LUND UNIVERSITY

Essays in Quantitative Finance

Karlsson, Patrik

2016

Document Version:
Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):
Karlsson, P. (2016). *Essays in Quantitative Finance*. [Doctoral Thesis (compilation), Department of Economics].

Total number of authors:
1

Essays in Quantitative Finance

Patrik Karlsson

LUND
UNIVERSITY

Essays in Quantitative Finance

Patrik Karlsson

Organization	Document name	
Lund University	Doctoral dissertation	
Department of Economics	**Date of issue**	
P.O. Box 7082	23 February 2017	
S-220 07 Lund, Sweden		
Author	**Sponsoring organization**	
Patrik Karlsson		

Title

Essays in Quantitative Finance

Abstract

This thesis contributes to the quantitative finance literature and consists of four research papers.

  Paper 1. This paper constructs a hybrid commodity interest rate market model with a stochastic local volatility function that allows the model to simultaneously fit the implied volatility of commodity and interest rate options. Because liquid market prices are only available for options on commodity futures (not forwards), a convexity correction formula is derived to account for the difference between forward and futures prices. A procedure for efficiently calibrating the model to interest rate and commodity volatility smiles is constructed. Finally, the model is fitted to an exogenously given cross-correlation structure between forward interest rates and commodity prices. When calibrating to options on forwards (rather than futures), the fitting of cross-correlation preserves the (separate) calibration in the two markets (interest rate and commodity options), whereas in the case of futures, a (rapidly converging) iterative fitting procedure is presented. The cross-correlation fitting is reduced to finding an optimal rotation of volatility vectors, which is shown to be an appropriately modified version of the "orthonormal Procrustes" problem. The calibration approach is demonstrated on market data for oil futures.

  Paper 2. This paper describes an efficient American Monte Carlo approach for pricing Bermudan swaptions in the LIBOR market model using the Stochastic Grid Bundling Method (SGBM) which is a regression-based Monte Carlo method in which the continuation value is projected onto a space in which the distribution is known. We demonstrate an algorithm to obtain accurate and tight lower–upper bound values without the need for the nested Monte Carlo simulations that are generally required for regression-based methods.

  Paper 3. The credit valuation adjustment (CVA) for over-the-counter derivatives are computed using the portfolio's exposure over its lifetime. Usually, future exposure is approximated by Monte Carlo simulations. For derivatives that lack an analytical approximation for their mark-to-market (MtM) value, such as Bermudan swaptions, the standard practice is to use the regression functions from the least squares Monte Carlo method to approximate their simulated MtMs. However, such approximations have significant bias and noise, resulting in an inaccurate CVA charge. This paper extend the SGBM to efficiently compute expected exposure, potential future exposure, and CVA for Bermudan swaptions. A novel contribution of the paper is that it demonstrates how different measures, such as spot and terminal measures, can simultaneously be employed in the SGBM framework to significantly reduce the variance and bias

  Paper 4. This paper presents an algorithm for simulation of options on Lévy driven assets. The simulation is performed on the inverse transition matrix of a discretised partial differential equation. We demonstrate how one can obtain accurate option prices and deltas on the variance gamma (VG) and CGMY model through finite element-based Monte Carlo simulations.

Keywords
Credit Valuation Adjustment (CVA), Derivative Pricing, Interest Rate Derivatives, Monte Carlo Simulation

Classification system and/or index terms (if any)	
JEL Classification:	
Supplementary bibliographical information	**Language**
	English
ISSN and key title	**ISBN**
0460-0029 Lund Economic Studies no. 199	978-91-7753-060-2 (print)
	978-91-7753-061-9 (pdf)

Recipient's notes	**Number of pages**	**Price**
	150	
	Security classification	

Signature   Date 2016-12-15

Essays in Quantitative Finance

Patrik Karlsson

LUND
UNIVERSITY

To my Mother

Contents

Paper 4: Finite Element Based Monte Carlo Simulation of Option Prices on Lévy Driven Assets   **97**

Epilogue : The Future of Quantitative Finance   **123**

List of Publications

1. Karlsson, P., Pilz, K. and Schlögl, E. (2016). Calibrating a Market Model with Stochastic Volatility to Commodity and Interest Rate Risk. *Quantitative Finance*. Forthcoming. `http://dx.doi.org/10.1080/14697688.2016.1254814`.

2. Feng, Q., Jain, S., Karlsson, P., Kandhai, D. and Oosterlee, C.W. (2016). Efficient computation of exposure profiles on real-world and risk-neutral scenarios for Bermudan swaptions. *Journal of Computational Finance* 20(1): 139-172. `http://dx.doi.org/10.21314/JCF.2017.337`.

3. Karlsson, P., Jain, S. and Oosterlee, C.W. (2016). Counterparty Credit Exposures for Interest Rate Derivatives using the Stochastic Grid Bundling Method. *Applied Mathematical Finance* 23(1): 175-196. `http://dx.doi.org/10.1080/1350486X.2016.1226144`.

4. Karlsson, P., Jain, S. and Oosterlee, C.W. (2016). Fast and accurate exercise policies for Bermudan swaptions in the LIBOR market model. *International Journal of Financial Engineering* 3(1): 1650005. `http://dx.doi.org/10.1142/S2424786316500055`.

5. Hofer, M. and Karlsson, P. (2016). Efficient Calibration for CVA using Multi-Level Monte Carlo. Submitted for publication. Available at SSRN: `http://ssrn.com/abstract=2776932`.

6. Jain, S., Karlsson, P. and Kandhai, D. (2016). KVA, Mind your P's and Q's!. Submitted for publication. Available at SSRN: `http://ssrn.com/abstract=2792956`.

7. Karlsson, P. (2014). Finite Element Based Monte Carlo Simulation of Option Prices on Lévy Driven Assets. Submitted for publication. Available at SSRN: `http://ssrn.com/abstract=2721095`.

Acknowledgment

First, I would like to thank my supervisor, Birger Nilsson, who has given me the freedom to realise my own ideas ever since my bachelor's thesis.

My sincere gratitude goes to Björn Hansson and Hans Bystöm, who have done so much for me. My sincere gratitude also goes to Cornelis Oosterlee and Erik Schlögl for their supervision, for giving me the opportunity to join their research, and for opening up so many doors for me.

I want to send a special thanks to my co-authors, Kay Pilz and Shashi Jain. I also want to thank Karl Larsson, Bujar Huskaj, Emanuel Alfranseder, Lech Grzelak, Mia Hinnerich and the former director for PhD studies, Jerk Holm. Erik Lindström and Magnus Wiktorsson, thank you for introducing me to quantitative finance.

Most of all, I would like to thank my Mother, my brothers Mikael and Ricky, and friends Marjan and Mohamed. You made this journey possible.

Patrik Karlsson

INTRODUCTION

Introduction

1 History of Quantitative Finance[1]

Louis Bachelier is considered to be the founding father of mathematical finance. His PhD thesis on Brownian motion, modelling stocks (assuming a normal distribution), and pricing options, Bachelier (1990), are the main building blocks for modern option pricing theory. However, the world was not ready for his work, which ended up in the shadows.

More than half a century later, Bachelier's work on option pricing was rediscovered and finally brought to light. The major drawback of the Bachelier model was the normality assumption because it allowed for negative stock values. Samuelson (1965) tackled the issue and assumed that underlying returns follow a geometric Brownian motion (log-normal distribution).

The breakthrough then came in Black and Scholes (1973) and Merton (1973). The hedging arguments and the Black-Scholes equation, a partial differential equation (PDE) that described the evolution of option prices over time, was revolutionary. It was founded on the concept that option risk can be fully eliminated by continuously maintaining a hedge consisting of the underlying asset and cash. The Black-Scholes-Merton model laid out a strong foundation for modern option pricing theory. Coupled with the establishment of the Chicago Board Options Exchange (CBOE) in 1973, the first marketplace for trading listed options, the option market expanded globally and everyone traded using the Black-Scholes equation.

The Black-Scholes-Merton model is based on simplistic assumptions such as geometric Brownian motion, constant volatility, deterministic interest rates, and no credit risk. October 19, 1987, is often referred to as Black Monday, the day on which compelling evidence was seen that the log-normal distribution with the constant volatility assumption was too crude because it neglected extreme events that cannot be hedged out. For instance, this phenomenon was illustrated using implied volatility. Prior to Black Monday, the implied volatility of an option on major indices demonstrated an almost flat surface. However, after Black Monday, to compensate for the Black-Scholes limitations, the implied surface became highly skewed, which pushed for more sophisticated models. Important landmarks are jump extension in Merton (1976), stochastic volatility in Hull and White (1987) and Heston (1993), and local volatility in Dupire (1994).

Most all of the work to date focused on equity derivatives; however, with its well-developed theory, modelling interest rate derivatives took off. Major milestones were the one-factor Gaussian mean-reverting models in Vasicek (1977)

---

[1]This section provides a brief overview of important landmarks in the history of option pricing. A more rigorous overview of the option pricing history can for instance be found in Jarrow (2010).

and Cox, Ingersoll and Ross (1985). The concept is to model the short rate (one point on the interest rate yield curve) to price derivatives containing interest rate risk. The drawback of not being able to subsequently match the market yield curve led to the development of the multifactor Heath, Jarrow and Morton (1992) (HJM) framework. Instead of only modelling one point on the yield curve, as is done in a short rate model, HJM enables modelling of the entire yield curve.[2]

HJM defines its dynamics in terms of an infinite set of instantaneous forward rates, which is inconsistent with the actual traded instruments, such as the discrete set of compounded London interbank offered rates (LIBORs). Furthermore, the inconsistency with the market practice of pricing fixed-income derivatives using the Black (1976) formula for vanilla securities was a major drawback.[3] However, the introduction of the LIBOR market model (LMM) in Jamshidian (1997), Miltersen, Sandmann and Sondermann (1997), and Brace, Gatarek and Musiela (1997) solved these issues. The LMM is formulated directly in terms of market observable LIBORs and their correlations and volatilities. The LMM's ability to price vanilla securities (e.g. caps) using Black (1976) as well as securities that rely strongly on correlations between forward rates (e.g. Bermudan swaptions[4]) are reasons for its popularity.

With all of these models in place, a maturity was reached and commodity, equity, foreign exchange (FX), and interest rate risks could finally be managed.[5] Now, the focus could finally be on one last piece: managing default risk. Although credit risk modelling was introduced in Merton (1974), it was too simplistic (it assumed that counterparties only issued one bond). Some remarkable results are in Jarrow and Turnbull (1992, 1995) and Lando (1998), who solved this problem by drawing analogues to the well-developed interest rate modelling.

Standardised options had then become so liquid (often even dominating the underlying asset) that they contained crucial information that could be used (for model calibration) to price complicated over-the-counter (OTC) exotic derivatives. As the models and tools developed, so did the derivatives. Exotic derivatives gained in popularity as investors explored complex payoff opportunities, such as barriers, Bermudans, cliquets, lookbacks, and variance swaps. Even trading in simple OTC derivatives, such as credit default obligations (CDS), collateralise debt obligations (CDOs), and interest rate swaps, exploded. However, the growth of OTC derivatives had a significant impact on the global fi-

---

[2] Spot rate models are a special case of HJM. Hull and White (1990) is an extension that allows a perfect fit to the initial yield curve by imposing a time-dependent drift.

[3] The log-normal formulation for interest rates violated the so-called arbitrage-free argument.

[4] The Bermudan swaption is one of the most traded exotic derivatives and gives the holder the right (but not the obligation) to enter into an underlying swap on a set of pre-specified dates.

[5] Commodities and FX derivatives pricing is based on similar concepts as for equity and interest rate derivatives.

nancial crisis (2007–2009) and its aftermath. Before the crisis, banks engaged in careless behaviour by taking on too much risk (e.g. in credit derivatives). There was also a general view that large companies were *"too-big-to-fail"* and, thus, an overall tendency to underestimate counterparty risk occurred. However, the bankruptcy of AIG and Lehman Brothers in 2008 demonstrated that, instead of being *"too-big-to-fail"*, they were *"too-big-to-be-allowed-to-fail"* (Gregory, 2010, 17).

Subsequently, the market's concern over counterparty risk regarding OTC derivatives increased. The Basel committee on banking supervision formulated regulatory standards for setting up capital requirements to cover for losses in the case of a counterparty default. For instance, the credit valuation adjustment (CVA) requirement – an adjustment to the derivative price to compensate for a possible counterparty default – was introduced. Next to the CVA are a number of related valuation adjustments (XVAs), such as debt valuation adjustments (DVA), funding valuation adjustments (FVA), and capital valuation adjustments (KVA). The importance of these XVAs has grown significantly in the aftermath of the financial crisis and has given rise to new areas within banks, such as the XVA (trading) desk.[6] The XVA desk has the overall responsibility for pricing and hedging XVAs; for example, in the case of a counterparty default, the bank should not take a major loss.

2   The Thesis

Today, a bank's derivatives portfolio is large and complex and requires efficient calibration and pricing methods. From a trading desk perspective, a simple workflow given market data from vanilla products consists of calibrating a number of models, such as the LMM, Hull-White, Heston, or a hybrid model, among others, such that consistent risk aggregation exists between the different asset classes. Next, the bank turns to pricing and hedging derivative portfolios, and having products of different complexity put significant pressure on the implementations. A trader hedging his portfolio cannot wait hours or even minutes for his current portfolio risk numbers to be updated. Instead, this update has to been done in real time or as quickly as possible, such that the market does not move against him when hedging using old data. The XVA trader faces an even worse situation because he is overlooking the bank's entire portfolio of derivatives, which can consist of millions of trades compared with a few thousand or tens of thousands for the desk level trader.

---

[6]Today, major banks have separated the XVA trading desk from the normal trading desks. The XVA desk is often centralised to overlook a bank's entire risk and all portfolios in terms of XVAs. The normal trading desk is divided into subgroups, such as by asset classes, vanilla or exotic products, regions, and others.

This thesis consists of four essays devoted to the most recent topics within quantitative finance and focuses on the calibration and pricing of derivatives.

Paper 1 - Calibrating a Market Model with Stochastic Volatility to Commodity and Interest Rate Risk

In the first essay, we develop a hybrid commodity interest rate market model with stochastic volatility, together with an efficient calibration routine, to be able to aggregate risk between asset classes in a consistent manner. This chapter is of particular interest to a desk level trader, such as a commodity derivatives trader with a portfolio consisting of commodity derivatives that are highly dependent on interest rates (e.g. a Bermudan oil option) and who seeks a model that incorporates both interest rates and commodity risk, to be able to hedge out variations in the underlying assets.

Paper 2 - Fast and Accurate Exercise Policies for Bermudan Swaptions in the LIBOR Market Model

In the second essay, we construct an efficient Monte Carlo scheme to price Bermudan swaptions in the LMM. This chapter is of particular interest to a desk level trader, such as an exotic interest rate derivatives trader who needs an efficient and accurate method for pricing and hedging Bermudan swaptions.

Paper 3 - Counterparty Credit Exposures for Interest Rate Derivatives using the Stochastic Grid Bundling Method

In the third essay, we construct an efficient Monte Carlo scheme to calculate credit exposures on interest rate derivatives, such as on a portfolio consisting of Bermudan swaptions. In particular, we study CVA. This chapter is of particular interest to an XVA trader who needs a fast but also accurate method to be able to hedge variations in XVA charges.

Paper 4 - Finite Element Based Monte Carlo Simulation of Option Prices on Lévy Driven Assets

In the fourth and final essay, we demonstrate a new technique for simulating option prices on a class of popular jump models by combining both Monte Carlo and numerical PDE methods. This chapter could be of particular interest to an FX trader who needs a method to price FX derivatives (that typically have a big smile).

PAPER 1

Calibrating a Market Model with Stochastic Volatility to Commodity and Interest Rate Risk

with Kay Pilz and Erik Schlögl

Abstract

Using the multi-currency LIBOR Market Model (LMM), this paper constructs a hybrid commodity interest rate market model with a stochastic local volatility function that allows the model to simultaneously fit the implied volatility surfaces of commodity and interest rate options. Because liquid market prices are only available for options on commodity futures and not on forwards, a convexity correction formula for the model is derived to account for the difference between forward and futures prices. A procedure for efficiently calibrating the model to interest rate and commodity volatility smiles is constructed. Finally, the model is fitted to an exogenously given correlation structure between forward interest rates and commodity prices (cross-correlation). When calibrating to options on forwards (rather than futures), the fitting of cross-correlation preserves the (separate) calibration in the two markets (interest rate and commodity options), whereas in the case of futures, a (rapidly converging) iterative fitting procedure is presented. The cross-correlation fitting is reduced to finding an optimal rotation of volatility vectors, which is shown to be an appropriately modified version of the "orthonormal Procrustes" problem in linear algebra. The calibration approach is demonstrated in an application of market data for oil futures.

1   Introduction

Modelling market risks to price derivative financial instruments has come a long way since the seminal paper of Black and Scholes (1973). In particular, it is widely recognised that such models need to be calibrated to all available liquid market prices, including options of various strikes and maturities, for all relevant sources of risk. For commodity derivatives, the approach presented in this paper represents a step closer to this ideal.

In addition to commodity prices and their stochastic dynamics, the valuation and risk management of positions in commodity derivatives also depend on market interest rates and the stochastic dynamics thereof. The market instruments to which a model should be calibrated include the swaption "cube" (swaptions of (1) various maturities, (2) various strikes, on (3) swaps of various lengths) and commodity options of various maturities and strikes. For commodities, futures are more liquid than forwards. Consequently (as well as to make the model more realistic), the correlation between commodity prices and interest rates becomes a relevant model input already at the calibration level.

The model presented in this paper, with its associated calibration method, is fitted to market prices for swaptions in a swaption cube, options on commodity futures for various maturities and strikes, and – of course – the underlying futures and interest rate term structures. Furthermore, it is fitted to exogenously estimated correlations between interest rates and commodity prices. The construction is based on a LMM[1] for the interest rate and commodity markets. The two markets linked in a manner analogous to the construction of the multicurrency LMM,[2] where the convenience yield takes the role of the interest rate in the commodity market (thus, convenience yields are assumed to be stochastic). To allow a fit to market-implied volatility smiles (and skews) of commodity and interest rate options, the model is equipped with a stochastic local volatility function (SLV), following Joshi and Rebonato (2003), Andersen and Brotherton-Ratcliffe (2005), and Piterbarg (2005a,b).[3]

Efficient calibration is achieved in two steps. The first step is to separately calibrate the model to the interest rate market, building on a synthesis of the calibration approaches for the LMM in Pedersen (1998) and the SLV-LMM in Piterbarg (2005a,b). Then, to be able to calibrate efficiently to commodity futures, we consider two approximations for calculating the difference between futures and forwards in the proposed model. The separate calibration in the interest rate market in the first step is then followed by an iterative, two-stage calibration to the commodity market. An orthonormal transformation of the commodity volatility vectors is applied in the second stage, rotating the commodity

---

[1] See the seminal papers by Miltersen, Sandmann and Sondermann (1997), Brace, Gatarek and Musiela (1997) and Jamshidian (1997).

[2] See Schlögl (2002b).

[3] Grzelak and Oosterlee (2011a) presented an extensions of Schlögl (2002b) with stochastic volatility.

volatilities relative to the interest rate volatilities in such a manner that achieves the desired correlations between the two markets. The calibration of this orthonormal transformation to the desired cross-correlations is cast in terms of a modified orthonormal Procrustes problem, permitting an effective solution algorithm to be applied. We illustrate the use of the model on real market data.

For the example, we chose U.S. Dollars (USD) as the "domestic" currency and Brent Crude Oil as the commodity (the "foreign currency"). The "exchange rate" is given by the Brent Crude Oil futures prices denoted in USD. These prices, when converted to forward prices using an appropriate convexity correction, can be interpreted as forward exchange rates between the USD economy and an economy in which value is measured in terms of units of Brent Crude Oil (where convenience yields are interpreted as the foreign interest rates). In the example, the model is calibrated to the USD swaptions volatility cube and the volatility smile of European-style options on Brent Crude Oil futures.

