. In particular, we extend the adjoint method from Itkin (2015) from two to three dimensions and use it for the forward component. For a more general introduction to finite difference methods for pricing financial derivatives, see for example Tavella and Randall (2000) and, specifically for interest rate derivatives Andersen and Piterbarg (2010).

As usual, a grid is defined with one dimension per risk factor, and the partial derivatives in the Kolmogorov forward (9) or backward PDE (3) are approximated by finite differences on this grid.

B.1. Backward equations

We use a combination of second order central differences in the centre of the spatial domain, and upwinding for stabilization of large drifts in outer parts of the domain (see appendix B.3), to obtain a system of ODEs in the time variable (in 't Hout and Welfert 2009). Let $U(\tau)$ be the solution to this semi-discrete Kolmogorov backward equation at time τ , where $\tau = T - t$, then we have the following initial value problem

$$\frac{dU}{d\tau} = F(\tau)U, \quad \tau \geq 0, \quad U(0) = U_0, \quad (B1)$$

with given matrix F , derived from the PDE, and initial vector U_0 given by the pay-off at the mesh points. The boundary conditions are also derived from the pay-off function.

We apply an Alternating Direction Implicit (ADI) splitting method. Let $F^n = F(\tau_n)$ be the discretization matrix, as per (B1), at time step $\tau_n = n\Delta\tau$, $n \in \mathbb{N}$, $\Delta\tau > 0$ a uniform step size, then this matrix is first decomposed into

$$F^n = F_0^n + F_1^n + F_2^n + F_3^n,$$

where the individual F_i^n , $1 \leq i \leq d$ contains the contribution to F stemming from the first and second order derivatives in the i th dimension. Following in 't Hout and Foulon (2010), we define one matrix F_0 which accounts for the mixed derivative terms, and treat that fully explicitly. Here we show the scheme for a three-dimensional problem, and the two-dimensional case is obtained as special case by setting $F_3 = 0$. The Hundsdorfer-Verwer (HV) scheme (Hundsdorfer and Verwer 2003),

$$\begin{aligned} Y_0 &= U_{n-1} + \Delta\tau F^{n-1} U_{n-1}, \\ (I - \theta\Delta\tau F_j^n) Y_j &= Y_{j-1} - \theta\Delta\tau F_j^{n-1} U_{n-1}, \quad j = 1, 2, 3, \\ \tilde{Y}_0 &= Y_0 + \frac{1}{2}\Delta\tau [F^n Y_3 - F^{n-1} U_{n-1}], \\ (I - \theta\Delta\tau F_j^n) \tilde{Y}_j &= Y_{j-1} - \theta\Delta\tau F_j^n Y_j, \quad j = 1, 2, 3, \end{aligned}$$