Hybrid modelling combining commodity and interest rate risk was initiated by Schwartz (1982), who modelled interest rate risk via the stochastic dynamics of the continuously compounded short rate without reference to a full model calibrated to an initial term structure. Subsequently, a number of authors proposed models for stochastic convenience yields, some of whom also incorporated the stochastic dynamics of the term structure of interest rates.[4] In these models, continuously compounded convenience yields (and possibly interest rates) typically are normally distributed because they are assumed to be driven by a Heath, Jarrow and Morton (1992) term structure model with generalised (possibly multi-factor) Ornstein/Uhlenbeck dynamics. In such a model, effective calibration to available commodity and interest rate options is difficult when only at-the-money options are considered and is not possible for the full range of available strikes. At-the-money calibration is a strength of log-normal LIBOR Market Models, and Pilz and Schlögl (2013) construct a hybrid model that exploits this and uses an orthonormal rotation of volatility vectors to fit the cross-correlations between the commodity and interest rate markets. By lifting the log-normality assumption, this paper goes beyond their work to allow calibration to the full swaption cube and commodity volatility surface. This paper also refines the correlation fitting procedure through rotation by casting it as a modification of the orthonormal Procrustes problem,[5] which can be solved by a fast numerical algorithm.[6]

From a practical perspective, whether it is worthwhile to move beyond simple models (e.g. in the present context, pricing commodity derivatives using an adaption of the Black and Scholes (1973) model in which the underlying

---

[4]See, for example, Gibson and Schwartz (1990), Cortazar and Schwartz (1994), Schwartz (1997), Miltersen and Schwartz (1998), and Miltersen (2003).

[5]See Golub and Van Loan (1996).

[6]This algorithm is given as Algorithm 8.1 in Gower and Dijksterhuis (2004).

commodity follows a geometric Brownian motion and interest rates are deter-
ministic) depends on whether the more sophisticated model produces substan-
tially different derivative prices (and hedges) when compared with a simpler
benchmark. Relative to a model along the lines of Black and Scholes (1973) or
Black (1976), a hybrid model of stochastic commodity prices and interest rates
based on a stochastic local volatility LMM represents extensions in two dimen-
sions: fitting the implied volatility surface observed in the market and integrat-
ing commodity and interest rate risk. Regarding the necessity of the former,
little debate exists among practitioners. The primary purpose of the calibrated
models is to price illiquid derivative products in a manner consistent with ob-
served prices for (typically simpler) products that liquidly trade in the market.
Thus, soon after implied volatility smiles (contradicting the Black/Scholes as-
sumption of geometric Brownian motion) appeared in the market, practitioners
recognised the need for models consistent with this observation.[7] For the latter,
the impact of adding stochastic interest rates to a (stochastic volatility) com-
modity model has been studied recently in a series of papers[8] by Cheng, Niki-
topoulos, and Schlögl (2016a,b,c,d). Cheng, Nikitopoulos and Schlögl (2016a)
find a noticeable impact of interest rate volatility and correlation between the
interest rate process and the (commodity) futures price process on the prices
of long-dated futures options, with this impact becoming less pronounced for
shorter maturities. Looking at the problem from a more traditional academic
perspective, Cheng, Nikitopoulos and Schlögl (2016b) find that allowing for
stochastic interest rates improves the out-of-sample empirical performance of
their model. The impact on the hedging effectiveness (as opposed to pricing
performance) of incorporating interest rate risk into a commodity derivatives
model is studied in Cheng, Nikitopoulos and Schlögl (2016c) and (2016d). The
first paper considers this question in a simulated model, and the second pa-
per conducts a back-test of hedging performance on empirical data. The latter
study finds that, in times of market turbulence such as, in particular, during
the Global Financial Crisis of 2007–2009, augmenting a commodity delta hedge
with an interest rate hedge consistently improves hedge performance – more
consistently than augmenting the delta hedge by vega or gamma. Thus, a hybrid
model provides benefits for the pricing and risk management of vanilla prod-
ucts, and this effect is more pronounced for more exotic products that explicitly
condition jointly on interest rate and commodity risk.

   As previously noted, the model considered in the present paper is assem-
bled from individual components based on the stochastic local volatility for-
mulation of the LMM by Piterbarg (2005a,b). This represents one major strand
of the literature that extends the LMM beyond at-the-money calibration. The

---

[7]This is reported, for example, by Derman (2003).

[8]The stochastic interest rate dynamics used in these papers are simpler than the SLV-LMM dynam-
   ics considered here, but nevertheless demonstrate the relevance of incorporating interest rate risk
   into a commodity derivatives model.

other major strand is based on the SABR model of Hagan et al. (2002), which is laid out in detail in Rebonato, McKay and White (2009). In the present context, the choice of the SLV–LMM over the SABR–LMM as the basis of the hybrid interest rate/commodity model is driven by three considerations.

1. Specification of the stochastic volatility dynamics along the lines of Piterbarg (2005a,b) avoids the mathematical problems associated with SABR,[9] i.e. the log-normality of the volatility process and the undesirable behaviour of the stochastic differential equation (SDE) of the underlying for certain values of the "constant elasticity of variance" (CEV) parameter β. The former implies the divergence to infinity of volatility almost surely in finite time. The latter involves non-uniqueness of the solution to the SDE (for $0 < \beta < \frac{1}{2}$) and/or the process of the underlying financial variable being absorbed at zero (for $0 < \beta < 1$).

2. The SLV–LMM is directly amenable to calibration by an appropriately modified Pedersen (1998) algorithm, in which we directly and exogenously control the correlation structure of the underlying financial variables (as opposed to the correlation structure of the driving Brownian motions). In particular, in the absence of liquid market instruments containing useable information on "implied" correlations,[10] correlation is subject to considerable "parameter uncertainty." Having directly interpretable correlation inputs assists in controlling for this source of "model risk."

3. In the present paper, the SLV–LMM for each market (interest rates and the commodity) is driven by a vector of independent Brownian motions. Thus, correlations are introduced by the manner in which volatility is distributed over these Brownian motions ("factors") by the vector-valued volatility functions. This permits a two-stage procedure of fitting cross-correlations between markets after fitting the models for the individual markets (although some coupling occurs when calibrating to futures) using orthonormal transformations.

   The paper is organised as follows. The basic notation, the results of the single– and multi-currency LMM and their interpretation in the context of commodities are presented in Section 2. In Section 3 the calibration of the commodity part of the Commodity LMM to plain vanilla options is discussed. In Section 4 the relationship between futures and forwards in the model is presented, which permits calibration of the model to futures as well as forwards.

---

[9]See Section 3.10 of Rebonato, McKay and White (2009).

[10]Although swaption prices depend in theory on correlations between forward rates, in practice this dependence is too weak for these correlations to be extracted in a meaningful way; see e.g. Choy, Dun and Schlögl (2004).

The calibration of the interest rate part of the hybrid Commodity LMM will not be discussed in detail in this paper, because this problem has already been addressed by many authors, e.g., Piterbarg (2005a,b), and most methods should be compatible with our model. However, in Section 5 we discuss how both separately calibrated parts of the model – the interest rate and the commodity part – can be merged in order to have one underlying d-dimensional Brownian motion for the joint model and still match the market prices used for calibration of the particular parts. Section 6 illustrates the application of the model to real market data.

2   The Commodity LIBOR Market Model

2.1   The LIBOR Market Model

For the construction of the LMM for the domestic interest rate market we assume a given probability space $(\Omega, \mathscr{F}, \mathbb{P})$, where the underlying filtration $\{\mathscr{F}_t, t \in [0, T_N]\}$ coincides with the $\mathbb{P}$-augmentation of the natural filtration of a d-dimensional standard Brownian motion W, and $\mathbb{E}_t^{\mathbb{P}}[\cdot] := \mathbb{E}^{\mathbb{P}}[\cdot | \mathscr{F}_t]$ denote the conditional expectation on the information at time t. Let T_N be a fixed time horizon, $P(t, T)$, $T \in [t, T_N]$ the bond price, i.e. the amount that has to be invested at time t to receive one unit of the domestic currency at time T, hence $P(T, T) = 1$ for every $T \in [0, T_N]$. Assuming the discrete-tenor structure, $0 = T_0 < T_1 < \ldots < T_N$, with intervals $\tau_n = T_{n+1} - T_n$, the forward LIBOR rate $L(t, T_n)$ with fixing period T_n as seen at time t is given by

$$L(t, T_n) = \tau_n^{-1} \left(\frac{P(t, T_n)}{P(t, T_{n+1})} - 1 \right), \quad q(t) \leq n \leq N - 1,$$

where $q(t)$ is the index function of the LIBOR rate with the shortest maturity not fixed at time t, defined as $T_{q(t)-1} \leq t < T_{q(t)}$. The price of the discounted bond maturing at time $T_n > t$ is then given by

$$P(t, T_n) = P\left(t, T_{q(t)}\right) \prod_{i=q(t)}^{n-1} \frac{1}{1 + \tau_i L(t, T_i)}.$$

The dynamics of the forward LIBOR rate $L(t, T_n)$ as seen at time $t \in [0, T]$,

under the $\mathbb{P}^{T_{n+1}}$-forward measure[11] is given by

$$dL(t, T_n) \quad = \quad \sigma_L(t, T_n)^\top dW^{T_{n+1}}(t), \tag{1}$$

where $\sigma_L(t, T_n)$ is a d-dimensional process, discussed later in this section. From Girsanov's theorem, the dynamics of $L(t, T_n)$ are

$$dL(t, T_n) = \sigma_L(t, T_n)^\top \left(\gamma_L(t, T_n)\, dt + dW^{T_n}(t)\right),$$

where W^{T_n} is a d-dimensional vector–valued Brownian motion[12] under the $\mathbb{P}^{T_n}$-forward measure and γ_L is determined by the volatility of the forward bond price process, i.e.

$$d\left(\frac{P(t, T_n)}{P(t, T_{n+1})}\right) = \frac{P(t, T_n)}{P(t, T_{n+1})}\gamma_L(t, T_n)dW^{T_{n+1}}(t) \qquad \text{with}$$

$$\gamma_L(t, T_n) = \frac{\tau_n \sigma_L(t, T_n)}{1 + \tau_n L(t, T_n)}, \tag{2}$$

relates $dW^{T_{n+1}}$ to dW^{T_n} by,

$$dW^{T_n}(t) = dW^{T_{n+1}}(t) - \gamma_L(t, T_n)\, dt. \tag{3}$$

Further results and the connection of this model to the framework of Heath, Jarrow and Morton (1992) can be found in the original LMM literature, commencing with Miltersen, Sandmann and Sondermann (1997), Brace, Gatarek and Musiela (1997), Jamshidian (1997) and Musiela and Rutkowski (1997a).

2.1.1   The Stochastic Local Volatility LMM

For most markets, implied volatilities calculated from traded option prices are strike dependent, i.e. exhibit a volatility smile and skew (slope of the at-the-money volatility). To capture the skew, we assume that the time-dependent volatility functions are of the separable form

$$\sigma_L(t, T_n) = \varphi_L\left(L(t, T_n)\right)\lambda_L(t, T_n), \tag{4}$$

---

[11] This forward measure is the equivalent martingale measure associated with taking the zero coupon bond $P(t, T_{n+1})$ as the numeraire, and under this measure (the existence of which is assured under the model assumptions below) forward LIBOR $L(t, T_n)$ is necessarily a martingale, i.e. driftless — see e.g. Musiela and Rutkowski (1997b).

[12] Thus W^{T_n} is a d-dimensional vector, each component $W_i^{T_n}$, $1 \le i \le d$, is a Brownian motion under the $\mathbb{P}^{T_n}$-forward measure, and the quadratic covariation between the components is zero: $dW_i^{T_n} dW_j^{T_n} = 0 \; \forall i \ne j$.

where $\lambda_L(t, T_n)$ is a bounded deterministic d-dimensional function and $\varphi_L :$ $\mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a time homogenous local volatility function. This is a fairly general setup, and one model allowing for skewed implied volatility is the displaced–diffusion model along the lines of Joshi and Rebonato (2003), where, following the notation of Andersen and Piterbarg (2010) (see their Remark 7.2.13), φ_L is given by

$$\varphi_L(L(t, T_n)) = b_L(t, T_n) L(t, T_n) + (1 - b_L(t, T_n)) L(0, T_n).$$

When $(1 - b_L(t, T_n))/b_L(t, T_n) < (L(0, T_n)\tau_n)^{-1}$, the existence of path wise unique solutions follow (Andersen and Piterbarg, 2010, Lemma 14.2.5).

To capture the volatility smile, we follow Andersen and Brotherton-Ratcliffe (2005) and scale the Brownian motions with a mean-reverting stochastic volatility process given by,

$$dz_L(t) \quad = \quad \theta\left(z_{L,0} - z_L(t)\right)dt + \eta\sqrt{z_L(t)}dZ_L(t), \tag{5}$$

where θ and η are positive constants, $z_L(0) = z_{L,0} = 1$ and Z_L is a Brownian motion under the spot measure $\mathbb{Q}^B$.[13] The quadratic covariation of Z_L and each component of W is assumed to be zero. Assuming the LIBOR dynamics in (1) with the separable volatility function in (4) and stochastic volatility (5), the stochastic local volatility LIBOR market model (SLV–LMM) specifies the dynamics of the forward LIBOR rates for $n = 1, \ldots, N$ by

$$dL(t, T_n) \quad = \quad \sqrt{z_L(t)}\varphi_L(L(t, T_n))\lambda_L^\top(t, T_n)dW^{T_{n+1}}(t). \tag{6}$$

When calibrated to interest rate option market data, the model matches at-the-money volatilities through λ_L, the skews (slope of the Black/Scholes implied volatilities) through b_L and curvatures of the volatility smiles through the volatility of variance η. The speed of mean reversion κ determines how fast the spot volatility converges to the forward volatility, or more specifically, how fast $z_L(t)$ is pulled back to its long–term mean level $z_{L,0}$.

The relationship between spot measure $\mathbb{Q}^B$ and forward measures is given by standard results for the LMM dating back to Jamshidian (1997) and Brace, Gatarek and Musiela (1997):

$$dW^{T_{n+1}}(t) \quad = \quad \sqrt{z_L(t)}\mu_n(t)dt + dW^B(t), \tag{7}$$

$$\mu_n(t) \quad = \quad \sum_{j=q(t)}^{n} \frac{\tau_j\varphi_L\left(L(t, T_j)\right)}{1 + \tau_j L(t, T_j)}\lambda_L(t, T_j). \tag{8}$$

Moreover, we assume that the Brownian motion $Z_L(t)$ of the variance pro-

---

[13]See for instance Section 4.2.3 of Andersen and Piterbarg (2010).

cess $z_L(t)$ is independent of the d-dimensional Brownian motion W and that all forward LIBORs and factors are driven by the same scaling $\sqrt{z_L(t)}$.

2.2   The Commodity Market

The approach incorporating a commodity market corresponds largely to the approach as described for the log-normal case in Section 2.2 of Pilz and Schlögl (2013), which in turn is based on the multi-currency extension of LIBOR market models introduced in Schlögl (2002b). Therefore, we focus on the aspects related to the stochastic local volatility extensions of the model. The setup for the commodity parallels the one for interest rates in the previous section, and the corresponding volatility functions are denoted by σ_F and γ_F. As explained in Pilz and Schlögl (2013), the commodity market can be seen as a "foreign interest market" with the commodity (e.g. crude oil) as currency. "Foreign bond prices" $C(t, T)$ can be interpreted as "convenience yield discount factors" for the commodity, defined as the amount of the commodity today which is equivalent to the discounted (using domestic interest rates) value of receiving one unit of the commodity (e.g. one barrel of crude oil) at time T, taking into account any storage costs and convenience yields.[14] The same logic as for the domestic interest rate market can be used to derive "forward rates" for the commodity market, but since such "convenience yield instruments" are not traded for commodities, we construct the model by specifying domestic interest rate dynamics on the one hand, and the dynamics of forward commodity prices on the other hand. Then, as noted in Schlögl (2002b), this implicitly determines the "foreign interest rate dynamics," i.e. the convenience yield dynamics in the present interpretation.

As in Schlögl (2002b), the existence of a spot price process $S(t)$ for the commodity is assumed, denoted in the local currency (e.g. USD per barrel crude oil). Then, its forward value is given by

$$F(t, T_n) = \frac{C(t, T_n)S(t)}{P(t, T_n)},\tag{9}$$

for all $n = 0, \ldots, N$. We assume the same tenor structure $\tau_n = T_{n+1} - T_n$ for interest and commodity markets. If this assumption needs to be lifted in order to reflect market reality, an interpolation on either of the forwards can be applied. Since LIBORs have typically 3-month or 6-month tenors, and exchange traded futures[15] often have expiries with 1-month or 3-month time difference, the interpolation has to be made for forward interest rates in most cases. See for instance Schlögl (2002a) on forward interest rate interpolation.

The forwards in (9) are necessarily martingales under the $\mathbb{P}^{T_n}$-forward mea-

---

[14] Thus, the $C(t, T)$ represent the effect of the convenience yield net of storage cost.

[15] The futures versus forward relation will be discussed in Section 4.

sure, i.e.

$$dF(t, T_n) = \sigma_F^\top(t, T_n)\, dW^{T_n}(t),\tag{10}$$

for all $n = 0, \ldots, N$. To account for a stochastic local volatility dynamics for the commodity prices, the dynamics for the commodity forward prices are set to

$$dF(t, T_n) = \sqrt{z_F(t)}\varphi_F(F(t, T_n))\,\lambda_F(t, T_n)\, dW^{T_n}(t),\tag{11}$$

where

$$dz_F(t) = \theta_F\left(z_{F_0} - z_F(t)\right)dt + \eta_F\sqrt{z_F(t)}dZ_F(t),\tag{12}$$

$$\varphi_F(F(t, T_n)) = b_F(t, T_n)F(t, T_n) + (1 - b_F(t, T_n))F(0, T_n),\tag{13}$$

and θ_F, η_F positive constants, $z_F(0) = z_{F_0} = 1$ and $b_F(t, T)$ a deterministic function mapping from $\mathbb{R}^+ \to \mathbb{R}^+$. The Brownian motions $W^{T_{n+1}}$, for $n = 0, \ldots, N-1$, are the same as in (1), $Z_F(t)$ is a Brownian motions under the spot measure $\mathbb{Q}^B$, and there is no correlation between the underlying drivers and volatility drivers in the sense that for all $n = 0, \ldots, N-1$

$$dW_i^{T_{n+1}}dZ_L(t) = dW_i^{T_{n+1}}dZ_F(t) = dZ_F(t)dZ_L(t) = 0 \qquad \forall 1 \le i \le d.$$

Note that the structure of the dynamics for the commodity forwards $F(t, T_n)$ is the same as for the interest forward rates $L(t, T_n)$, except that they are martingales under different forward measures.

As demonstrated in Schlögl (2002b), this fully specifies the hybrid model: Denote by $\gamma_F(t, T_n)$ the volatility of the quotient $C(t, T_n)/C(t, T_{n+1})$ of convenience yield discount factors, then $\gamma_F(t, T_n)$ is determined by the no-arbitrage relation of the multi-currency LMM,[16]

$$\sigma_F(t, T_n) = \gamma_F(t, T_n) - \gamma_L(t, T_n) + \sigma_F(t, T_{n+1}).\tag{14}$$

where $\gamma_L(t, T_n)$ is defined by (2) and $\sigma_F(t, T_n)$ (and analogously $\sigma_F(t, T_{n+1})$) by (10) above.

In addition to the no-arbitrage condition (14), cross-correlations specify linkages between the interest rates and commodity forwards markets. Their form and calibration will be discussed in detail in Section 5.

3   Calibration with Time Dependent Parameters

This section discusses aspects of the calibration of the hybrid model, which consists in our approach of two parts. The first part calibrates the (LIBOR) in-

---

[16]See Equation (11) in Schlögl (2002b).

terest rate forward market and the commodity market separately to their market instruments. The second part merges the two separate calibrations with due regard to cross-correlations and the no-arbitrage condition. Readers who have their own preferred individual calibration routines for the stochastic local volatility LMM of Section 2.1.1, as well as for the stochastic local volatility commodity model of Section 2.2, may skip this and the following section and continue directly with Section 5. As mentioned in the introduction, we focus on the calibration of the commodity leg, since a calibration of the LIBOR market model in the context of stochastic local volatility has already been addressed by many other authors, for instance Joshi and Rebonato (2003) and Piterbarg (2005a).

Since the Commodity LMM is based on commodity forwards, we have to calibrate to forward implied volatilities or plain vanilla option prices written on forwards. However, commodities futures rather than forwards are most liquidly traded (consider, for example, the Brent Crude Oil futures in the market data example in Section 6) and thus forward prices have to be deduced from futures prices. As we are working in a hybrid model that is integrating commodity and interest rate risk, it is not adequate to equate forward prices with futures prices, as is still common among practitioners. Section 4 describes how to take into account the distinction between futures and forwards when applying the calibration methods proposed in the present section.

The calibration of the model to commodity forward instruments follows the ideas of Piterbarg (2005a,b) and is split into two parts. First, a pre-calibration is performed to determine a globally constant speed of mean reversion θ_F and volatility of variance η_F such that the volatility smile given from market option quotes is matched as closely as possible. Second, the volatility term structure λ_F and the volatility skew structure b_F are fitted to option prices.

3.1 Step 1 – Calibrating the level of mean reversion and volatility of variance

To obtain an efficient calibration algorithm we follow Piterbarg (2005a) and project the full dynamics of the commodity forwards $F(\cdot, T_n)$ in Equation (11) with time-dependent parameters onto a model with constant parameters using the parameter averaging technique. Formally, the SDE with time-dependent parameters is replaced by an SDE with constant parameters for each maturity, where both have the same marginal distribution. These parameters are called effective parameters, and let $\overline{\lambda}_{F,n}$ denote the effective volatility and $\overline{b}_{F,n}$ denote the effective skew, for all maturity times T_n. The dynamics of $F(t, T_n)$ is then given by

$$dF(t, T_n) \quad = \quad \sqrt{z_F(t)}\Big(\overline{b}_{F,n}F(t, T_n) + (1 - \overline{b}_{F,n})F(0, T_n))\Big)\overline{\lambda}_{F,n}\, dW^{T_n}(t). \quad (15)$$

We assume to have forward processes $F(\cdot, T_1), \ldots, F(\cdot, T_N)$ with expiries $T_1, \ldots, T_N$ and we further think of T_0 as "now." Times-to-maturity for an arbitrary calendar time $t \geq 0$ are given by $x_n = T_n - t$ for $n = 0, 1, \ldots, N$. For the commodity calibration, market prices for call options[17] on $F(\cdot, T_n)$, with payoff $(F(T_n, T_n) - K_i)^+$, and for several strikes K_i, $i = 1, \ldots, k_n$, are assumed to be available and are denoted by $C_{n,i}^{\mathrm{mkt}}$.

Vanilla options on (15) can be calculated efficiently by the Fourier method.[18] We denote the resulting model call prices by $C_{n,i}^{\mathrm{mod}}$.

The calibration problem of the first step is then to find parameters θ_F, η_F and $\bar{b}_{F,n}$, $\bar{\lambda}_{F,n}$ for $n = 1, \ldots, N$, such that

$$\sum_{\substack{n = 1, \ldots, N, \\ i = 1, \ldots, k_n}} \left(C_{n,i}^{\mathrm{mod}} - C_{n,i}^{\mathrm{mkt}} \right)^2 \longrightarrow \min.$$

For the global parameters θ_F for mean reversion level of the variance process, and η_F for volatility of variance, this optimisation yields their final values in the calibration. The term structure and skew parameters, $\bar{b}_{F,n}$ and $\bar{\lambda}_{F,n}$ respectively, will be adapted in the next step to fit the market as closely as possible.

3.2   Step 2 – Calibrating the Volatility Term- and Skew-Structure

The calibration of the volatility term structure follows the approach given in Pilz and Schlögl (2013), but we include the calibration of the skew structure into this procedure.

The term structure of volatility levels is assumed to be piecewise constant for a specified grid of calendar times $t_i = (0, t_1, \ldots, t_{n_c})$ and times to maturity $x_j = (x_0, x_1, \ldots, x_{n_f})$, which defines a $(n_c \times n_f)$ matrix of volatilities $V = (v_{i,j})_{1 \leq i \leq n_c, 1 \leq j \leq n_f}$. The relation of the d-dimensional model volatility vectors $\lambda(t, T_n)$ (for all n) and matrix V is given by

$$\| \lambda(t, T_n) \| = \sum_{\substack{0 \leq i \leq n_c - 1 \\ 0 \leq j \leq n_f - 1}} \mathbb{1}_{\{t_i \leq t < t_{i+1}, \, x_j \leq T_n - t < x_{j+1}\}} v_{i,j}. \tag{16}$$

The number of forward times n_f in the volatility matrix do not need to coincide with the number of traded forwards N, and especially in regions of large

---

[17] For notational simplicity, we assume that the option expires at the same time the futures does. In most cases the option expires a few days before the futures expiry. In some cases, like for EUA carbon emission futures, the option can even expire several months before the underlying futures.

[18] For details, see e.g. Andersen and Piterbarg (2010) (Chapter 9).

forward times a rougher spacing can be chosen for x_j, since volatilities tend to flatten out with increasing forward time. To be able to price options on all of the forwards, the maturity of the longest available forward has to be smaller or equal to the latest calendar time and the longest time to maturity, $T_N \leq \min\{t_{n_c}, x_{n_f}\}$.

We refer to Pilz and Schlögl (2013) for a more detailed discussion on the setup with piecewise constant volatilities, and how to compute total variances efficiently for given calendar and forward times.

In the next section we will use the correlations to obtain a map from the volatility levels $\|\lambda_F(t, T_n)\|$ to the components of the volatility vectors $\lambda_F(t, T_n)$ that are multiplied by the d-dimensional Brownian motion in (11).

In a manner analogous to the volatility levels, we define a $(n_c \times n_f)$-dimensional matrix $B = (b_{i,j})_{1 \leq i \leq n_c, 1 \leq j \leq n_f}$ for the matrix of piecewise constant skews. To keep notation simple, we use the same grid as for the volatility term structure. The entry $b_{i,j}$ represents the skew corresponding to forward $F(t, T_n)$ with $t_{i-1} \leq t < t_i$ and $x_{j-1} \leq T_n - t < x_j$.

The optimisation of Step 2 is defined with respect to a set of calibration criteria. The first calibration criterion measures the *quality of fit* and is as in Step 1 defined by the sum of squared differences between market and model prices,

$$q \quad = \sum_{\substack{n = 1, \ldots, N, \\ i = 1, \ldots, k_n}} \left(C_{n,i}^{\mathrm{mod}} - C_{n,i}^{\mathrm{mkt}}\right)^2. \tag{17}$$

Since the number of parameters is potentially quite large (and larger than the number of market prices in (17)), we follow Pedersen (1998) and specify for the volatility two *smoothness criteria* s_λ given by,

$$s_\lambda \quad = \quad \eta_{\lambda,1} \sum_{j=1}^{n_f} \sum_{i=1}^{n_c-1} (v_{i+1,j} - v_{i,j})^2 + \eta_{\lambda,2} \sum_{i=1}^{n_c} \sum_{j=1}^{n_f-1} (v_{i,j+1} - v_{i,j})^2. \tag{18}$$

The first term measures departures from time–homogeneity; it demands that volatilities with different calendar times but the same time to maturity do not deviate from each other too much. The second term, the forward time smoothness, forces the volatility term structure to be smooth in time to maturity for each fixed calendar time. The larger the weight $\eta_{\lambda,1}$, the more volatility and skew become (calendar) time homogeneous. The larger the weights $\eta_{\lambda,2}$, the flatter the volatility and skew becomes in forward time direction.

We specify an analogous smoothness function for the skew term structure

s_b with corresponding weights $\eta_{b,1}$ and $\eta_{b,2}$

$$s_b \quad = \quad \eta_{b,1} \sum_{j=1}^{n_f} \sum_{i=1}^{n_c-1} (b_{i+1,j} - b_{i,j})^2 + \eta_{b,2} \sum_{i=1}^{n_c} \sum_{j=1}^{n_f-1} (b_{i,j+1} - b_{i,j})^2. \qquad (19)$$

The smile contribution to the implied volatility coming from the parameters fixed in the first calibration step of Section 3.1 is unaffected by the smoothness criteria here.

Remark 1 *Although the number of parameters is potentially quite large, the optimisation, for instance using a Levenberg–Marquardt approach for minimising the objective value $q + s_\lambda + s_b$, usually gives stable calibration results, since the smoothness criteria force the parameters to a non-parametric but structured form.*

Remark 2 *The reason why we have not used the effective parameters $\bar{b}_{F,n}$ and $\bar{\lambda}_{F,n}$ from the global calibration of 3.1 as target values in our second calibration 3.2, as for example proposed in Piterbarg (2005a), is that we here assumed the more complex case of calibrating to options on futures. This requires to compute the convexity correction in each optimisation step, which changes the relation between options prices and the effective volatility and skew parameters. Therefore, the advantage of using pre-computed effective parameters for calibration is not applicable in our general case.*

Remark 3 *As pointed out by Andersen and Piterbarg (2010), the degree of freedom is potentially quite large here, and obtaining the volatility and skew term-structure simultaneously is computationally inefficient. One could assume the skew b_F and volatility λ_F to be two almost orthogonal problems (changing the volatility has an minor impact on the skew and vice versa) and solve for them separately. Step 2 can therefore be divided into two parts. First, solve for the skew term structure by fixing the volatility parameters, e.g., to the ones obtained in the first step, $\bar{\lambda}_{F,n}$, together with η and θ. And since the time-dependent skew can be solved for explicitly as in Piterbarg (2005a), one can target the implied skews from Step 1 instead of option prices and reduce the computational time significantly by avoiding Fourier pricing within each iteration. Then, given the skew term-structure and the parameters η and θ, the volatility term structure is calibrated. However, as mentioned in Remark 2, due to the futures-forward convexity correction we cannot target the implied volatilities directly here.*

3.3   Volatility Factor Decomposition

The method for constructing a map from $V = \|\lambda_F(t, T_n)\|$ in Equation (16) to $\lambda_F(t, T_n)$ via volatility factor decomposition or PCA described in Pilz and Schlögl (2013) remains applicable in the present setting, and we provide a brief summary for the reader's convenience. Due to the dependence of the convexity correction for futures contracts on interest rate dynamics (see Sections 4 and 5, below), the volatility factor decomposition has to be included in the calibration process when calibrating to futures and options on futures. As in Pedersen (1998), there is a separate factor decomposition for every calendar time step to be covered by the calibrated volatility function, thus in the following we fix this to an arbitrary t_i. Let v_i denote the ith row (corresponding to calendar time t_i) of V, written as column vector. The matrix C of commodity forward correlations is exogenously given and assumed to be constant over calendar time. For each t_i, the covariance matrix then is calculated by

$$\Sigma = (v_i v_i^\top) \odot C, \tag{20}$$

where $\odot$ means component–wise multiplication (Hadamard product). Decomposing

$$\Sigma = RD^{1/2}(RD^{1/2})^\top$$

results in orthonormal eigenvectors of Σ in the columns of $R = (r_{j,k})_{1 \le j, k \le n_f}$ and the corresponding eigenvalues on the diagonal of the matrix $D = (\xi_{j,k})_{1 \le j, k \le n_f}$. Choosing the largest d eigenvalues, R and D can be reduced to matrices $R \in \mathbb{R}^{n_f \times d}$ and $D \in \mathbb{R}^{d \times d}$ by retaining only the corresponding d columns in R and $d \times d$ sub-matrix in D. For the factor–reduced volatilities we then have

$$v_{i,j}^2 \quad = \quad \sum_{k=1}^d r_{j,k}^2 \xi_k \quad (j = 1, \dots, n_f). \tag{21}$$

and the stepwise constant volatility function for the forwards is $\lambda_{ijk} = r_{j,k}\sqrt{\xi_k}$.

Remark 4 *Under the convexity adjustments suggested in Section 4, the forward returns correlation matrix C in (20) is also the correlation matrix of futures returns.*

4   Futures/Forward Relation and Convexity Correction

The calibration method in Section 3 is applicable only when forwards and options on forwards are available. This section presents a model–consistent ap-

proximate conversion of futures prices to forward prices for all relevant data for calibration, in order to apply the methods of the previous section.

We introduce the notation $G(t, T)$ for a futures price at time t with maturity T, and, as before, $F(t, T)$ will be the corresponding forward price. From no-arbitrage theory we know $F(T, T) = G(T, T)$ and that prices of plain vanilla options on forwards and futures must coincide, whenever the maturities of option, forward and futures are the same. This allows us to use the call prices of options on futures for calibration of forwards, and we only have to assure that the (virtual) forwards have the same maturities as the futures. Due to equation (10) the forward $F(\cdot, T_n)$ is an exponential martingale under the T_n-forward measure.

Denoting by $\mathbb{E}^B$ the expectation under the spot risk–neutral measure $\mathbb{Q}^B$, futures follow the general relation

$$G(t, T) = \mathbb{E}_t^B [S(T)], \tag{22}$$

see Cox, Ingersoll and Ross (1981), where $S(t)$ is the spot price, which satisfies by no-arbitrage constraints $S(t) = F(t, t) = G(t, t)$ for all t. Integrating (11) and using (7)-(8) gives,

$$
\begin{aligned}
F(T_n, T_n) &= F(t, T_n) + \int_t^{T_n} \sqrt{z_F(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) dW^{T_n}(s) \\
&= F(t, T_n) + \int_t^{T_n} \sqrt{z_F(s)} \sqrt{z_L(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) \mu_{n-1}(s) ds \\
&\quad + \int_t^{T_n} \sqrt{z_F(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) dW^{\mathbb{B}}(s).
\end{aligned}
$$

Putting these relations together and taking the $\mathbb{Q}^B$-expectation for the futures, as in (22), results in

$$
\begin{aligned}
G(t, T_n) &= \mathbb{E}_t^B [F(T_n, T_n)] \\
&= F(t, T_n) + \mathbb{E}_t^B \left[\int_t^{T_n} \sqrt{z_F(s)} \sqrt{z_L(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) \mu_{n-1}(s) ds \right] \\
&\quad + \mathbb{E}_t^B \left[\int_t^{T_n} \sqrt{z_F(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) dW^{\mathbb{B}}(s) \right] \\
&= F(t, T_n) + \mathbb{E}_t^B \left[\int_t^{T_n} \sqrt{z_F(s)} \sqrt{z_L(s)} \varphi_F(F(s, T_n)) \lambda_F^\top(s, T_n) \mu_{n-1}(s) ds \right] \\
&\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad (23) \\
&= F(t, T_n) + D(t, T_n), \tag{24}
\end{aligned}
$$

the third equality follows since the last expectation is that of an Itô integral, further letting $D(t, T_n)$ denote the convexity correction.

From here we consider two alternative ways to proceed.

4.1 Approximation 1 - Freeze all risk factors

The first crude approximation would be to freeze all random variables by setting $\overline{L}_i = L(0, T_i)$, $\overline{F}_j = F(0, T_j)$ and $\bar{z}_L = z_L(t) = 1$, $\bar{z}_F = z_F(t) = 1$ such that the convexity correction $D(t, T_n)$ can be approximated by,

$$D(t, T_n) \approx \int_t^{T_n} \varphi_F(\overline{F}_n) \sum_{j=q(s)}^{n-1} \frac{\tau_j \varphi_L(\overline{L}_j)}{1 + \tau_j \overline{L}_j} \lambda_F^\top(s, T_n) \lambda_L(s, T_j) ds. \tag{25}$$

Pilz and Schlögl (2013) demonstrate how the integrals in Equation (25) can be computed when the volatility functions are piecewise constant. Since $\overline{F}_n$ is unknown, it is necessary to solve equation (24) for $F(t, T_n)$ using the approximation (25) above. This is straightforward for displaced diffusions with $\varphi_F(\overline{F}_n)$ piecewise constant subject to the integration variable s.

4.2 Approximation 2 - Freeze LIBORs and Commodity forwards

For the second approximation we choose to freeze the LIBORs and commodity forwards but keep the volatility stochastic. Set $\overline{L}_i = L(0, T_i)$ and $\overline{F}_j = F(0, T_j)$ and use conditioning for the stochastic volatility processes. If the expectation and the integration can be interchanged it follows from the independence of the stochastic volatility processes that the convexity correction $D(t, T_n)$ can be approximated by,

$$D(t, T_n) \approx \int_t^{T_n} \varphi_F(\overline{F}_n) \sum_{j=q(s)}^{n-1} \frac{\tau_j \varphi_L(\overline{L}_j)}{1 + \tau_j \overline{L}_j} \lambda_F^\top(s, T_n) \lambda_L(s, T_j) \mathbb{E}_t^B \big[\sqrt{z_F(s)}\big] \mathbb{E}_t^B \big[\sqrt{z_L(s)}\big] ds. \tag{26}$$

Grzelak and Oosterlee (2011b) show that the first moments of the squared volatility process can be represented as

$$\mathbb{E}_t^B \big[\sqrt{z(T)}\big] = \sqrt{\frac{2e^{-\theta(T-t)}}{n(t, T)}} e^{-\frac{n(t,T)}{2}} \sum_{j=0}^{\infty} \frac{(n(t, T)/2)^j}{j!} \frac{\Gamma\big(d/2 + j + 1/2\big)}{\Gamma\big(d/2 + j\big)}, \tag{27}$$

where $\Gamma(x)$ is the Gamma function, and the parameters n and d are given by

$$n(t, T) = \frac{4\theta e^{-\theta(T-t)}}{\eta^2(1 - e^{-\theta(T-t)})}, \quad d = \frac{4\theta z_0}{\eta^2}. \tag{28}$$

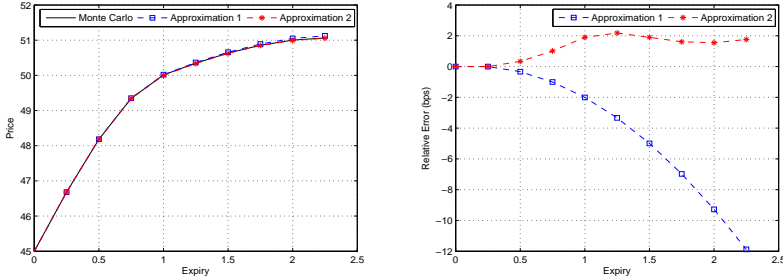We illustrate the accuracy of the two approximations in Figure 1. Clearly

Figure 1: Left: Futures prices from Monte Carlo simulation, Approximation 1 and Approximation 2. Right: Relative error in basis points (bps). With $\lambda_L = 0.20$, $\lambda_F = 0.30$, $b_L = 0.9$, $b_F = 0.8$, $\theta_L = \theta_F = 1$, $\eta_L = 0.5$, $\eta_F = 0.4$, $\text{Corr}\big(dL\big(t, T_j\big), dL\big(t, T_k\big)\big) = \text{Corr}\big(dF\big(t, T_j\big), dF\big(t, T_k\big)\big) = 0.9$.

the stochastic volatility factors have too much impact on the futures–forward convexity to be neglected. As demonstrated, Approximation 2 performs exceptionally well and produces an error less than 2 basis points, which is more than acceptable within a calibration routine.

5 Merging Interest Rate and Commodity Calibrations

The calibrations for the interest rate market and the commodity market have so far been considered separately, with exception of the convexity correction in Section 4. This section focuses on linking the calibrations of the two asset classes in order to get a joint commodity and LMM calibration. The relation between these two asset classes is determined by their cross-correlation matrix, which is assumed to be constant over time. Before addressing the problem of merging the calibrations in Section 5.2, we briefly discuss how to generate the cross-correlation matrix from historical data in Section 5.1.