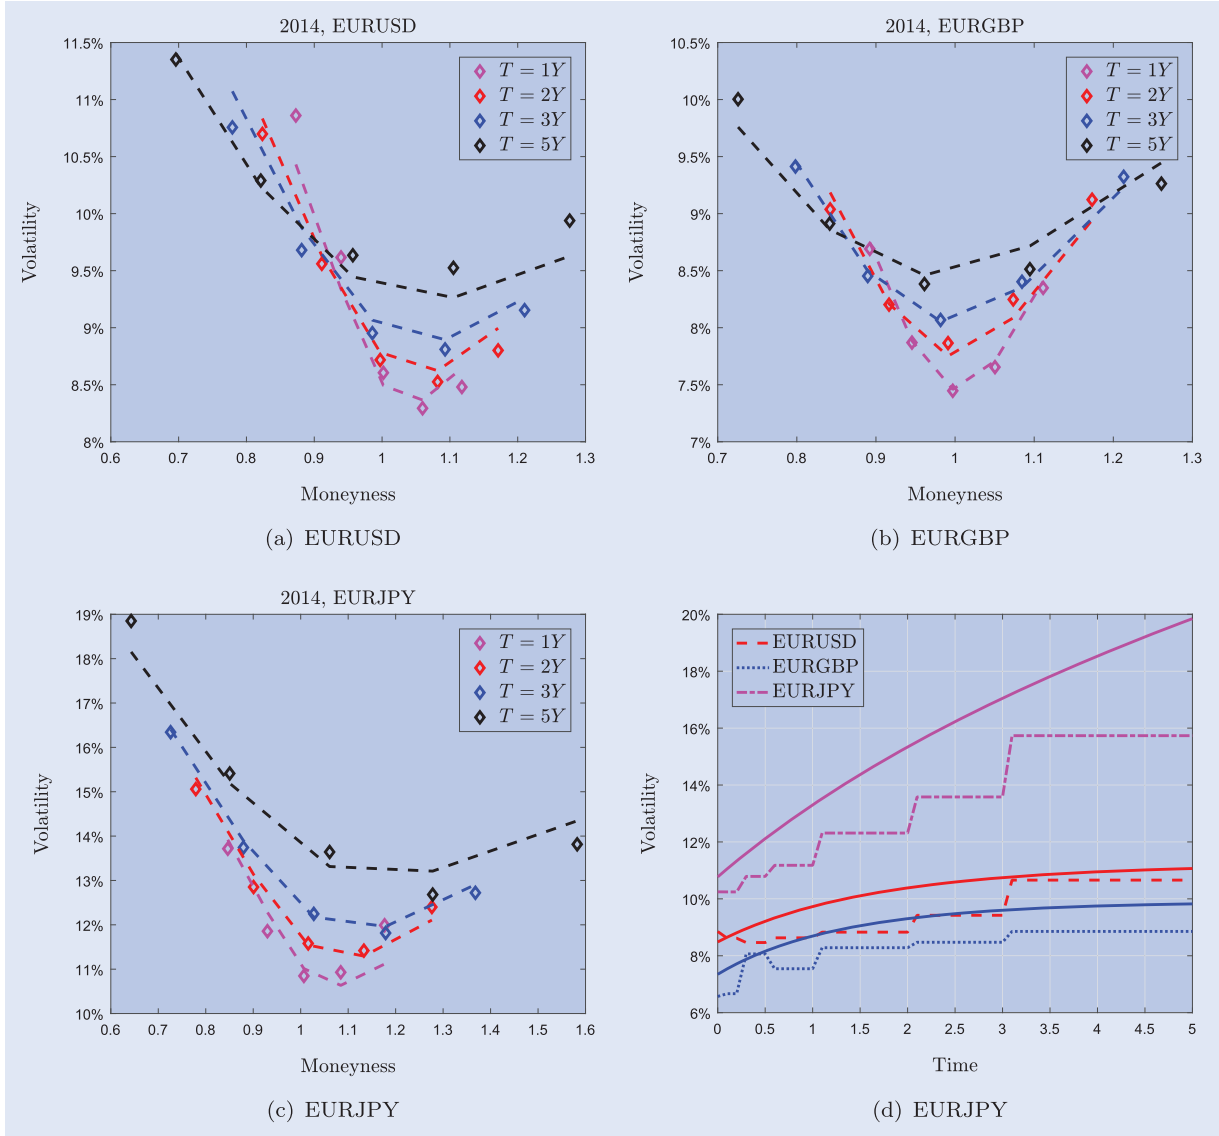


Figure A1. Implied volatility smiles of FX pairs. The lines are the model fit and the diamonds are the market implied volatilities. The associated errors are presented in table A3. In figure 1(d), shown are the piecewise constant volatilities as a function of time, based on bootstrapped ATM volatilities (dashed lines) and the expectations of the model stochastic volatilities (solid lines).

$$U_n = \tilde{Y}_3,$$

defines a second order consistent ADI splitting for all θ , and can be shown to be von Neumann stable for $\theta \in \left[\frac{1}{2} + \frac{1}{6}\sqrt{3}, 1\right]$, see Haentjens (2012). We use $\theta = 0.8 > \frac{1}{2} + \frac{1}{6}\sqrt{3} \approx 0.789$ in the computations. Accuracy and stability do not appear to be very sensitive to the choice of θ within the above range.