5.1 The Exogenous Cross-Correlations

The cross–correlations between commodity forwards and interest rate forwards are much less pronounced than the correlations between the futures and forwards, respectively, within the asset classes. This makes the estimation for cross-correlations from historical data less stable than for the correlations

within the asset classes. For example, the structure of the cross-correlation matrix between Brent Crude Oil futures and USD interest rate forwards in Figure 8 of Section 6 can hardly be explained by obvious rationales. Therefore, one might decide to specify exogenously a very simple cross-correlation structure in practice. The simplest case would be to assign a single value to all cross-correlation in the matrix R_{LF}^{exogen}. If this seems too crude, a linear relationship can be fitted by regressing the empirical cross-correlations for the forward interest rates and commodity futures on a 2-dimensional plane.

5.2   The Cross-Calibration

The quadratic (cross-)covariation process for commodity forwards with maturity T_m and forward interest rates with settlement T_n is given by the dynamics

$$dL(t, T_n)dF(t, T_m) = \left(\sqrt{z_L(t)}\varphi_L(L(t, T_n))\lambda_L(t, T_n)\right)^{\top}\left(\sqrt{z_F(t)}\varphi_F(F(t, T_m))\lambda_F(t, T_m)\right)dt,$$

where we have used the fact that

$$dW^{T_{n+1}}(t) = dW^{T_m}(t) + \sqrt{z_L(t)}\left(\mu_n(t) - \mu_{m-1}(t)\right)dt,$$

and $dt\,dW^{T_m} = 0$. Similar to the pure interest forward rate correlations, the cross-correlations of the increments are given by

$$\text{Corr}\left(dL(t, T_n), dF(t, T_m)\right) = \frac{\lambda_L(t, T_n)\,\lambda_F^{\top}(t, T_m)}{\|\,\lambda_L(t, T_n)\,\|\,\|\,\lambda_F(t, T_m)\,\|}, \tag{29}$$

since only the factorised λ_F and λ_L are (column-)vectors, and all other values in the equation above are scalars. Abbreviate

$$\Lambda_L = (\lambda_L^{\top}(t, T_n))_{1\leq n\leq N} = (\lambda_{L;n,k}(t))_{1\leq n\leq N, 1\leq k\leq d}, \tag{30}$$

$$\Lambda_F = (\lambda_F^{\top}(t, T_m))_{1\leq m\leq M} = (\lambda_{F;m,k}(t))_{1\leq m\leq M, 1\leq k\leq d}, \tag{31}$$

$$U_{LF} = \left(\|\,\lambda_L(t, T_n)\,\|\,\|\,\lambda_F(t, T_m)\,\|\right)_{1\leq n\leq N, 1\leq m\leq M}, \tag{32}$$

where m and n are indices for different expiries and forward (settlement) times, respectively, and k is the index relating to the stochastic factor. We skip the calendar time dependence, since it could be considered as fixed throughout the rest of this subsection. Then, the model cross-correlation in (29) can be written as

$$R_{LF}^{\text{model}} = (\Lambda_L\Lambda_F^{\top}) \oslash U_{LF}, \tag{33}$$

where $\oslash$ denotes the element-wise division. Merging the individual interest rate and commodity calibrations in order to get a joint calibration means to modify the calibrated matrices Λ_L and Λ_F such that equation (33) is matched (maybe only as closely as possible) for an exogenously given correlation matrix R_{LF}^{exogen},

$$R_{LF}^{\text{exogen}} \approx (\Lambda_L \Lambda_F^{\top}) \oslash U_{LF} = R_{LF}^{\text{model}}.$$

Note that the skew parameter, being part of $\varphi_L(\cdot)$ and $\varphi_C(\cdot)$, as well as the smile parameter, being part of the specification of the square root processes $z_F(t)$ and $z_L(t)$, need not to be modified, as they do not determine the cross-correlations, as equation (29) demonstrates.

Clearly, there are many possible approaches to improve the relation $R_{LF}^{\text{exogen}} \approx R_{LF}^{\text{model}}$. The approach we propose in this section will be guided by the following. Firstly, the quality of the individual calibrations to market instruments relating to one asset class only – interest rates or commodities – is more important that the cross-correlation fit. This is because these market instruments are much more liquidly traded, and therefore provide reliable information. Secondly, the adjustment step needs to be carried out in an efficient way, having a common fit criterion for the approximation.

An approach that combines these considerations is to find a rotation matrix Q, that is applied to one of the calibrated matrices, say Λ_F, with the objective (following Rebonato and Jäckel (2000) in our choice of metric) to minimise the Frobenius distance between the model and the exogenous covariance matrices. Formally, a matrix Q is sought, such that

$$\| R_{LF}^{\text{exogen}} \odot U_{LF} - \Lambda_L(\Lambda_F Q)^{\top} \|_F \longrightarrow \min, \quad \text{subject to } QQ^{\top} = I_d, \qquad (34)$$

with I_d the $(d \times d)$ identity matrix and $\| \, . \, \|_F$ the Frobenius norm. From the theory of normal distributions it follows that an orthonormal rotation changes the cross-covariances, but not the covariance matrices $\Sigma_L = \Lambda_L \Lambda_L^{\top}$ and $\Sigma_F = \Lambda_F \Lambda_F^{\top}$, hence, the individual calibrations remain unaffected.

For the problem of (34) it is quite reasonable to assume $d \leq N$, i.e., the number of stochastic factors has to be equal or less than the number of forward interest rates.[19]

The problem of finding the Q satisfying (34) is similar to the so-called "orthonormal Procrustes" problem: For given matrices $A, B \in \mathbb{R}^{m \times p}$, find an orthonormal matrix $Q \in \mathbb{R}^{p \times p}$ that minimises the distance between A and B,

$$\| A - BQ \|_F \longrightarrow \min, \quad \text{such that } QQ^{\top} = I_d, \qquad (35)$$

---

[19]Note that the roles of Λ_L and Λ_F can be interchanged, and the sufficient assumption actually is $d \leq \max\{M, N\}$. Furnishing the model with more stochastic factors than $\max\{M, N\}$ contributes only spurious complexity to the model. From a practical point of view, the aim is to keep the number of stochastic factors small.

In Golub and Van Loan (1996, Section 12.4.1) it is shown that the solution is given by $Q = UV^\top$, where U and V result from the singular value decomposition (SVD) of $B^\top A$, i.e. $B^\top A = UDV^\top$.

Unfortunately, (34) is more complicated than the Procrustes problem in (35), because in our case Q comes under a transposition. As stated in Gower and Dijksterhuis (2004), Section 8.3.3, "there seems to be no algebraic solution to the problem" (34), to which they refer as a "scaled orthonormal Procrustes" problem. However, the authors discuss a numerical solution, based on an algorithm of Koschat and Swayne (1991), that works well in our case, as we will demonstrate in Section 6 below.

Algorithm 1 Orthogonal scaled procrustes

1: **procedure** OSPROCRUSTES(Q_0, X_1, S, X_2, TOL, MAXITER)
2: $\quad r \leftarrow \| X_1 \|_F$
3: $\quad \text{iter} \leftarrow 0$
4: $\quad \delta \leftarrow 2 \cdot \text{TOL}$
5: $\quad Q \leftarrow Q_0$
6: $\quad$ **while** $\delta > \text{TOL}$ **and** iter $< \text{MAXITER}$ **do**
7: $\quad\quad Z \leftarrow S(X_2^\top X_1 + S^\top Q^\top (rI - X_1^\top X_1))$
8: $\quad\quad$ Compute singular value decomposition $Z = UDV^\top$
9: $\quad\quad Q_{\text{new}} \leftarrow VU^\top$
10: $\quad\quad \delta \leftarrow |\, \| X_1 QS - X_2 \|_F - \| X_1 Q_{\text{new}} S - X_2 \|_F \,|$
11: $\quad\quad \text{iter} \leftarrow \text{iter} + 1$
12: $\quad\quad Q \leftarrow Q_{\text{new}}$
13: $\quad$ **end while**
14: $\quad$ **return** Q
15: **end procedure**

Proposition 1 *The problem* (34) *is equivalent to the problem of finding a T satisfying*

$$\| D_L T D_F^\top - U_L^\top \Sigma_{LF} U_F \|_F \longrightarrow \min, \quad \text{such that } T T^\top = I_d, \tag{36}$$

where the matrices used are from the singular value decompositions[20]

$$\Lambda_L = U_L D_L V_L^\top \qquad U_L \in \mathbb{R}^{N \times N}, D_L \in \mathbb{R}^{N \times d}, V_L \in \mathbb{R}^{d \times d}$$
$$\Lambda_F = U_F D_F V_F^\top \qquad U_F \in \mathbb{R}^{M \times M}, D_F \in \mathbb{R}^{M \times d}, V_F \in \mathbb{R}^{d \times d},$$

and

$$\Sigma_{LF} := R_{LF}^{\text{exogen}} \odot U_{LF}.$$

---

[20] Note that $\text{rank}(\Lambda_L) = d \le N$ and $\text{rank}(\Lambda_F) = d \le M$

The orthonormal solution Q of (34) *is given by*

$$Q = (V_L T V_F^\top)^\top \tag{37}$$

Proof. Substituting the SVD yields

$$
\begin{aligned}
\| \Lambda_L(\Lambda_F Q)^\top - \Sigma_{LF} \|_F = &\| U_L D_L V_L^\top Q^\top V_F D_F^\top U_F^\top - \Sigma_{LF} \|_F \\
= &\| U_L D_L T D_F^\top U_F^\top - \Sigma_{LF} \|_F \qquad T := V_L^\top Q^\top V_F \\
= &\| D_L T D_F^\top U_F^\top - U_L^\top \Sigma_{LF} \|_F \qquad U_L \text{ is orthonormal} \\
= &\| D_L T D_F^\top - U_L^\top \Sigma_{LF} U_F \|_F \qquad U_F \text{ is orthonormal}
\end{aligned}
$$

∎

The solution for T in (36) can be numerically approximated by Algorithm 8.1 in Gower and Dijksterhuis (2004). See also Kercheval (2006) for a nice description of the algorithm with an application in the context of portfolio risk management. For completeness we reproduce the algorithm here (see Algorithm 1).

The objective is to find a Q that minimises

$$\| X_1 Q S - X_2 \|_F, \quad \text{subject to } QQ^\top = I,$$

where S has entries on its diagonal only.

The stopping of the numerical procedure is controlled by the maximal number of iterations (MAXITER) and the tolerance in change of the Frobenius norm (TOL) in the objective above.

As initial guess the trivial transformation $Q_0 = I$ can be used.

5.3   Cross-Calibration Computational Complexity

We conclude the section with a note on the computational effort required by the proposed algorithm. From a theoretical point of view, one would expect the Procrustes method to be faster than the Levenberg–Marquardt optimisation as used in Pilz and Schlögl (2013), because the number of parameters to be optimised is given by $n_{opt} = d(d-1)/2$, which means the dimension of the optimisation problem grows quadratically in the number of stochastic factors d.

The Levenberg-Marquardt optimisation requires at least $O(n^3)$ or $O(mn^2)$ operations in each iteration step (see Nocedal and Wright (2006), Section 10.3), where n is the number of optimisation parameters, and m is the number of residuals in the objective function. In our case we get $O(n_{opt}^3) = O(d^6)$ operations per iteration for the unconstrained optimisation problem. The constraint $QQ^\top = I$ is non-linear, hence cannot be represented by simple bounds for the

parameters, though it is possible to re-parametrise the problem expressing the orthonormal transformation Q by n_{opt} successive rotations, where the angles are the new parameters in an unconstrained Levenberg–Marquardt optimisation, see Anderson, et al. (2005) for a detailed description.

In contrast, the Procrustes problem requires only $O(d^3)$ operations per iteration for the matrix multiplications and the singular value decomposition.

We applied the Procrustes method to several randomly chosen 100×100 cross-correlation matrices, and it converged in all cases within a tolerance of less than 10^{-5} after 200 seconds and about $120,000$ iterations. Our results are in line with the numerical experiments in Kercheval (2006), in which the Procrustes problem turned out to be much faster than the least-squares optimisation using Levenberg–Marquardt. That paper reports that the latter method failed to optimise a 65×65 transformation even "after several hours," whereas the Procrustes method found an approximate solution in "about 5 minutes."

The transformation problems we encounter in the model calibration might be of much smaller size than dimension 100, typically somewhere between $4 - 20$. However, in the case of calibrating the model to commodity futures, the computation of the optimal transformation matrix has to be iterated, which led us to prefer the much more efficient Procrustes algorithm over the "brute force" Levenberg–Marquardt optimisation used in Pilz and Schlögl (2013).

6   Applying the calibration to market data

6.1   Summary of the Calibration Procedure

The following scheme summarises the steps of the calibration procedure, bringing together the steps discussed in the previous sections and assuming the calibration of the (domestic) interest rate LMM has already been carried out.[21]

I. Preliminary calculations applied to the LMM calibration.

---

    1. For each calendar time t_i ($1 \le i \le n_L$):

       (a) Compute the covariance matrix Σ_i^L as in (20).

       (b) Decompose Σ_i^L into λ_i^L using PCA as described in Section 3.

II. Hybrid Calibration.

---

[21] Note that we must iterate over repeated calibration to the commodity market and to the cross–correlations, as the conversion of commodity futures into forwards depends on cross–correlations.

1. Calibrate on all commodity data to obtain the global θ_F, and η_F.
2. Calibrate on all data to obtain the skew term-structure B_F, see Remark 3.
3. For each calendar time t_i ($1 \le i \le n_F$):
   (a) Compute the covariance matrix Σ_i^F as in (20).
   (b) Decompose Σ_i^F into U_i^F using PCA as described in Section 3.
   (c) Compute the rotation matrix Q using Algorithm 1
4. Compute forward prices from futures prices using (26).
5. Compute model options prices on forwards.
6. Continue with Step 3 until a sufficiently close fit to commodity market instruments and assumed cross–correlations is reached.

The dynamics of the thus calibrated hybrid Commodity LMM with SLV can then be written as follows. Let W be a d-dimensional Brownian motion and denote by $\lambda_{i,j}^L$ and $\lambda_{i,j}^F$ the d-dimensional vectors in Λ_L and Λ_F of volatilities for calendar times $t \in [t_{i-1}, t_i)$ and times to maturity $x \in [x_{j-1}, x_j)$. Then, the dynamics of the forward interest rates $L(t, T)$ and the dynamics of the commodity forwards $F(t, T)$ are

$$
\begin{aligned}
dL(t, T_k) &= \sqrt{z_L(t)} \varphi_L \left(L(t, T) \right) \lambda_{i,j}^L dW^{T_{k+1}}(t), \\
dF(t, T_k) &= \sqrt{z_F(t)} \varphi_F \left(F(t, T) \right) \lambda_{i,j}^F dW^{T_k}(t),
\end{aligned}
$$

for all maturity times satisfying $T_k = t_i + x_j$ (for some $1 \le i \le n_c$ and $1 \le j \le m_f$ or $1 \le j \le n_f$, respectively) and all calendar times $t_{i-1} \le t < t_i$.

6.2   Calibration Setup

The Brent Crude Oil is selected as commodity, and the USD forward rate as interest rate.

The calibration date is January 13$^{\text{th}}$, 2015. Figure 2 illustrates the historical Brent crude oil futures and implied volatilities for a selected number of contracts. As one can observe, the futures price (implied volatility) are close to their lowest (highest) levels on the selected calibration date; thus we are using a somewhat "stressed" market scenario to best illustrate the effectiveness of the calibration.

We assume that both the instantaneous volatility functions and skew functions are piecewise constant. The LIBOR forward rates and Brent crude oil futures correlation matrix as used for calibration is historically estimated from the
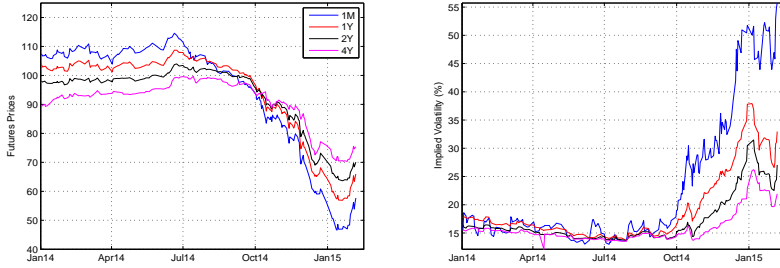
Figure 2: Left: Historical Brent Crude Oil 1M, 1Y, 2Y and 4Y Futures contracts between January 2014 and March 2015. Historical Brent Crude Oil 1M, 1Y, 2Y and 4Y at-the-money implied volatilities between January 2014 and March 2015

time series of LIBOR forward rates and Brent futures prices covering 3 months before the calibration date.

For the correlations and cross-correlations we use a parametric form given for the correlation of the underlying (LIBOR, oil futures) with fixings/maturities T_k and T_j, which is given by[22]

$$\text{Corr}\left(dL\left(t, T_j\right), dL\left(t, T_k\right)\right) = q\left(T_j - t, T_k - t\right),$$

where

$$
\begin{aligned}
q\left(x, y\right) &= \rho_\infty + \left(1 - \rho_\infty\right) \exp\left(-a\left(\min\left(x, y\right)\right)\left|x - y\right|\right), & (38) \\
a\left(z\right) &= a_\infty + \left(a_0 - a_\infty\right) e^{-\kappa z}, & (39)
\end{aligned}
$$

and subject to $0 \le \rho_\infty \le 1$, $a_0, a_\infty, \kappa \ge 0$. The two calibrated correlation matrices for the LIBOR forward rates and the Brent crude oil futures are illustrated in Figure 3.

6.3   The Interest Rate Calibration

The calibration of the SLV–LMM for the USD interest rate market was conducted based on Piterbarg (2005a) by first performing a pre–calibration to fit the SLV–LMM, with constant parameters for each tenor and expiry, to the swaption cube.

---

[22]The origins of this parametric form can be traced back to Rebonato (1999); see Chapter 22 of Rebonato (2004) for a detailed discussion of the rationale behind a parametric form of this type. Our notation follows Andersen and Piterbarg (2010).
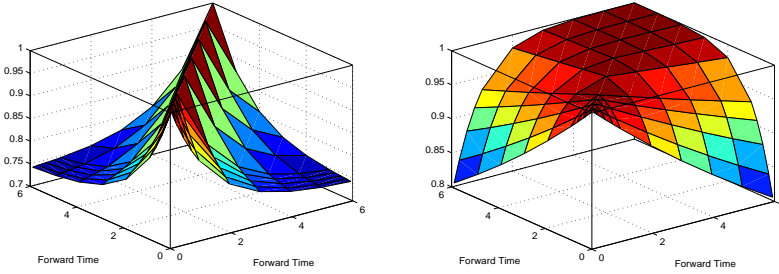
Figure 3: Left: The historically estimated LIBOR forward rate correlation matrix. Right: The historically estimated Brent crude oil futures correlation matrix.
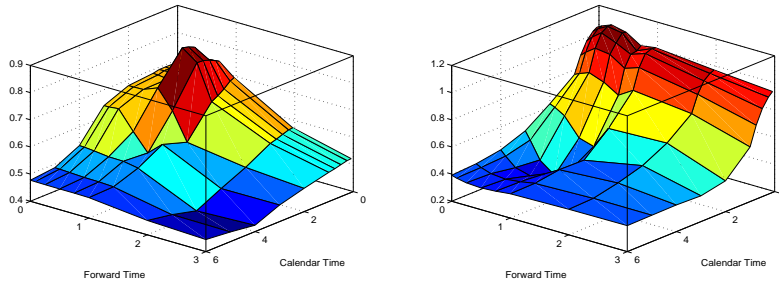

Figure 4: Left: The calibrated implied volatility term structure for the swaption cube. Right: The calibrated implied skew term structure for the swaption cube.

Figure 5 illustrates the obtained effective at–the–money volatility $\bar{\lambda}_{n,m}^{\mathrm{mkt}}$ and effective skew $\bar{b}_{n,m}^{\mathrm{mkt}}$ for each tenor T_n and expiry T_m.

These obtained model parameters then serve as target values in the main calibration when calibrating the volatility and skew term–structure as described in Piterbarg (2005a). Figure 4 shows the resulting volatility and skew term-structure obtained by the main calibration. Figure 5 shows the quality of fit of the effective at–the–money volatility $\bar{\lambda}_{n,m}^{\mathrm{mkt}}$ and effective skew $\bar{b}_{n,m}^{\mathrm{mkt}}$. The overall model fit to the implied volatility is very good as demonstrated by Figure 6 for selected expiries and tenors. Moreover, we chose $d_L = 4$ factors, which typically explain about 99.99% of the overall variance.
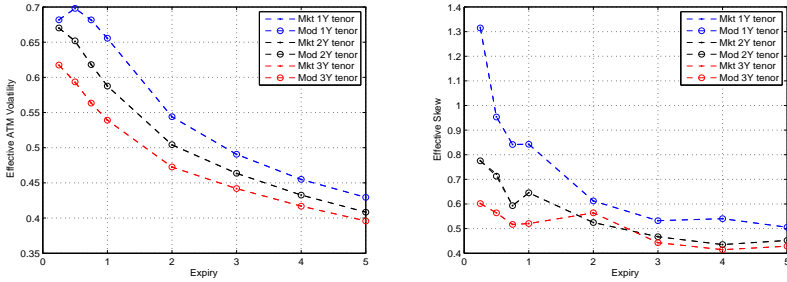
Figure 5: Left: The calibrated effective volatility $\bar{\lambda}_{n,m}^{\mathrm{mkt}}$ for the swaption cube. Right: The calibrated effective skew $\bar{b}_{n,m}^{\mathrm{mkt}}$ for the swaption cube

6.4   The Cross-Calibration

The calibration of the Brent Crude Oil futures was achieved by the method described in Section 3. The market instruments are futures and options on futures. Figure 7 shows the calibrated volatility and skew surface. Calendar and forward times go out to 3 years, and although on the exchange futures with expiries in every month are traded, we chose the calendar and forward time vectors to be unequally spaced (while still calibrating to all traded instruments), [0, 1M, 3M, 6M, 9M, 1Y, 2Y, 3Y]. This setup speeds up the calibration without losing much structure in the volatility surface, since the market views futures with long maturity to have almost flat volatility. For weighting in the calibration objective function we have chosen $\eta_{\lambda,1} = \eta_{b,1} = 0.1$ (time homogeneity, i.e. smoothness in calendar time direction), $\eta_{\lambda,2} = \eta_{b,2} = 0.01$ (smoothness in forward time direction).

The exogenously given target cross-correlation matrix (calculated from the historical time series covering 3 months before the calibration time) and the estimated cross-correlation from historical futures returns and the cross-correlation matrix for the calendar time that fitted worst are illustrated in Figure 8.

Finally, we link both separately calibrated volatility matrices to one set of stochastic factors. Table 1 shows how much of the overall variance, i.e. of the sum of variances over all factors, can be explained by the leading factors, when the factors are sorted according to decreasing contribution to total variance of the commodity forwards. The first two factors already account for more than 99% of the overall variance.

Note that if it is necessary to interpolate the forward interest rate volatility matrix in order to match the calendar times of the commodity volatility matrix, the forward rate covariance matrix will change and, hence, eigenvalue decom-

Number of Factors	1	2	3	4	5	...	19
Percentage of Overall Variance	98.14	99.84	99.96	99.98	99.99	...	100

Table 1: The percentage of overall variance that can be generated by the first i factors, obtained from PCA of the commodity forward covariance matrix for the first calendar time t_1.

positions of the calendar time adjusted covariance matrices yield different results than an eigenvalue decompositions of the original covariance matrices as used for calibration. However, these differences should not be substantial as long as the calibrated volatility matrix is sufficiently smooth in calendar time.

The model fit to the commodity implied volatility is illustrated in Figure 9.

6.5   Pricing Performance and Accuracy

The calibration routine was written in MATLAB and timing the execution was performed on a Intel Core i7-2600K 3.40GHz.

The interest rate volatility and skew term-structure calibration takes roughly 15 seconds, and the cross-correlation calibration (including fitting the commodity volatility) takes just roughly 10 seconds, where the cross-calibration takes about roughly 2 seconds. This would be fast enough for production purposes in most instances, though the fact that this prototype implementation runs in a serial, interpreted language environment leaves scope for substantial speed-ups through through parallelisation and the use of a compiled programming language (such as C++).

We conclude this section by studying the accuracy of the pricing derived in this paper and used within the calibration routine. We focus only on the overall hybrid calibration since Piterbarg (2005b) already performed extensive numerical tests on the accuracy of the parameter averaging approximation for the interest rate calibration itself, which is input to our routine. To summarize Piterbarg (2005b), in terms of (interest rate) Black volatility, the author reported a difference between the Monte Carlo simulation and the semi-analytic approximations for interest rate swaptions, which was within 15 basis points for at-the-money swaptions, and about 20 basis points for in-the-money and out-of-the-money swaptions, for various tenors, expiries and strikes.

For the overall hybrid calibration, we let *Market* refer to the market implied commodity volatilities, *MC* to the implied commodity volatilities from the Monte Carlo simulation of the calibrated model and *Approx* to the implied com-

modity volatilities obtained from the semi-analytic approximations used within the calibration.

Table 2 and Figure 10 provide full a comparison of (correlated hybrid) commodity implied Black volatilities, for the market volatilities, Monte Carlo simulated commodity implied Black volatilities and those obtained from the semi-analytic approximations (used within the calibration), for different expiries and strike offsets. For each expiry 0.16, 0.24, 0.32, 0.58, 0.82, 1.79, 2.28 and 3.79 years,[23] we price the options written on commodity futures with strikes offsets -15%, -7.5%, 0%, +7.5% and +15% (the +15% implied volatility was not available for the 0.16Y option). The model parameters are the ones obtained from the calibration in this section. For the Monte Carlo simulation we use 1,000,000 paths with monthly time steps.

The semi-analytical approximation demonstrates similar accuracy as in Piterbarg (2005b), this can be seen by the reported *MC-Market* error. All reported errors are below 25 bps of Black volatilities. For options with expiry longer than 1Y, the approximation error is about or below 10 bps. As in Piterbarg (2005b) the approximation quality seems to improve for options with longer maturity. This an effect due to the Monte Carlo discretisation error rather than an approximation error, due to the low number of time steps used within the simulation. We noticed that by increasing the number of steps per year, this error decreased.

MC-Market, illustrates the differences in Black volatility between the given market implied volatilities and the ones obtained by the Monte Carlo simulation. This numerical test is performed to demonstrate that given the full calibrated model, which was calibrated using the semi-analytical approximation, we are able to regenerate the market implied volatilities without loss of essential information. The biggest differences are for the short dated options where we observed a difference of 40 bps for the 0.16Y ATM option. As the expiry increases, this error, for all strikes, decreases, where we obtain an error of 20-30 bps for the 0.24Y and 0.32Y expiries, an error of 10-20 bps for the 0.58Y and 0.82Y expires. For the longer expiries this error decreases to about or below 10 bps.

7 Conclusion

As the market data example in the previous section demonstrates, the LMM approach to term structure modelling remains is one of the most flexible for good calibration of the model to market data, even when it is extended to allow for market quotes across multiple strikes (volatility "skews" and "smiles") and the integration of (and correlation between) multiple sources of risk —

---

[23]These are a selection of current expiries, i.e., the time between the calibration date 13 January 2015 and the various expiry dates, for the standardised commodity options.

Expiry	Offset	Market	Approx	MC	Approx-Market	MC-Approx	MC-Market
0.16y	-15%	50.37	50.11	50.36	-0.26	0.25	-0.01
	-7.5%	47.01	47.05	47.30	0.04	0.25	0.29
	0%	44.12	44.27	44.52	0.15	0.25	0.40
	7.5%	41.98	41.72	41.97	-0.26	0.25	-0.01
0.24y	-15%	46.59	46.10	46.35	-0.49	0.25	-0.24
	-7.5%	43.83	43.93	44.18	0.10	0.25	0.35
	0%	41.82	41.97	42.21	0.15	0.24	0.39
	7.5%	40.13	40.20	40.42	0.07	0.22	0.29
	15%	39.13	38.57	38.77	-0.56	0.20	-0.35
0.32y	-15%	44.25	44.22	44.42	-0.04	0.20	0.16
	-7.5%	42.24	42.11	42.32	-0.12	0.21	0.09
	0%	40.35	40.42	40.64	0.07	0.22	0.29
	7.5%	39.21	39.10	39.32	-0.11	0.22	0.12
	15%	38.05	38.09	38.32	0.04	0.23	0.27
0.58y	-15%	39.63	39.60	39.78	-0.03	0.18	0.15
	-7.5%	37.86	37.88	38.04	0.02	0.16	0.19
	0%	36.48	36.45	36.61	-0.04	0.16	0.13
	7.5%	35.23	35.27	35.43	0.04	0.16	0.20
	15%	34.32	34.30	34.46	-0.02	0.16	0.14
0.82y	-15%	37.03	37.02	37.18	0.00	0.16	0.15
	-7.5%	35.54	35.55	35.71	0.01	0.16	0.17
	0%	34.32	34.31	34.47	-0.02	0.16	0.15
	7.5%	33.24	33.26	33.42	0.02	0.16	0.18
	15%	32.41	32.38	32.53	-0.03	0.15	0.11
1.79y	-15%	30.05	30.05	30.17	0.00	0.12	0.12
	-7.5%	28.88	28.88	28.99	0.00	0.11	0.11
	0%	27.93	27.93	28.03	-0.01	0.10	0.09
	7.5%	27.15	27.15	27.26	0.01	0.10	0.11
	15%	26.56	26.54	26.64	-0.02	0.10	0.08
2.28y	-15%	29.12	29.15	29.23	0.03	0.08	0.11
	-7.5%	27.93	27.91	27.99	-0.01	0.08	0.06
	0%	26.97	26.97	27.04	0.00	0.07	0.07
	7.5%	26.25	26.27	26.34	0.02	0.07	0.09
	15%	25.82	25.79	25.86	-0.03	0.07	0.04
3.79y	-15%	25.08	25.07	25.10	-0.02	0.04	0.02
	-7.5%	24.44	24.45	24.48	0.01	0.03	0.04
	0%	23.96	24.00	24.03	0.03	0.03	0.06
	7.5%	23.75	23.68	23.71	-0.07	0.03	-0.04
	15%	23.42	23.49	23.51	0.07	0.03	0.09

Table 2: Pricing results for implied Black commodity volatilities, for commodity options with expiries 0.16, 0.24, 0.32, 0.58, 0.82, 1.79, 2.28 and 3.79 years, and strikes offsets -15%, -7.5%, 0%, +7.5% and +15% . The values are reported in percent (%).

commodity and interest rate risk in our present example. This is due to actual market observables (in particular forward LIBORs) being modelled directly and model prices for calibration instruments (e.g. caps/floors, swaptions, commodity futures and options) being available either in exact or accurate approximate closed form.

The dynamics of all market variables can be expressed in terms of the same, vector–valued Brownian motion and correlation between market variables is obtained via the sum products of the respective vector–valued volatilities. As a consequence, the calibration across multiple sources of risk can be broken down into stages, simplifying the high–dimensional optimisation problems to be solved at each stage. The interest rate market can be calibrated separately using well established procedures. We chose to base our interest rate volatility calibration on the robust and widely used method of Pedersen (1998), but this could easily be replaced by a different method without impacting the remainder of our calibration approach. The model is fitted to an exogenously given correlation structure (typically estimated statistically from historical data) between forward interest rates and commodity prices (cross–correlation) — without impacting the interest rate volatility calibration already obtained — via a modified version of the "orthonormal Procrustes" problem in linear algebra, for which an efficient numerical solution exists.

Finally, it is worth noting that a model which is calibrated to the term structures of commodity futures and options will implicitly reflect any seasonality and/or mean reversion of commodity prices. If seasonality is present and priced by the market, this information will be contained in the initial term structure of commodity futures prices, and integrated into the model by the fact that it fits the initial observed term structure by construction. If mean reversion is present and priced by the market, this information will be contained in the term structure of commodity volatilities, or, more completely, the commodity option implied volatility surface, to which the model is calibrated.[24]

---

[24]Specifically, mean reversion would manifest itself in the market as a downward sloping term structure of commodity option implied volatilities.
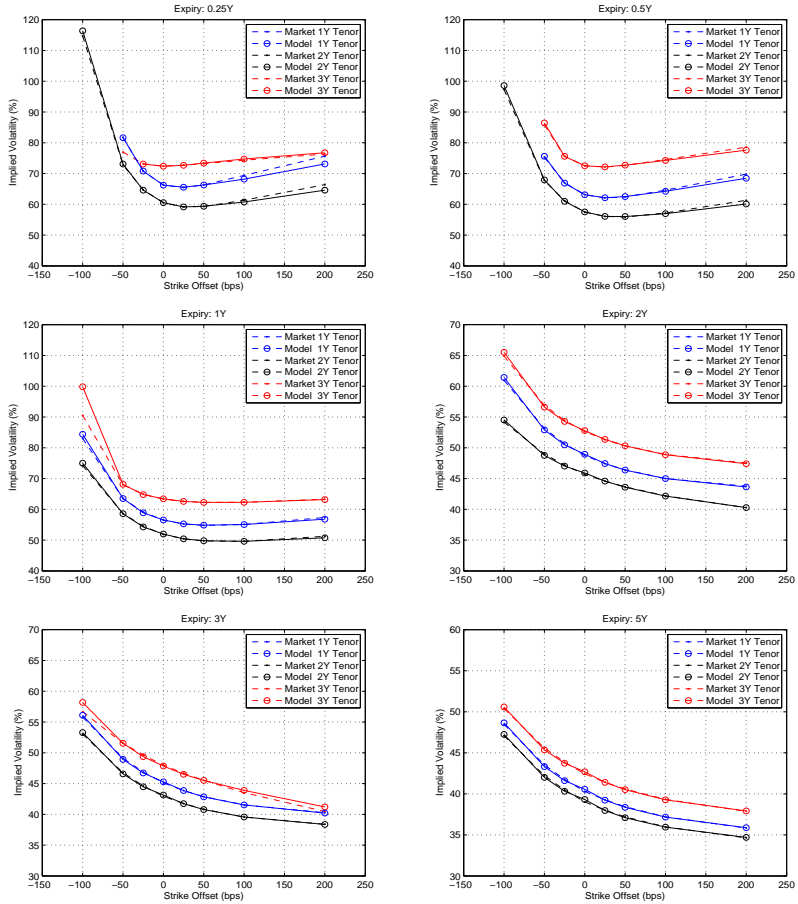
Figure 6: The calibrated interest smile with maturity 0.25, 0.5, 1, 2, 3, 5 years and tenors 1, 2, 3 years.
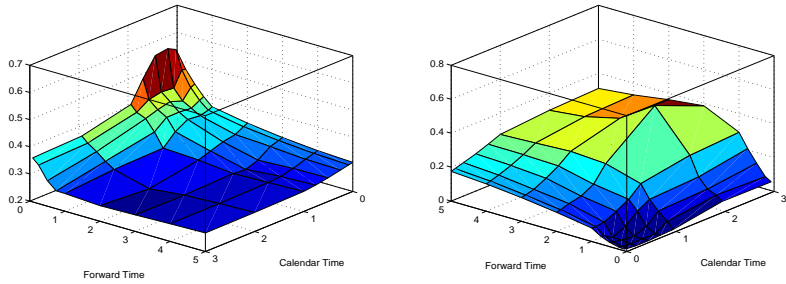
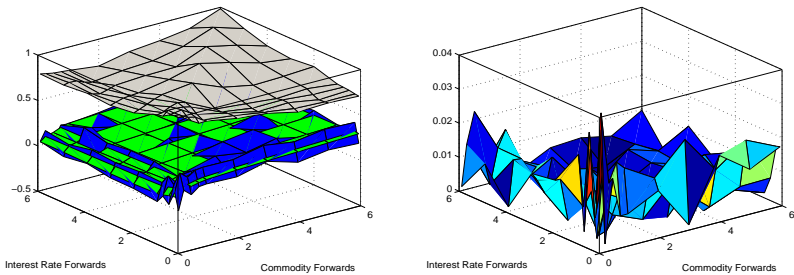Figure 7: Left: The calibrated commodity volatility surface. Right: The calibrated commodity skew surface.

Figure 8: Left: The target cross-correlation matrix (green) estimated from historical futures returns, the rotated cross correlation (blue) and the un-rotated cross correlation (grey) matrix, for the first calendar time. Right: The absolute differences between target and rotated cross-correlations.
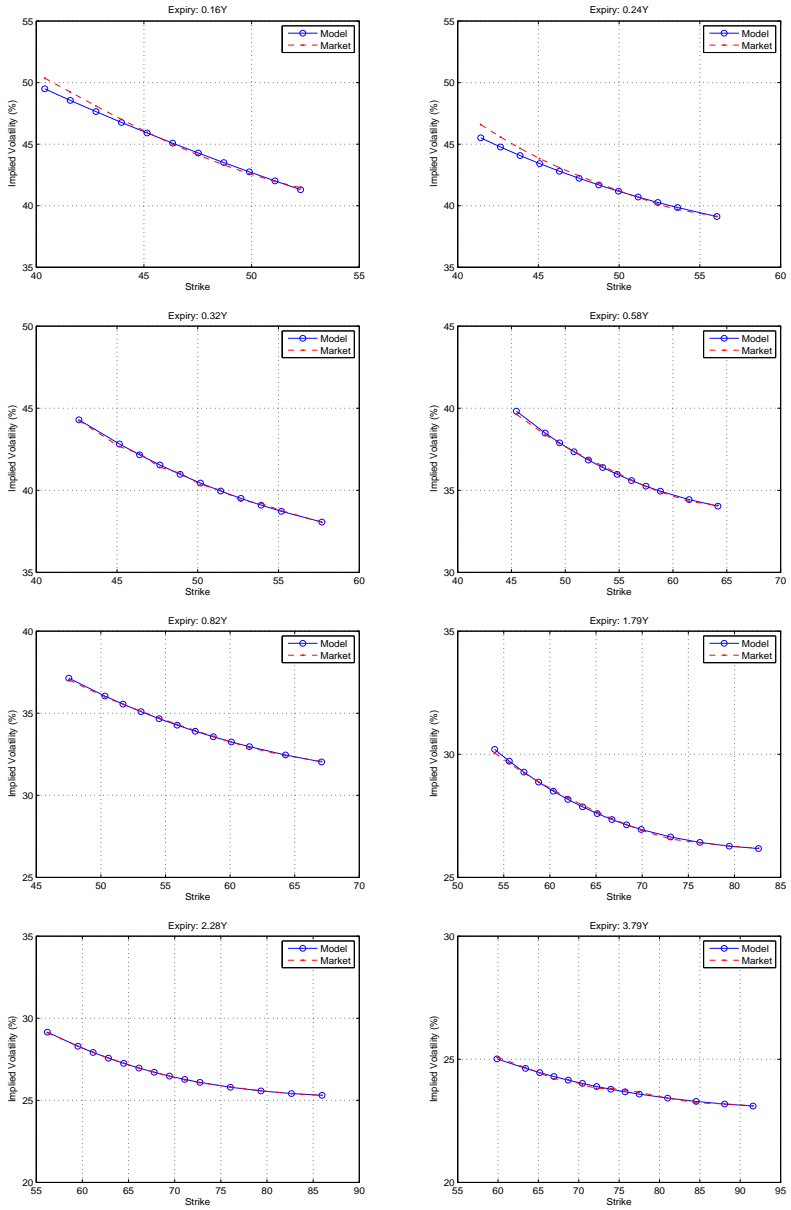
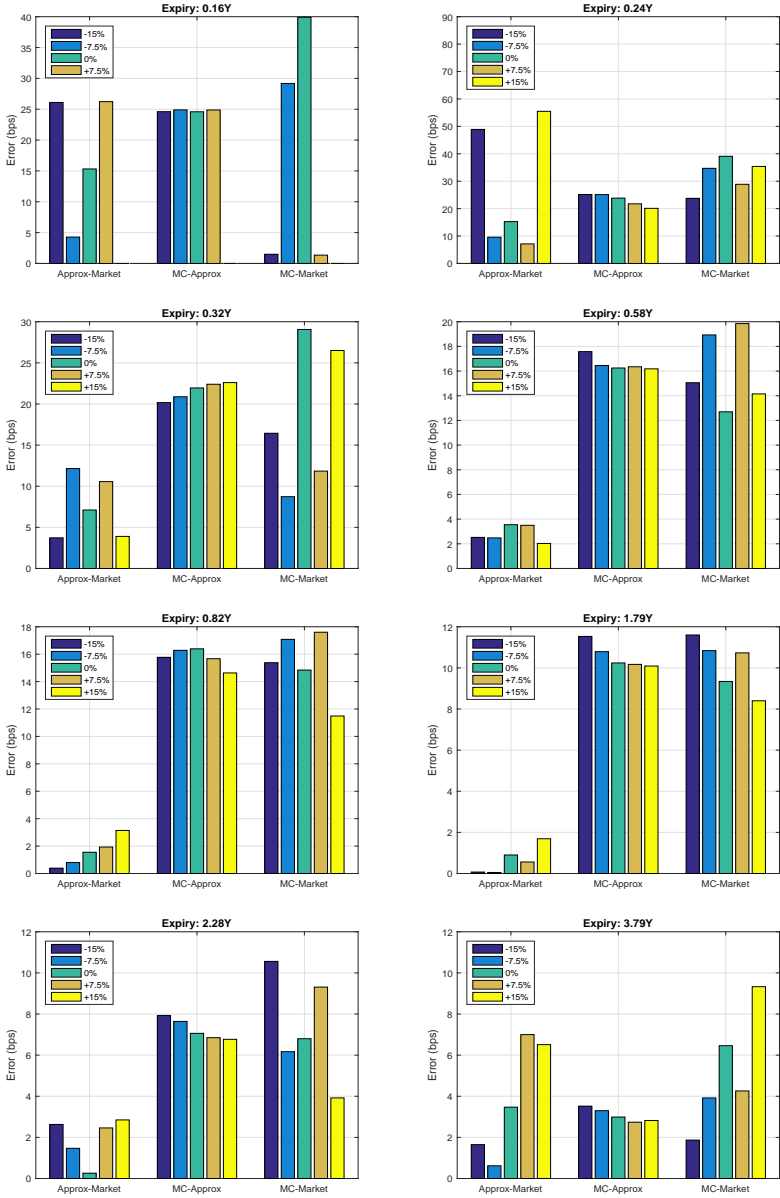Figure 9: The calibrated commodity smile.