B.2. Forward equations

We use the adjoint relation between the Kolmogorov forward and the backward PDE. It can easily be seen that (for sufficiently smooth coefficients)

$$\frac{dP}{dt} = F^T(t)P, \quad t \geq 0, \quad P(0) = P_0, \quad (B2)$$

with F from above, is a consistent scheme for the Kolmogorov forward equation. The initial datum for the discrete density function P is given by an approximation P_0 to the Dirac delta. We choose a mesh such that a mesh point coincides with the location of the Dirac delta. Then P_0 is set to zero for all other mesh points and to a large value at this particular point, chosen such that a numerical quadrature

rule applied to P_0 gives 1. At the spatial boundaries, we apply zero Dirichlet conditions.

Hence, following Andreasen and Huge (2010) and Capriotti et al. (2015), we can first set up matrices that approximate the partial derivatives in the Kolmogorov Backward PDE, and use the transpose of these matrices for the Kolmogorov forward PDE.

Now we could apply HV splitting to (B2) and obtain a second order (in time) consistent approximation to the forward PDE. Instead, we follow Itkin (2015) to calculate the exact adjoint of the HV scheme for the backward equation, which results in a different scheme, i.e. the transposition and approximate factorization do not commute. The forward scheme, adapted to the three-dimensional case from the two-dimensional case analysed in Itkin (2015), then is,

$$\begin{aligned} (I - \theta \Delta t F_3^n)^T Y_0 &= P_{n-1}, \\ (I - \theta \Delta t F_2^n)^T Y_1 &= Y_0, \\ (I - \theta \Delta t F_1^n)^T Y_2 &= Y_1, \\ \tilde{Y}_0 &= P_{n-1} + \Delta t \left(\left(\frac{1}{2} (F^n)^T - \theta (F_1^n)^T \right) \right. \\ &\quad \left. Y_2 - \theta (F_2^n)^T Y_1 - \theta (F_3^n)^T Y_0 \right) \\ (I - \theta \Delta t F_3^n)^T \tilde{Y}_1 &= \tilde{Y}_0. \end{aligned}$$

Table B1. Finite difference grid parameters.

	min	max	ξ
FX rate (F_t)	0	$8F_0$	20
Domestic IR (R_t^d)	-0.5	0.8	100
Foreign IR (R_t^f)	-0.5	0.8	100

$$\begin{aligned}
(I - \theta \Delta t F_2^n)^T \tilde{Y}_2 &= \tilde{Y}_1, \\
(I - \theta \Delta t F_1^n)^T \tilde{Y}_3 &= \tilde{Y}_2, \\
P_n &= \tilde{Y}_3 + \Delta t \left(((F^{n-1})^T - \theta (F_1^{n-1})^T) (\tilde{Y}_3 - Y_2) \right. \\
&\quad \left. - \theta (F_2^{n-1})^T (\tilde{Y}_2 - Y_1) \right. \\
&\quad \left. - \theta (F_3^{n-1})^T (\tilde{Y}_1 - Y_0) + \frac{1}{2} (F^{n-1})^T \tilde{Y}_3 \right).
\end{aligned}$$

Using this scheme, we have the adjoint relation,

$$P_0^T U_n = P_n^T U_0,$$

i.e. using the backward equation to compute the conditional expectation function and then evaluating it at the location of the Dirac measure (left-hand side), or computing the density and then integrating over the “payoff” function, gives exactly the same result. Apart from the mathematical aestheticism of using the exact adjoint, this scheme has the following stability advantages. As discussed in Hout (2016), splitting schemes do not give stable, and hence convergent, solutions for Dirac initial data. A direct splitting of (B2) is prone to wildly oscillating densities and may give meaningless expectations. Using the above scheme instead, we can be sure that even if we have no guarantee for the stability of the solution, the derived quantity of interest, e.g. EPE, is the same as for the backward equation with regular (for EPE, Lipschitz and piecewise smooth) data.† The adjoint equation also ensures conservation of discrete probability in the interior of the mesh. The loss of mass at the boundaries is minimized by choosing the domain large enough.