Figure 10: Reported errors between the volatilities implied from the Market, the Monte Carlo (MC) simulation and the semi-analytic approximations (Approx) used in within the calibration. The values are reported in basis points (bps).

PAPER 2

Fast and Accurate Exercise Policies for Bermudan Swaptions in the LIBOR Market Model

with Shashi Jain and Cornelis W. Oosterlee

Abstract

This paper describes an efficient American Monte Carlo approach for pricing callable LIBOR Exotics (e.g. Bermudan swaptions) in the LIBOR market model using the Stochastic Grid Bundling Method (SGBM). SGBM is a regression-based Monte Carlo method in which the continuation value is projected onto a space in which the distribution is known. We demonstrate an algorithm to obtain accurate and tight lower–upper bound values without the need for the nested Monte Carlo simulations that are generally required for regression-based methods. The computational results for Bermudan swaptions demonstrate the simplicity and efficiency of the SGBM.

1   Introduction

A Bermudan receiver (payer) swaption (i.e. a Bermudan option on an interest rate swap) is currently one of the most liquid and important exotic derivatives. This swaption gives the owner the right (but not the obligation) to enter into a receiver (payer) interest rate swap at a discrete set of dates (exercise dates). This contrasts with the European swaption, which only can be exercised on a single exercise date, and American options, which can be exercised at any time before maturity.

The LIBOR market model (LMM) is popular for modelling and pricing interest rate derivatives; see, for instance, Miltersen, Sandmann and Sondermann (1997), Brace, Gatarek and Musiela (1997), and Jamshidian (1997). The LMM dynamics are specified as non-overlapping sets of discretely compounded Libor rates. The LMM's consistency with the market practice of pricing fixed-income derivatives allows for pricing to be reduced to standard market formulae such as, for example, the Black and Scholes (1973) formula. Its ability to price securities that rely strongly on correlations between forward rates is a reason for its popularity. Pricing Bermudan swaptions in the LMM is a more complex problem than pricing corresponding European options. First, the holder of a Bermudan swaption is in a position in which, at each exercise date, he needs to determine whether it is optimal to exercise or hold onto the option. Second, given the high dimensionality of LMM,[1] only Monte Carlo methods are feasible for the valuation of exotic fixed-income securities, such as Bermudan swaptions.

Pricing American-style derivatives using Monte Carlo simulation has been actively studied. The industrial standard Longstaff and Schwartz (2001) method (LSM) uses a regression to approximate the continuation value for a set of simulated paths. The fact that LSM is easy to implement and robust, and generates accurate lower bound Bermudan swaptions values for a careful choice of regression variables, are reasons for its popularity. Moreover, lower bounds have been studied in Andersen (2000), in which pre-simulation is performed to estimate a parameterised exercise policy which is then used in a larger simulation for valuing Bermudan swaptions. Generally, American Monte Carlo techniques such as LSM (for lower bound values) are divided into two passes: a first and a second pass. In the first pass, in which the exercise strategy is estimated, conditional discounted option values are projected onto basis functions of the state variables. The projected value is then used as the approximate continuation value, which is compared with the intrinsic value for determining the optimal exercise strategy. This is then followed by a second pass, in which low-biased option values are obtained by simulating a new set of simulation paths and are exercised according to the sub-optimal exercise strategy obtained in the first pass.

---

[1]E.g. pricing a Bermudan swaption on a swap with 10 year maturity and frequency 3 months requires 40 Libor rates.

To validate the pricing models and the lower bound values generated from the second pass, we need a third pass. These are referred to as the upper bound values, and the closer they are to the lower bounds, the better. This phenomenon was previously studied in Rogers (2002), Haugh and Kogan (2004), and Andersen and Broadie (2004), in which the upper bound is approximated using a duality approach. Generally, upper-bounds algorithms such as in, for example, Andersen and Broadie (2004), require nested Monte Carlo simulations, making them computationally expensive. The quality of the upper bound produced by the algorithm depends on the quality of the estimated exercise policy in the first pass, and a better policy provides tighter upper–lower bounds.

The Stochastic Grid Bundling Method (SGBM) was introduced in Jain and Oosterlee (2015) to price equity Bermudan options on geometric Brownian motions. They show that SGBM increases the efficiency of Monte Carlo simulation by reducing the variance of the simulation estimates on the basis of conditional expectations and the use of regressions, as in Milstein and Tretyakov (2009). The method is based on the Stochastic Grid Method (SGM) by Jain and Oosterlee (2012), LSM, and the bundling approach by Tilley (1993). The concept behind SGBM is that neighbouring simulated paths will have similar continuation values for a large set of paths. Therefore, regression can be used to perform local averaging to compute a continuation value for grid points within a bundle. The main difference between LSM and SGBM is that, in SGBM, the option values are projected onto a set of basic functions of the state variables, where the distribution is analytically (or approximately) known. In LSM, the distribution is not taken into account.

This paper is more than a re-interpretation of Jain and Oosterlee (2015). First, because interest rates are stochastic, compared with the fixed interest rates in Jain and Oosterlee (2015), the trivial expectations for the continuation values need to be carefully calculated. One of the interesting aspects of SGBM employed in the present paper is that, because of a formulation in terms of an inner and outer expectation, to calculate the continuation value, we can benefit from the flexibility of using different pricing measures within the same problem. For Bermudan swaptions, we can use the spot measure, which is useful for simulating paths, and the forward measure, which allows the discounting term to be taken out of the expectation, giving rise to an analytic expression for the outer expectation. Second, we also present an efficient way to obtain upper bound values for Bermudan swaptions in LMM by avoiding nested Monte Carlo simulations and, therefore, reducing the required computational time. Third, we demonstrate that SGBM provides more accurate results and is computationally more attractive than LSM.

The paper is organized as follows. Section 2 introduces notations, the general framework and formulates the Bermudan swaption pricing problem in the LMM. Section 3 describes LSM and the SGBM algorithm for pricing Bermudan

swaptions (both a lower and an upper bound method). In Section 4 we present various numerical examples to illustrate the method and finally we conclude in Section 5.

2   Notation and General Framework

In this section, we introduce notation, give a short introduction to the LMM and define the Bermudan swaption pricing. We follow the notation in Andersen and Piterbarg (2010).

The LIBOR Market Model

For the LMM we start with a fixed discrete-tenor structure $0 = T_0 < T_1 < \cdots < T_N$. The intervals over the time horizon are given by $\tau_n = T_{n+1} - T_n$ and are typically three or six calendar months. Let $P(t, T_n)$ denote the time-t price of a zero-coupon bond delivering one unit of currency at some maturity time $T_n \geq t$. The discrete LIBOR forward rate $L_n(t)$ with fixing date T_n as seen at time t is

$$L_n(t) = \tau_n^{-1}\left(\frac{P(t, T_n)}{P(t, T_{n+1})} - 1\right), \quad N - 1 \geq n \geq q(t),$$

where $q(t)$ is the index function of the bond with the shortest maturity, defined as $T_{q(t)-1} \leq t < T_{q(t)}$. The price of the discounted bond maturing at time $T_k > t$ is then given by

$$P(t, T_n) = P\left(t, T_{q(t)}\right) \prod_{n=q(t)}^{n-1} \frac{1}{1 + \tau_n L_n(t)}.$$

For the set of Libor rates $L(t) = \left(L_{q(t)}, L_{q(t)+1}, \ldots, L_{N-1}(t)\right)$ we choose to work under the spot Libor measure, denoted by $\mathbb{Q}^B$, in which the discrete money market account $B(t)$ is the numeraire, given by

$$B(t) = P\left(t, T_{q(t)}\right) \prod_{n=0}^{q(t)-1} (1 + \tau_n L_n(t)).$$

The no-arbitrage dynamics of the forward Libor rates $L_n(t)$ under the spot Libor measure $\mathbb{Q}^B$ for $n \geq q(t)$ are given by

$$dL_n(t) = L_n(t)\lambda_n(t)^\top \left(\mu_n(t)\,dt + dW^B(t)\right), \tag{1}$$

$$\mu_n(t) = \sum_{i=q(t)}^{n} \frac{\tau_i \lambda_i(t)}{1 + \tau_i L_i(t)}, \tag{2}$$

where $W^B(t)$ is an m-dimensional Brownian motion under measure $\mathbb{Q}^B$ and λ_n for $n \geq q(t)$, is a bounded m-dimensional deterministic function. Let $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot|\mathcal{F}_t]$ be denoting the conditional expectation at time t under the spot Libor measure and where $\mathcal{F}_t$ is the filtration at time t generated by W^B. Then by standard arbitrage-free arguments the time-t price of a security paying $V(T)$ at time T is

$$V(t) = \mathbb{E}_t\left[V(T)\frac{B(t)}{B(T)}\right].$$

Further details on the LMM, such as derivations of the bond equations, connection to HJM etc., is out of the scope of this paper and can be found in Andersen and Piterbarg (2010).

Bermudan Swaptions

Given a lockout, i.e., a no-call period up to time T_1, the Bermudan swaption gives the holder the right, but not the obligation, on a set of fixing dates T_n in $\mathcal{T} = \{T_1, T_2, \ldots, T_{m-1}\}$, for $m \leq N-1$, to enter into a fixed for floating swap with fixing date T_n and last payment date T_m. The holder of a payer Bermudan will pay the fixed swap leg and receive the floating swap leg. If exercise at T_n the payout is given by

$$U(T_n) = \phi \mathcal{N} \sum_{i=n}^{m-1} \tau_i P(T_n, T_{i+1})(L_i(T_n) - k),$$

where k is the fixed coupon, $\mathcal{N}$ the notional, and $\phi \in \{-1, +1\}$ is the payer or receiver factor ($+1$ for payer swaption and -1 for a receiver swaption). The payoff is also equivalent to

$$U(T_n) = \phi \mathcal{N} A_{n,m}(T_n)\left(S_{n,m}(T_n) - k\right),$$

where $S_{n,m}(t)$ is the value of the fixed-for-floating swap with payments at times $T_{n+1}, \ldots, T_m$, see for instance (Andersen and Piterbarg, 2010, Chapter 19). The value of the forward swap rate $S(t)$ and swap annuity $A(t)$ at time t are given by.

$$S(t) := S_{n,m}(t) = \frac{P(t, T_n) - P(t, T_m)}{A_{n,m}(t)}, \quad A(t) := A_{n,m}(t) = \sum_{i=n}^{m-1} P(t, T_{i+1})\tau_i. \quad (3)$$

The present value $V(T_0)$ of a Bermudan swaption at time T_0 is the supremum taken over all discrete stopping times of all conditional expected dis-

counted payoffs, that is

$$V(T_0) \quad = \quad B(T_0) \sup_{\tau \in \mathcal{T}} \mathbb{E}_0 \left[\frac{U(\tau)}{B(\tau)} \right] \tag{4}$$

$$= \quad B(T_0) \mathbb{E}_0 \left[\frac{U(\tau^*)}{B(\tau^*)} \right], \tag{5}$$

where $\tau* \in \mathcal{T}$ is the optimal stopping time taking values in the finite set of allowed discrete exercise dates $\mathcal{T}$. For the American swaption, the holder is allowed to exercise on any date within $[T_1, T_{m-1}]$. And for European swaption case, we have only one exercise date, i.e., $T_1 = T_{m-1}$.

3   Monte Carlo Simulation of Bermudan Swaptions

In this section, we define the Bermudan option pricing problem, summarize SGBM, present a bundling algorithm suitable for the pricing of Bermudan swaptions in the LMM and discuss our implementation of the LSM and the upper and lower bounds via SGBM.

The present value $V(0)$ of a Bermudan swaption in (4) is usually solved via backward induction starting from the last exercise date T_{m-1}. The holder of the option receives $U(T_n)$ if the contract is exercised at time T_n. The option value at $V(T_n)$ at time T_n is the maximum of the intrinsic value $U(T_n)$ and the conditional continuation value $H(T_n)$, that is

$$V(T_n) = \max(U(T_n), H(T_n)), \tag{6}$$

where $H(T_{m-1}) = 0$. The conditional continuation value $H(T_n)$ is the conditional expected time T_{n+1} option value given by,

$$H(T_n) = B(T_n) \mathbb{E}_{T_n} \left[\frac{V(T_{n+1})}{B(T_{n+1})} \right]. \tag{7}$$

The problem is solved by recursively repeating Equations (6) and (7) for each T_n until we reach time T_0, where we find the value $V(T_0)$ of the contract.

As mentioned in the introduction, lower bound American Monte Carlo methods as LSM and SGBM are divided into two phases, a first and a second pass. In the first pass the conditional discounted option values are projected onto basis functions of the state variables. The projected value is then used as the approximate continuation value, which is compared with the intrinsic value for determining the optimal exercise strategy. This is followed by a second pass where the low-biased option values are obtained by simulating a new set of simulation paths, and exercising according to the sub-optimal exercise strategy

obtained in the first pass.

3.1   The Least Squares Method (LSM)

In the LSM the problem is solved by recursive value iteration, by the dynamic programming approach, starting from the last exercise date and working backwards as given by Equations (6) and (7). As pointed out by Clement, Lamberton, and Protter (2002), the main problem with dynamic programming is the evaluation of the conditional expectation. The LSM method is based on approximation of the conditional expectation of $H(T_n)$ at time T_n by an ordinary least squares estimate,

$$H(T_n) = \sum_{i=0}^{q} \beta_{i,n} \zeta_i(T_n), \tag{8}$$

for a set of q basis-functions $\zeta_i : \mathbb{R}^d \to \mathbb{R}$, $i = 1, 2, \ldots, q$, e.g., function of the underlying swap rates, and where $\beta_{i,n}$ are constants. The regression is usually performed using the simulated in-the-money paths and the basis functions are usually polynomials of the state variables. The optimal stopping time derived using this approximation, denoted by τ, can be written as

$$\tau_n = t_n \mathbf{1}\{H(T_n) \le U(T_n)\} + t_{n+1} \mathbf{1}\{H(T_n) > U(T_n)\}, \quad n < m-1, \tag{9}$$

having $\tau_{m-1} = T_{m-1}$. The option price is then computed using Equation (5).

A rigorous mathematical justification and proof of the almost sure convergence of the method can be found in Clement, Lamberton, and Protter (2002).

3.2   The Stochastic Grid Bundling Method (SGBM)

SGBM is a simulation-based dynamic programming method, which first generates Monte Carlo paths, forward in time, followed by finding the optimal early-exercise policy moving backwards in time. The main difference between LSM and SGBM is that in SGBM one projects the option values onto a set of basis functions of the state variables where the distribution is analytically (or approximately) known, whereas in LSM this is not taken into account.

The discounted continuation value, $H(T_n)$ in Equation (7), is computed using the law of iterated expectations, i.e.,

$$\mathbb{E}[X|\mathcal{H}] = \mathbb{E}[\mathbb{E}[X|\mathcal{G}]|\mathcal{H}], \tag{10}$$

where $\mathcal{H}$ is a sub-σ algebra of $\mathcal{G}$. Using Equation (10), the continuation value

at time T_n can be written as

$$
\begin{aligned}
H(T_n) &= B(T_n)\mathbb{E}\left[\left.\frac{V(T_{n+1})}{B(T_{n+1})}\right| S(T_n)\right] \\
&= B(T_n)\mathbb{E}\left[\left.\mathbb{E}\left[\left.\frac{V(T_{n+1})}{B(T_{n+1})}\right| \zeta(T_{n+1}), S(T_n)\right]\right| S(T_n)\right],
\end{aligned}
\tag{11}
$$

where $\zeta(T_n) = \left(\zeta_1(T_{n+1}),\ldots,\zeta_q(T_{n+1})\right)^\top$ is a q-dimensional vector of regression variables, for example the q first monomials

$$
\zeta_i(T_{n+1}) = S(T_{n+1})^i, \quad i = 1,\ldots,q,
\tag{12}
$$

and where S is the swap rates defined in (3).

Writing the continuation value as in Equation (11) decomposes the problem into two steps. The first step involves computing the inner conditional expectation,

$$
Z(T_n) = \mathbb{E}\left[\left.\frac{V(T_{n+1})}{B(T_{n+1})}\right| \zeta(T_{n+1}), S(T_n)\right].
\tag{13}
$$

It is followed by the computation of the outer expectation,

$$
H(T_n) = B(T_n)\mathbb{E}[Z(T_n)| S(T_n)].
\tag{14}
$$

With a smart choice of basis functions ζ and simulation measure, Equation (14) can generally be computed in "closed-form". However, numerical approximations are involved in the computation of $Z(T_n)$ in Equation (13).

Consider the conditional expectation without the extra conditioning on $S(T_n)$ as in Equation (13),

$$
\mathbb{E}\left[\left.\frac{V(T_{n+1})}{B(T_{n+1})}\right| \zeta(T_{n+1})\right].
\tag{15}
$$

Equation (15) can be approximated by regressing $V(T_{n+1})/B(T_{n+1})$ onto the first $q < \infty$ basis functions, $\zeta_1,\ldots,\zeta_q$. For example, by using the polynomials of the conditioning function as the basis, e.g., polynomials up to order 2-4 constructed by the monomials of the explanatory variable.