With the grid of discrete transition probabilities, obtained by the numerical solution of the Kolmogorov forward PDE, and a grid of portfolio values from the backward PDE approximations, the exposures can be extracted by numerical quadrature applied to (8).

B.3. Grid construction and numerical parameters

The grids used for the approximation are non-uniform as in Haentjens (2015), where a sinh transformation is used to place a high density of grid points around the spot values. Moreover, the grid is shifted to include the initial spot value where the non-smooth Dirac delta function for the forward equation is located.

In table B1 the computational domain in F_t , R_t^d and R_t^f is chosen. The parameter ξ controls the fraction of points that lie close to initial spot values (see Haentjens 2015).

Moreover, we use upwind differences for large absolute interest rates, i.e. $R \in [-0.5, -0.1] \cup [0.2, 0.8]$, otherwise central differences.

The number of grid points m_1, m_2, m_3 in the three directions is chosen as $m_1 = 2m_2 = 2m_3$; see also appendix 4.

For the F grid, an interval $[F_{\text{left}}, F_{\text{right}}] \subseteq [F_{\text{min}}, F_{\text{max}}]$ is defined wherein the grid is uniform and dense, similar to Haentjens (2012), while outside the sinh transformation is used. We set this interval as $[F_{\text{left}}, F_{\text{right}}] = [0.95F_0, 1.02F_0]$.

Appendix 3. Regression-based Monte Carlo algorithm

We use constant, linear and bi-linear basis functions

$$\{\psi_j : j = 1, \dots, 6\} = \{F, R^d, R^f, FR^d, FR^f, R^d R^f\},$$

†But see Remark 1 why we are not using the backward equation directly to compute exposures.

Table D1. Finite difference errors of exposures (EE) of a portfolio with three CCYS. The nomenclature of the terms is explained in Remark 4. The errors are measured in e_{L_2} (see equation (34)) in percentages for different number of mesh points. The number of time steps is fixed at 500. Within brackets (for $m_1 = 40$) the error for a FD solution with 1000 time steps. The Monte Carlo benchmark is computed with $4 \cdot 10^6$ paths and 1000 time steps.

	MC SE	$m_1 = 40$	$m_1 = 60$	$m_1 = 80$	$m_1 = 100$
BS EU	0.15	0.22 (0.18)	0.064	0.12	0.077
BSBS EU-EG	0.22	0.23 (0.20)	0.072	0.13	0.087
BSBS EU-EJ	0.22	0.22 (0.19)	0.050	0.10	0.057
BSHW EU-RE	0.21	0.24 (0.15)	0.15	0.20	0.17
BSHW EU-RU	0.22	0.22 (0.20)	0.23	0.18	0.21
BSHW EU-RG	0.22	0.22 (0.18)	0.058	0.11	0.071
BSHW EU-RJ	0.22	0.21 (0.17)	0.061	0.11	0.075
BSBSBS EU-EG-EJ	0.21	0.24 (0.20)	0.088	0.14	0.10
BSBHW EU-RE-RU	0.21	0.16 (0.14)	0.14	0.098	0.11
BSBSHW EU-EG-RE	0.21	0.23 (0.15)	0.16	0.21	0.18
BSBSHW EU-EJ-RE	0.21	0.22 (0.15)	0.13	0.18	0.15
BSBSHW EU-EG-RU	0.22	0.24 (0.23)	0.21	0.17	0.19
BSBSHW EU-EJ-RU	0.22	0.22 (0.21)	0.20	0.16	0.18
BSBSHW EU-EG-RG	0.22	0.25 (0.22)	0.088	0.14	0.096
BSBSHW EU-EJ-RJ	0.22	0.23 (0.19)	0.063	0.12	0.074

to account for the correlation between risk factors. The algorithm determines option values $Y_i(t) = Y(t; \omega_k)$ along the sample path ω_i , and can be summarized by the following steps:

- (i) Generate all the paths $(F_t(\omega_i), R_t^d(\omega_i), R_t^f(\omega_i))$ at all the time points $t = 0, \Delta t, \dots$
- (ii) Start from maturity and set $t \leftarrow T$, $Y_i(T) \leftarrow \phi(F_T(N\omega_i))$.
- (iii) Set $Y_i(t) \leftarrow \exp(-\Delta t R_t^d(\omega_i)) Y_i(t)$.
- (iv) Solve the linear regression problem

$$\mathbb{E} \left[\left(\mathbb{E} \left[Y(t) | X_{t-\Delta t} \right] - \sum_{j=1}^6 \beta_j \psi_j(X_{t-\Delta t}) \right)^2 \right] \rightarrow \min_{\beta}$$

by approximating the expectations on the paths ω_i to find a minimizer $\hat{\beta}$.

- (v) Set $t \leftarrow T - \Delta t$ and

$$\begin{aligned}
Y_i(t) &\leftarrow \hat{\beta}_0 + \hat{\beta}_1 F_t(\omega_i) + \hat{\beta}_2 R_t^d(\omega_i) + \hat{\beta}_3 R_t^f(\omega_i) \\
&\quad + \hat{\beta}_4 F_t(\omega_i) R_t^d(\omega_i) + \hat{\beta}_5 F_t(\omega_i) R_t^f(\omega_i) \\
&\quad + \hat{\beta}_6 R_t^d(\omega_i) R_t^f(\omega_i)
\end{aligned}$$

- (vi) Repeat backwards from 2. for all time steps.

Appendix 4. Finite difference errors and individual terms for Case B

Here we show the individual FD errors for all terms involved for Case B, compared to a Monte Carlo estimate, for both EE (table D1) and EPE (Table D2).

Remark 4 In the following tables, for instance, BSHW EU-RU refers to $V_{\{1,5\}}(F^1, R^{f,1})$ from Section 5, i.e. a Black–Scholes–Hull–White model where the EURUSD rate F^1 follows a Black–Scholes-type model but with Hull–White foreign short rate $R^{f,1}$ and deterministic domestic rate, and similarly for the other terms.

The results in table D1 (especially) and table D2 show that the FD errors are already in the same order of magnitude as the MC standard error for a larger number of samples (i.e. significantly more samples than are used in practice). We therefore, use $m_1 = 60$ mesh points and 500 time steps in most of the computations (unless otherwise stated).

Table D2. Finite difference errors of positive exposures (EPE) of a portfolio with three CCYS. The nomenclature of the terms is explained in Remark 4. The errors are measured in e_{L_2} (see equation (34)) in percentages for different number of mesh points. The number of time steps is fixed at 500. Within brackets (for $m_1 = 40$) the error for a FD solution with 1000 time steps. The Monte Carlo benchmark is computed with $4 \cdot 10^6$ paths and 1000 time steps.

	MC SE	$m_1 = 40$	$m_1 = 60$	$m_1 = 80$	$m_1 = 100$
BS EU	0.073	0.30 (0.34)	0.099	0.042	0.030
BSBS EU-EG	0.070	0.23 (0.25)	0.11	0.078	0.071
BSBS EU-EJ	0.084	0.24 (0.27)	0.070	0.038	0.048
BSHW EU-RE	0.092	0.31 (0.37)	0.11	0.050	0.041
BSHW EU-RU	0.094	0.45 (0.42)	0.24	0.17	0.14
BSHW EU-RG	0.093	0.29 (0.32)	0.090	0.040	0.043
BSHW EU-RJ	0.093	0.30 (0.34)	0.10	0.054	0.048
BSBSBS EU-EG-EJ	0.064	0.21 (0.23)	0.090	0.056	0.045
BSHWHW EU-RE-RU	0.091	0.53 (0.51)	0.31	0.24	0.21
BSBSHW EU-EG-RE	0.069	0.22 (0.25)	0.12	0.10	0.10
BSBSHW EU-EJ-RE	0.083	0.25 (0.29)	0.080	0.054	0.062
BSBSHW EU-EG-RU	0.071	0.33 (0.30)	0.18	0.13	0.11
BSBSHW EU-EJ-RU	0.085	0.44 (0.41)	0.26	0.20	0.17
BSBSHW EU-EG-RG	0.070	0.26 (0.28)	0.12	0.076	0.058
BSBSHW EU-EJ-RJ	0.084	0.29 (0.32)	0.11	0.051	0.030

Table D3. Errors of decomposed approximations to exposures of a portfolio with three CCYS (see table 2, Case B), computed with different numbers of mesh points. The errors are measured in e_{L_2} (see equation (34)) in percentages for different number of grid points in space. The number of time steps is fixed at 500. The Monte Carlo estimator for these approximation errors is computed with $4 \cdot 10^6$ paths and 1000 time steps. Within brackets for MC the error for a MC simulation with identical paths for all estimators, as explained in remark 5. Within brackets for $m_1 = 40$ the error for a FD computation with 1000 timesteps.

		MC	$m_1 = 40$	$m_1 = 60$	$m_1 = 80$	$m_1 = 100$
EE	2D corrections	1.27 (1.27)	1.35 (1.31)	1.41	1.36	1.39
	3D corrections	0.19 (0.24)	0.25 (0.27)	0.20	0.18	0.20
EPE	2D corrections	2.39 (2.38)	2.20 (2.21)	2.29	2.32	2.36
	3D corrections	0.41 (0.24)	0.31 (0.29)	0.26	0.26	0.26

Table D4. Contribution of individual corrections for EE and EPE of a portfolio with three CCYS. The nomenclature of the terms is explained in remark 4. The differences are measured in e_{L_2} (see equation (34)) in percentages for $m_1 = 60$ grid points in space.

2D terms	EE	EPE	3D terms	EE	EPE
BSBS EU-EG	14.58	40.35	BSBSBS EU-EG-EJ	0.000	0.50
BSBS EU-EJ	8.74	18.40	BSHWHW EU-RE-RU	1.19	2.37
BSHW EU-RE	1.40	2.73	BSBSHW EU-EG-RE	0.31	1.28
BSHW EU-RU	0.52	0.51	BSHWHW EU-RE-RU	1.19	2.37
BSHW EU-RG	0.0003	0.0022	BSBSHW EU-EJ-RE	0.18	0.72
BSHW EU-RJ	0.000	0.0093	BSBSHW EU-EG-RU	0.000	0.083
			BSBSHW EU-EJ-RU	0.000	0.58
			BSBSHW EU-EG-RG	0.046	0.066
			BSBSHW EU-EJ-RJ	0.12	0.30

In table D3, we report the accuracy of the complete approximations for increasing number of mesh points m_1 . For reference, we also compute the approximation with a standard Monte Carlo estimator.

Remark 5 In the first column of Table D3, we report the results for an estimator where the same Brownian paths are used for the estimation of all V_u in (20) for a given ΔV_u , and in brackets the results if the same

paths are also used across all u . It does not seem generally clear which of the estimators has the smaller variance. Using independent paths for different u results in a summation of the variances of ΔV_u . Using the same paths for different u is expected to increase the variance if the terms are predominantly positively correlated, and decrease the variance if they are predominantly negatively correlated.