But in order to compute $Z(T_n)$ in Equation (13), we also need to condition $V(T_{n+1})$ on $S(T_n)$, which can be done in two ways. In the first approach, with nested Monte Carlo simulation, the paths are simulated until the next time T_{n+1} with $S(T_n)$ as the source, the option values for these paths are used to approximate Equation (13). The fitted value of this regression will converge in mean square and probability, when the number of paths in this sub-simulation goes to infinity. However, this approach will be computationally intractable as the number of paths grows exponentially with each time step. The second approach, is

to condition $V(T_{n+1})$ on $S(T_n)$ and then use bundling.

Bundling as introduced by Tilley (1993) is a method to partition the state space into non-overlapping regions, so that each point in the space can be identified to lie in exactly one of the bundled regions. The idea behind bundling is that for a large set of paths, the neighbouring paths will have similar continuation values and one can therefore perform local-averaging. The key step is to construct bundles, by first generating K paths, $\omega_1, \ldots, \omega_K$, of the underlying asset, $S(T_n, \omega_k)$, and bundle them at each time, T_n, into $a_n(K)$ non-overlapping sets, $\mathscr{B}^s(T_n) = \left(\mathscr{B}^1(T_n), \ldots, \mathscr{B}^a(T_n)\right)$. This is done by defining at each time, T_n, representative states μ_n^s for $s = 1, \ldots, a_n(K)$. The s-th-bundle at time T_n is thus defined as

$$\mathscr{B}^s(T_n) = \left\{ S(T_n, \omega_k) : \left\| S(T_n, \omega_k) - \mu_n^s \right\|_2 \le \left\| S(T_n, \omega_k) - \mu_n^\ell \right\|_2, \ \forall \ 1 \le \ell \le a_n(K) \right\}, \tag{16}$$

for $k = 1, \ldots, K$ and where μ_n^s is the mean of the points in $\mathscr{B}^s(T_n)$.

The continuation value (7) for a general path ω_k at time T_n is then approximated by,

$$\widehat{H}(T_n, \omega_k) = B(T_n, \omega_k) \mathbb{E}\left[\left. \frac{\widehat{V}(T_{n+1}, \omega_k)}{B(T_{n+1}, \omega_k)} \right| \mathscr{B}(T_n, \omega_k) \right], \tag{17}$$

where bundle $\mathscr{B}(T_n, \omega_k)$ is the set of path-indices of paths that lie in the bundle containing $S(T_n, \omega_k)$.

SGBM employs a *recursive bifurcation algorithm* to bundle the grid points at each time step, the number of partitions, or bundles, after p iterations, equals 2^p. The algorithm is explained in detail in Appendix A and Figure 1 illustrates the idea behind the bundling from simulated swap rates and continuation values, using 2 respectively 4 bundles.

As explained, SGBM computes the continuation value in two steps. First, we compute the expected option value, conditioned on a finer information set, given by Equation (13), which is followed by the computation of the outer expectation, given by Equation (14). Let $\mathscr{B}(T_n, \omega_k)$ denote the set of path-indices of paths that share the bundle containing the k-th grid point $S(T_n, \omega_k)$ at time T_n. Second, we approximate Z in Equation (13) by regressing the option values at T_n for those paths that originate from the bundle containing $S(T_n, \omega_k)$, that is

$$\widehat{Z}(T_n, \omega_k) = \sum_{i=1}^{q} \beta_{i,n} \zeta_i(T_n, \omega_k), \tag{18}$$

where $k \in \mathscr{B}(T_n, \omega_k)$ so that the following residual is minimized

$$\min_{\beta} \sum_{l \in \mathscr{B}(T_n, \omega_k)} \left(\widehat{Z}(T_n, \omega_l) - V(T_n, \omega_l) \right)^2.$$
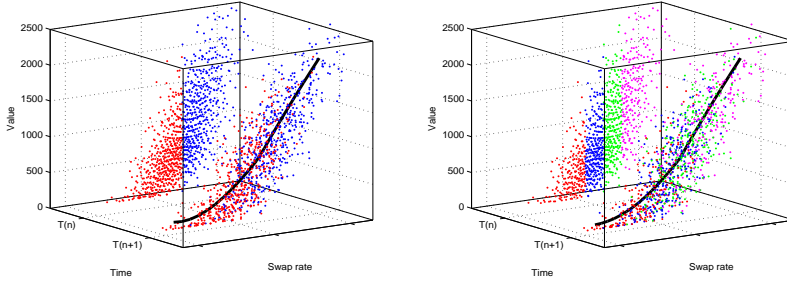
Figure 1: Simulated continuation values as a function of the swap value for a Bermudan swaption at one of the exercise dates. Continuation values approximated by a second order polynomial, with the swaps as basis functions. Left: Regression with 2 bundles. Right: Regression with 4 bundles.

The continuation value for grid point $S(T_n, \omega_k)$ in bundle $\mathscr{B}(T_n, \omega_k)$ is then given by,

$$\widehat{H}(T_n, \omega_k) = B(T_n, \omega_k) \sum_{i=0}^{q} \beta_{i,n} \mathbb{E}_{T_n} \left[\frac{\zeta_i(T_{n+1}, \omega_k)}{B(T_{n+1}, \omega_k)} \middle| \mathscr{B}(T_n, \omega_k) \right]. \quad (19)$$

Remark 5 *SGBM requires significantly fewer paths and basis functions than LSM. The reason for this is that LSM uses the regressed continuation values to make early exercise decision directly. The quality of the early exercise policy is inaccurate when a small number of paths and basis functions are used, one therefore need a large number of paths and basis functions to reduce the regression noise. In SGBM, the regressed function is just an inner expectation. The outer expectation, which can be calculated analytically, gives the continuation value and is used for decision making. Since the regression error is normally distributed with a zero mean, the noise of outer expectation of is zero. Therefore, the continuation value surface generated by SGBM is much smoother, compared to the one generated by LSM*

Further details of SGBM, such as convergence and different bundling algorithms can be found in Jain and Oosterlee (2015).

Some of the difficulties in the pricing of the Bermudan swaptions lie in the choice of regression variables. Choosing a suitable set of explanatory variables and parametric functions is crucial. In our framework this can be considered as a combination of both art and science. An estimation of the exercise boundary

close to the true boundary gives an estimated price closer to the true value. One significant problem with regression is ease of overfitting. One should not therefore use too many regression variables and high-order polynomials since they are easily affected by outliers in the simulation. One needs to focus on finding significant explanatory variables. Glasserman and Yu (2004) showed that for the simplest case of Bermudan swaptions a second-order polynomial with the underlying swap values as basis appears sufficient to obtain accurate Bermudan swaption values. More generally, the choice of basis functions is usually product dependent and needs to be carefully investigated for complicated Bermudan swaptions, for example for products with exotic coupons.

We conclude this section by emphasizing the choices of measures used in order to allow for efficient simulation. The T-forward measure, with corresponding expectation $\mathbb{E}^T$ and the T-maturity zero coupon bond $P(t, T)$ as the numeraire have the advantage that it allows for decoupling the payoff $V(T)$ from the numeraire and to take out the discount factor from the expectation, i.e,

$$
\begin{aligned}
V(t) &= B(t)\mathbb{E}_t\left[\frac{V(T)}{B(T)}\right] \\
&= P(t, T)\mathbb{E}_t^T[V(T)].
\end{aligned}
$$

One benefit however of the spot measure compared to the T-forward measure is that the numeraire asset $B(t)$ is alive throughout the tenor and therefore allows for simulating paths irrespective of tenor. We employ *hybrid measures* to obtain efficient Monte Carlo simulation. The inner expectation is simply approximated by regression calculated in the spot measure. To express the outer expectation in closed form, we compute the expectation under the T-forward measure. Since the spot measure $\mathbb{Q}^B$ coincides with the T_{n+1}-forward measure $\mathbb{Q}^{T_{n+1}}$ over the interval $[t_n, t_{n+1}]$ this allows us to write the continuation value for grid point $S(T_n, \omega_k)$ in bundle $\mathscr{B}(T_n, \omega_k)$ as,

$$
\begin{aligned}
\widehat{\mathrm{H}}(T_n, \omega_k) &= B(T_n, \omega_k)\mathbb{E}\left[\frac{\sum_{i=0}^q \beta_{i,n}\zeta_i(T_{n+1}, \omega_k)}{B(T_{n+1}, \omega_k)}\middle| \mathscr{B}(T_n, \omega_k)\right] \\
&= P(T_n, T_{n+1}, \omega_k)\mathbb{E}^{T_{n+1}}\left[\sum_{i=0}^q \beta_{i,n}\zeta_i(T_{n+1}, \omega_k)\middle| \mathscr{B}(T_n, \omega_k)\right] \\
&= P(T_n, T_{n+1}, \omega_k)\sum_{i=0}^q \beta_{i,n}\mathbb{E}^{T_{n+1}}[\zeta_i(T_{n+1}, \omega_k)|\mathscr{B}(T_n, \omega_k)].
\end{aligned}
$$

By this we can efficiently simulate the exposure. We refer to (Andersen and Piterbarg, 2010, Chapter 4) for additional information on available fixed income probability measures.

Remark 6 *Valuation of Bermudan swaptions with American Monte Carlo techniques such as LSM and SGBM requires an estimation of the exercise boundary. The option can then be seen as a barrier option (knock-in) with the estimated exercise boundary as the barrier.*

3.2.1 Algorithm for Lower Bound

For clarity we summarise the steps of the complete SGBM pricing process for the Bermudan swaptions. We first simulate a first pass with K_1 paths and then estimate an exercise policy. Second, we simulate a second pass with K_2 paths using the exercise policy estimated in the first pass. Usually $K_2 \approx 10,000$ to $100,000$ and $K_1 \approx K_2/4$.

I. FIRST PASS: Exercise Policy.

---

1. Generate K_1 paths $\omega_1,\ldots,\omega_{K_1}$, using (1). Each ω_k represents one simulated path of all core LIBOR rates.

2. For each path ω_k and time T_n, for $k = 1,\ldots,K_1$ and $n = 1,\ldots,N-1$, calculate the numeraire $B(T_n,\omega_k)$, swap rates $S(T_n,\omega_k)$ and exercise values $U(T_n,\omega_k)$.

3. Compute the option value for the grid points at the terminal time T_{m-1},
$$V(T_{m-1}) = \max(U(T_{m-1}),0). \qquad (20)$$

4. For each $n = m-2\ldots,1$

   (a) Bundle the grid points at T_{n-1}, into a distinct bundles (except at T_0, where there is only one point and hence only one bundle corresponding to $S(T_0)$) using the bundling algorithm in Appendix A.

   (b) Compute the regression functions, Z_n^s, $s = 1,\ldots,a$, as given by Equation (18), using the option values at T_n for the paths originating from the s-th bundle, $\mathscr{B}^s(T_{n-1})$, at T_{n-1}.

   (c) Compute the continuation value for the grid points in the s-th bundle at T_{n-1}, using Equation (19), for those paths for which $S(T_n,\omega_k)$ belongs to the bundle $\mathscr{B}^s(T_{n-1})$, for $s = 1,\ldots,a$.

   (d) Compute the option values at T_{n-1}, as

$$\widehat{V}(T_n) = \max(U(T_n),H(T_n)). \qquad (21)$$

5. The option value, $\widehat{V}(T_0)$, at T_0, is defined to be the *direct estimator* value.

II. SECOND PASS: Lower Bound.

---

1. In order to compute lower bounds and an unbiased price, generate a fresh set of K_2 paths, $\omega'_1, \ldots, \omega'_{K_2}$, and bundle at each time step, using the same algorithm as in the first pass.

2. The continuation values for the grid points in bundle s, at time step T_{n-1}, are computed using the Z_n^s function, obtained for the direct estimator. The option is exercised when the continuation value is less than the immediate payoff. The lower bound can then be computed by determining the earliest time to exercise at each path, $\widehat{\tau}\left(\omega'_k\right) = \min\left\{T_n : H\left(T_n, \omega'_k\right) < U\left(T_n, \omega'_k\right)\right\}$. The lower bound of the option value is given by,

$$\underline{V}(T_0) = \frac{1}{K_2} \sum_{k=1}^{K_2} \frac{B\left(T_0, \omega'_k\right)}{B\left(\widehat{\tau}, \omega'_k\right)} U\left(\widehat{\tau}, \omega'_k\right).$$

Remark 7 *In the case of performance calculation issues, the direct estimate generated from the first pass will often be close to the lower bound values generated from the second pass. In this case one can neglect the second pass but should keep in mind that the estimated values are biased in an unknown direction.*

Remark 8 *One should also keep in mind that both LSM and SGBM are lower bound methods, basically because the conditional expectation is approximated by a regression technique that projects the high-dimensional continuation value onto a limited set of regression variables. The approximation can often be improved, for example, by having a richer and better set of regression variables, but with the risk of overfitting.*

3.3 Upper Bound Using Dual Formulation

One problem with the lower bound algorithm presented in the previous section is to determine how close the generated option prices are to the true value. One way to determine its goodness it to simulate both lower and upper bounds of the option values, the closer they are to each other the better. Haugh and Kogan (2004) and Rogers (2002) independently proposed the dual formulation for

Bermudan options, later extended to the primal-dual simulation algorithm in Andersen and Broadie (2004). The primal problem is given by Equation (4), for an arbitrary adapted super-martingale process $M(t)$ we have that,

$$
\begin{aligned}
V(T_0) &= \sup_{\tau \in \mathcal{T}} \mathbb{E}_{T_0} \left[\frac{U(\tau)}{B(\tau)} \right] \\
&= \sup_{\tau \in \mathcal{T}} \mathbb{E}_{T_0} \left[\frac{U(\tau)}{B(\tau)} + M(\tau) - M(\tau) \right] \\
&= M(0) + \sup_{\tau \in \mathcal{T}} \mathbb{E}_{T_0} \left[\frac{U(\tau)}{B(\tau)} - M(\tau) \right] \\
&\leq M(0) + \mathbb{E}_{T_0} \left[\sup_{\tau \in \mathcal{T}} \left(\frac{U(\tau)}{B(\tau)} - M(\tau) \right) \right].
\end{aligned}
\tag{22}
$$

The inequality follows from the fact $M(t)$ is a super-martingale. The dual formulation of the option pricing problem is then to minimize the upper bound with respect to all adapted super-martingale processes $\mathcal{K}$, that is,

$$
\overline{V}_0 = \inf_{M \in \mathcal{K}} \left\{ M(0) + \mathbb{E}_{T_0} \left[\sup_{\tau \in \mathcal{T}} \left(\frac{U(\tau)}{B(\tau)} - M(\tau) \right) \right] \right\}.
\tag{23}
$$

Haugh and Kogan (2004) showed that when the super-martingale process $M(t)$ in Equation (22) coincides with the discounted option value process $V(t)/B(t)$, the upper bound $\overline{V}_0$ equals the true value. This suggests that a tight upper bound can be obtained by approximation $\widehat{V}(t)$, when defining $M(t)$ such that when the approximate option price $\widehat{V}(t)$ coincides with the exact price $V(t)$, $M(t)$ equals the discounted process $V(t)/B(t)$. An obvious choice for $M(t)$ is then given by

$$
M(T_{n+1}) - M(T_n) = \frac{\widehat{V}(T_{n+1})}{B(T_{n+1})} - \frac{\widehat{V}(T_n)}{B(T_n)} - \mathbb{E}_{T_n} \left[\frac{\widehat{V}(T_{n+1})}{B(T_{n+1})} - \frac{\widehat{V}(T_n)}{B(T_n)} \right],
\tag{24}
$$

for $M(T_0) = \widehat{V}(T_0)$. Equation (24) can also be written as

$$
M(T_{n+1}) - M(T_n) = \frac{\widehat{V}(T_{n+1})}{B(T_{n+1})} - \mathbb{E}_{T_n} \left[\frac{\widehat{V}(T_{n+1})}{B(T_{n+1})} \right].
\tag{25}
$$

Then the upper bound, $\overline{V}_0$, corresponding to Equation (24) is given by

$$
\overline{V}(0) = \widehat{V}(0) + \Delta \geq V(0),
\tag{26}
$$

where the duality gap Δ is defined as

$$\Delta = \max_n \left(\frac{U(T_n)}{B(T_n)} - M(T_n) \right), \quad n = 1, \ldots, m-1. \tag{27}$$

Generally, upper-bounds algorithm as e.g., in Andersen and Broadie (2004) require nested Monte Carlo simulation and the quality of the upper bound produced by the algorithm depends on the quality of the estimated exercise policy in the first pass, better policy gives tighter upper-lower bounds. This makes it computational expensive and requires in worst cases a workload of $K \times K_{\text{nest}} \times m^2$ operations, where K is the number of outer simulations, K_{nest} the number of nested simulations and m the number of exercise dates. The workload is often less than this because the nested simulation can be stopped whenever the contract is exercised. This in comparison with the workload in the second pass where an exercise policy already is given and where the lower bound simulation has a workload of $K \times m$. This has further been improved by Broadie and Cao (2008) who showed that nested simulations are not needed on dates where it is suboptimal to exercise the option, which can lead to reduced workload, especially for out-of-the-money options. When the policy obtained from LSM is used, a sub-simulation with K_{nest} sub-paths is required. $K_{\text{nest}} \approx 100$ is often sufficient to find upper bounds with sufficient quality. Moreover, the upper bound bias introduced from the Monte Carlo simulation is positive and a decreasing function in the number of nested simulations.

We conclude this section by emphasising one important remark allowing one to avoid nested Monte Carlo simulations for upper bound values when estimating the exercise policy using SGBM.

Remark 9 *For pure regression based algorithms like LSM, Equation* (25) *cannot be estimated directly by regression since it will introduce an unknown bias and therefore destroys the martingale property of M and the inequality in Equation* (22). *Therefore one has to rely on nested Monte Carlo simulations to obtain an upper bound when LSM is used. But, as mentioned in Remark 5, the regressed function in SGBM is just the inner expectation, and it is not used for decision making. The outer expectation can be computed in closed form and we can therefore calculate the upper bounds without nested simulations. This reduces the workload of the upper bounds significantly, to the workload of the second pass and therefore we can obtain a speed-up factor of K_{nest}. The computational time for SGBM is comparable to Longstaff and Schwartz (2001)*

3.3.1   Algorithm for Upper Bound

We summarise the simulation procedure for obtaining duality based upper bounds via SGBM below (once the optimal exercise policy has been obtained). Let $\widehat{H}(T_n)$ be the holding value estimated from the exercise strategy η, given by the simulation in the first pass. The upper bound can then be obtained by the following algorithm,

---

III. THIRD PASS: Upper Bound.

---

1. Simulate K_u paths $\omega_1, \dots, \omega_{K_U}$

2. For each exercise time T_n and each path ω_k, compute $\widehat{H}(T_n, \omega_k)$ and $B(T_n, \omega_k)$, and update $M(T_n, \omega_k)$ in Equation (27).

    (a) Approximate $\widehat{H}(T_n, \omega_k) / B(T_n, \omega_k)$ using Equation (17). [2]

3. For each path ω_K, compute the pathwise duality gaps, as follows,

$$\widehat{D}(\omega_k) = \max_n \left(\frac{U(T_n, \omega_k)}{B(T_n, \omega_k)} - M(T_n, \omega_k) \right), \quad n = 1, \dots, m.$$

4. Estimate the upper bound given by Equation (27) as

$$\widehat{\Delta} = \frac{1}{K_U} \sum_{k=1}^{K_U} \widehat{D}(\omega_k). \tag{28}$$

4   Numerical Results

In this section we study the performance of SGBM for lower and upper bound values by means of numerical experiments. For a consistency check we use the same setup and reproduce the results in Andersen (2000) and Andersen and Piterbarg (2010).

---

[2] In order to have an accurate approximation of $H(T_n, \omega_k)$, when the policy obtained from LSM is used, a sub-simulation with K_{nest} sub-paths is required. $\widehat{H}(T_n, \omega_k)$ represents the discounted average cashflows from these paths when they are exercised following the policy obtained in the first pass.

4.1   Bermudan Swaption Prices

For the continuation value in LSM and SGBM we use a second-order polynomial with the swap rate as the basis. The swap rate moments in Equation (19) can, for example, with high accuracy be calculated by the convexity adjustment approach in Belomestny, Kolodko and Schoenmakers (2009).

We use the bundling scheme described in Section A, with 8 bundles and the same number of bundles at each time step, except at time T_0, where there is only one point, $S(T_0)$.

We consider Bermudan swaptions on 3 months LIBORS ($\tau = 0.25$) with 10% spot rate level and with two different volatility settings. First, a one-factor LMM with fixed volatility, $\lambda_n(t) = 0.2$ for all n and t. Second, more realistically, a two-factor LMM, with a time-to-maturity dependent volatility of the form[3]

$$\lambda_n(t) \quad = \quad \left[0.15, 0.15 - \sqrt{0.009\,(T_n - t)} \right]^\top.$$

We report values obtained from the second pass. First, we simulate a first pass with 10,000 seeds using an antithetic Monte Carlo random number generator and then estimate the exercise policy for both LSM and SGBM. Subsequently, we simulate 20,000 second pass paths with a quasi Monte Carlo random number generator (e.g., Sobol sequence) with the previously obtained exercise policy to estimate the unbiased Bermudan swaption value. These two steps are repeated iteratively $K' = 100$ times with different seeds in the first simulation, to remove the overall influence of the first simulation. The prices are reported in basis points, with the notional $\mathcal{N} = 10,000$ and the numbers in parentheses are sample standard deviations.

Duality-based upper bounds, together with the lower bound computed using the path estimator give valid confidence intervals within which the true option price lies. The $100\,(1 - \gamma)\,\%$ confidence interval is constructed as

$$\left[\underline{V}_0(T_0) - q_{\gamma/2} \frac{\hat{s}_L}{\sqrt{K'}},\ \overline{V}_0(T_0) + q_{\gamma/2} \frac{\hat{s}_H}{\sqrt{K'}} \right],$$

where $\hat{s}_L$, is the sample standard deviation for the path estimator and $\hat{s}_H$, is the sample standard deviation for the duality-based upper bound estimator and $q_{\gamma/2}$ the normal distributed quantile function.

Tables 1 and 2 report the lower bound value estimates for the Bermudan swaption via LSM and SGBM, the duality gap and 95%-confidence interval with one-factor, respectively two-factor LMM. Our reported values for SGBM differ at most 3 bps compared to the reported values in Andersen and Piterbarg (2010).

---

[3]Usually, a one-factor LMM already accounts for more than 98 percent and the two factor for more than 99.5 percent of the overall variance.

Type	Strike	LSM Lower		SGBM Lower		$\hat{\Delta}_{SGBM}$	$\hat{\Delta}_{AP}$	95% CI
15M/3M	8%	184.61	(0.01)	184.62	(0.00)	0.0022	0.02	184.62 - 184.63
15M/3M	10%	49.11	(0.01)	49.11	(0.00)	0.0008	0.02	49.111 - 49.114
15M/3M	12%	8.73	(0.02)	8.73	(0.00)	0.0001	0.004	8.7322 - 8.7346
3Y/1Y	8%	355.08	(0.08)	355.06	(0.02)	0.0133	0.07	355.05 - 355.07
3Y/1Y	10%	157.13	(0.11)	157.45	(0.03)	0.0030	0.2	157.45 - 157.46
3Y/1Y	12%	60.96	(0.07)	60.97	(0.02)	0.0011	0.04	60.97 - 60.98
6Y/1Y	8%	806.61	(0.41)	808.11	(0.08)	0.0186	0.23	808.09 - 808.14
6Y/1Y	10%	415.35	(0.82)	418.58	(0.13)	0.0088	0.63	418.55 - 418.61
6Y/1Y	12%	212.13	(0.48)	214.16	(0.12)	0.0041	0.33	214.13 - 214.19
11Y/1Y	8%	1377.00	(1.07)	1383.10	(0.26)	0.0307	1.3	1383.00 - 1383.10
11Y/1Y	10%	805.93	(1.00)	811.13	(0.23)	0.0188	1.3	811.08 - 811.20
11Y/1Y	12%	495.16	(0.69)	499.20	(0.27)	0.0120	0.7	499.15 - 499.27
6Y/3Y	8%	493.91	(0.15)	494.12	(0.04)	0.0235	0.08	494.11 - 494.15
6Y/3Y	10%	291.84	(0.22)	293.03	(0.05)	0.0092	0.65	293.02 - 293.05
6Y/3Y	12%	169.22	(0.19)	169.79	(0.04)	0.0040	0.53	169.79 - 169.80

Table 1: Lower bound estimate of Bermudan payer swaptions in a one-factor LMM. Prices are in basis points and standard deviations within parentheses.

The computational time for SGBM is roughly the same as for the LSM. The first conclusion, the standard deviation for SGBM lower bounds is much smaller than the ones obtained from LSM. On average, the ratio of variance of LSM and SGBM is around 16, meaning that on average, one would need 16 times fewer Monte Carlo seeds in order to obtain the same pricing accuracy. The second conclusion, the duality gap $\hat{\Delta}_{SGBM}$ obtained by SGBM is significantly smaller than the duality gap $\hat{\Delta}_{AP}$ as reported in Andersen and Piterbarg (2010). As one can observe we obtain significantly smaller duality gaps, the largest duality gap for the one-factor LMM is 0.0307 basis points, compared to 1.3 basis points in Andersen and Broadie (2004). The conclusion here is that the duality gap obtained by SGBM gives rise to really tight lower-upper bounds.

Figure 2 and 3 illustrate the exercise boundary and the exercise frequency for the one- and two-factor LMM, respectively. The pictures demonstrate the sensitivity of the product regarding the exercise policy. We can clearly see that the Bermudan swaption is mostly canceled directly at the first or last exercise date. The figures also demonstrate that the stopping times are not very differ-

Type	Strike	LSM Lower		SGBM Lower		$\hat{\Delta}_{SGBM}$	$\hat{\Delta}_{AP}$	95% CI
15M/3M	8%	183.83	(0.01)	183.83	(0.00)	0.0003	0.05	183.83 - 183.83
15M/3M	10%	42.17	(0.02)	42.24	(0.02)	0.0009	0.06	42.238 - 42.247
15M/3M	12%	5.21	(0.01)	5.22	(0.01)	0.0001	0.01	5.2183 - 5.2204
3Y/1Y	8%	339.15	(0.05)	339.35	(0.02)	0.0102	0.4	339.34 - 339.36
3Y/1Y	10%	125.12	(0.06)	125.58	(0.02)	0.0024	0.7	125.57 - 125.58
3Y/1Y	12%	35.76	(0.05)	35.87	(0.02)	0.0004	0.2	35.866 - 35.875
6Y/1Y	8%	747.23	(0.19)	751.88	(0.06)	0.0128	3.7	751.86 - 751.9
6Y/1Y	10%	315.73	(0.40)	319.18	(0.10)	0.0054	5.0	319.16 - 319.21
6Y/1Y	12%	126.41	(0.31)	129.14	(0.08)	0.0020	2.6	129.12 - 129.16
11Y/1Y	8%	1237.80	(0.63)	1253.40	(0.20)	0.0191	18.1	1253.4 - 1253.5
11Y/1Y	10%	610.34	(0.65)	628.93	(0.26)	0.0142	20.8	628.88 - 628.99
11Y/1Y	12%	322.55	(0.67)	335.18	(0.17)	0.0071	14.8	335.15 - 335.22
6Y/3Y	8%	444.83	(0.16)	446.15	(0.03)	0.0194	0.8	446.14 - 446.17
6Y/3Y	10%	225.67	(0.14)	227.24	(0.04)	0.0054	1.2	227.23 - 227.25
6Y/3Y	12%	106.16	(0.11)	107.27	(0.03)	0.0019	0.8	107.26 - 107.27

Table 2: Lower bound estimate of Bermudan payer swaptions in a two-factor LMM. Prices are in basis points and standard deviations within parentheses.

ent from each other, although the few scenarios in which SGBM in comparison with LSM exercises earlier have a significant effect on the price. This is the reason why the SGBM prices are superior to the LSM prices in these tests and why SGBM demonstrates tight lower-upper bounds.

4.2 The Effect of Number of Bundles

In this section we study how the number of bundles affect the lower, upper bound values and the duality gap. In particular, we study Bermudan swaptions with the same setup as in the previous section for a 10% coupon with 2^p bundles for $p = 1, 2, 3, 4$. The duality gap is illustrated in Figure 4. We observe an almost log-linear relationship between the duality gap and the number of bundles. Increasing the number of bundles will make the duality gap much smaller.
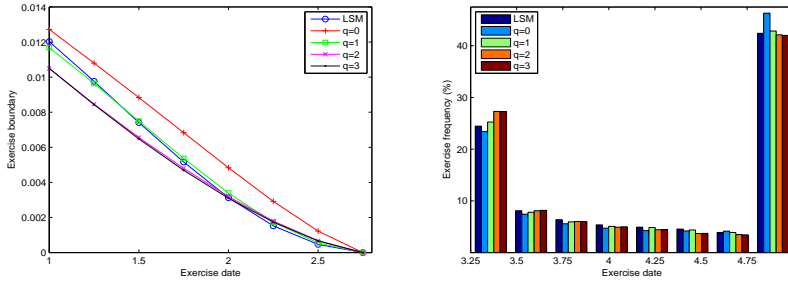
Figure 2: Exercise boundary for LSM and SGBM with bundles, 2^q, for $q = 0, 1, 2, 3$ for a 3Y/1Y Bermudan Payer Swaptions in a one-factor LIBOR market model.
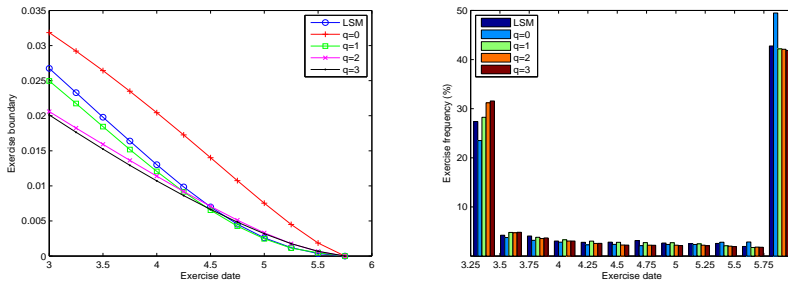
Figure 3: Exercise boundary for LSM and SGBM with bundles, 2^q, for $q = 0, 1, 2, 3$ for a 6Y/3Y Bermudan Payer Swaptions in a two-factor LIBOR market model.

5   Conclusion

This paper presented the application of the Stochastic Grid Bundling Method (SGBM) for approximating the values of Bermudan style options on the LMM by simulation. SGBM is a regression-based Monte Carlo method where the continuation value is projected onto a space where the distribution is known. In the method, a practical bundling algorithm is employed which completes the algorithm and performs very well for the test cases considered. We also demonstrate how to obtain upper bounds without the need for nested Monte Carlo simulations as generally required for regression based methods. The upper-lower bounds obtained by SGBM are much tighter compared to the bounds obtained by traditional methods. We illustrate SGBM's performance using a number of realistic examples. The computational time for the method is comparable to
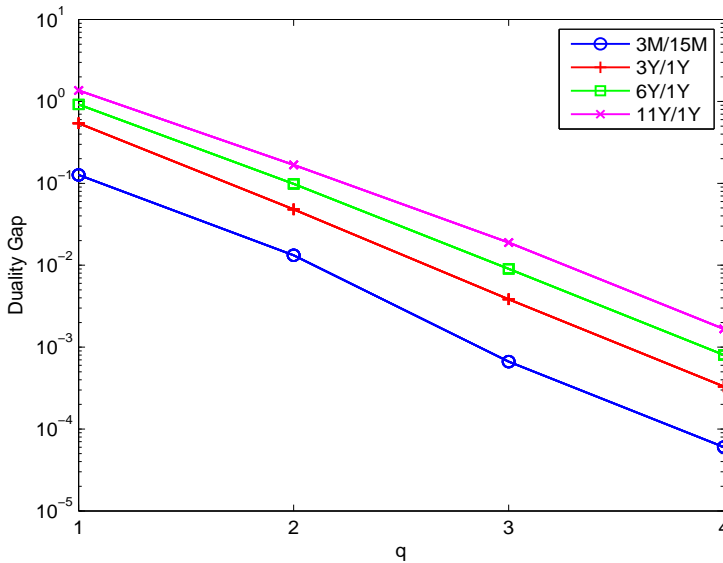
Figure 4: The duality gap from SGBM with varying numbers of bundles, 2^q, for $q = 1, 2, 3, 4$ for Bermudan payer swaptions in a one-factor LIBOR market model.

Longstaff and Schwartz (2001), but a higher accuracy is achieved as demonstrated by the reduced Monte Carlo variance. The SGBM method is easy to implement and accurate. Variance reduction, based on iterated conditioning, in combination with the bundling technique form the necessary ingredients for accurate Bermudan swaptions valuation with a relative small number of paths and basis functions. One should however keep in mind that we need to know the conditional expected value of the basis functions.

A Bundling

Suppose we need to bundle K_s grid points at epoch T_n, given by $S(T_n, \omega_k)$, where $k = 1, \ldots, K_s$. The following steps are performed recursively.

1. Compute the *mean* of the given set of grid points,

$$\mu_n^s = \frac{1}{K_s} \sum_{k=1}^{K_s} S(T_n, \omega_k).$$

2. Bundling the grid points is performed by dividing the grid points into two groups, depending on whether the asset price for the grid point is greater or less than the mean of the asset prices for the given set of grid points:

$$\begin{aligned}
\mathscr{B}^1(T_n, \omega_k) &= \mathbf{1}\left(S(T_n, \omega_k) > \mu_n^s\right), \\
\mathscr{B}^2(T_n, \omega_k) &= \mathbf{1}\left(S(T_n, \omega_k) \le \mu_n^s\right),
\end{aligned}$$

for $k = 1, \ldots, K_s$. $\mathscr{B}^1(T_n, \omega_k)$ returns 'true', when the asset price $S(T_n, \omega_k)$ is greater than the mean, μ_n^s and belongs to bundle 1. $\mathscr{B}^2(T_n, \omega_k)$ returns 'true', if it less than the mean and belongs to bundle 2. Formally, $\mathscr{B}^s(T_n, \omega_k)$ returns 'true', if the grid point $S(T_n, \omega_k)$ belongs to bundle s.

3. Bundles $\mathscr{B}^1(T_n)$ and $\mathscr{B}^2(T_n,)$ can be split again, returning to step 1.

PAPER 3

Counterparty Credit Exposures for Interest Rate Derivatives using the Stochastic Grid Bundling Method

with Shashi Jain and Cornelis W. Oosterlee

Abstract

The regulatory credit valuation adjustment (CVA) for an outstanding over-the-counter (OTC) derivative portfolio is computed using the portfolio's exposure over its lifetime. Usually, future portfolio exposure is approximated using a Monte Carlo simulation because the portfolio value can be driven by several market risk factors. For derivatives that lack an analytical approximation for their mark-to-market (MtM) value, such as Bermudan swaptions, the standard market practice is to use the regression functions from the least squares Monte Carlo method to approximate their MtM along simulated scenarios. However, such approximations have significant bias and noise, resulting in an inaccurate CVA charge. In this paper, we extend the Stochastic Grid Bundling Method (SGBM) for the one-factor Gaussian short rate model to efficiently and accurately compute expected exposure, potential future exposure, and CVA for Bermudan swaptions. A novel contribution of the paper is that it demonstrates how different measures, such as spot and terminal measures, can simultaneously be employed in the SGBM framework to significantly reduce the variance and bias of the solution.

1   Introduction

The notional of outstanding over-the-counter (OTC) derivatives has grown exponentially over the last two decades, a rapid growth mainly resulting from the increase in interest rate derivatives. Figure 1 illustrates the Bank for International Settlements' semi-annual market survey of outstanding OTC derivatives from June 1998 through December 2013. As of December 2013, the total outstanding notional value for OTC derivatives was 710.2 trillion USD, with 584.4 trillion USD in interest rate derivatives. Any trading desk that enters an OTC deal will face the risk that the counterparty at a future date may default and cannot fulfil its payment obligations. Therefore, the bank needs to estimate the total risk it faces with respect to a particular counterparty and to maintain a capital buffer, i.e. the capital requirement, to cover for potential losses attributable to a default.

Before the financial crisis of 2007, the general market view was that large companies were *"too-big-to-fail"* and, thus, an overall tendency existed to underestimate counterparty risk. *"A too-big-to-fail firm is one whose size, complexity, interconnections, and critical functions are such that, should the firm unexpectedly go into liquidation, the financial system and the economy would face severe adverse consequences"*, to quote Federal Reserve Chair Ben Bernanke in 2010. However, the bankruptcy of AIG and Lehman Brothers in 2008 demonstrated that, instead of being *"too-big-to-fail"*, they were instead *"too-big-to-be-allowed-to-fail"* (Gregory, 2010, 17). These events increased the markets' concern regarding counterparty risk and the need for better risk management when trading OTC derivatives. The Basel Committee on Banking Supervision has formulated in the Basel II and III accords regulatory standards for setting up capital requirements to cover for losses in the case of a counterparty default.

In the Basel II accord, the requirements consist of computing what is generally referred to as counterparty credit exposure, or the amount of money that can be lost if default occurs. Examples of such quantities are expected exposures (EE) and potential future exposures (PFE). In the Basel III accord, the requirements are more stringent and require the estimation of Credit Valuation Adjustment (CVA) charges.[1] CVA is an adjustment to derivatives' prices to compensate for a possible counterparty default. The value of an OTC deal that considers counterparty risk is the value without counterparty risk, the risk-free price, and a positive adjustment – the CVA charge.

Estimating CVA charges requires an underlying model and, therefore, makes it a model-dependent quantity. Products that initially were model independent, such as plain interest rate vanilla swaps, become model dependent because one needs an interest rate model to price the future portfolio exposure at simulated (also model dependent) default times of the counterparty.

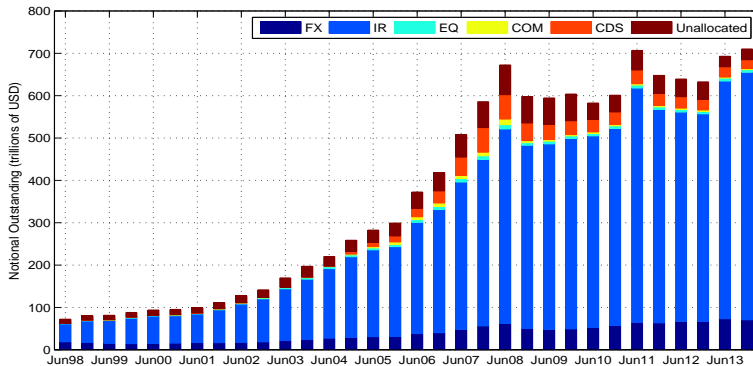[1] http://www.bis.org/publ/bcbs189_dec2010.pdf

Figure 1: Global OTC derivative markets. The notional amounts (in trillions of US dollars) outstanding of OTC derivatives by risk category from the Bank for International Settlements' semi-annual market survey, June 1998 through December 2013. For, foreign exchange (FX), interest rate (IR), equity-linked (EQ), commodity (COM) derivatives, and credit default swaps (CDS).

Moving to exotic derivatives, the situation becomes even more complex because some of them are priced using Monte Carlo simulations, and in the context of measuring counterparty risk, EE and PFE are computed using Monte Carlo simulations. Nested Monte Carlo simulations are not an option in this context for performance reasons. Rather than calculating CVA as an overnight job, a trading desk wants real-time CVA estimations for each counterparty. Additionally, to be able to hedge CVA and restructure portfolios to reduce CVA, the challenge exists of estimating risks and first-order derivatives for all input parameters.

American Monte Carlo methods, such as the well-known least squares method (LSM) as introduced by Longstaff and Schwartz (2001), for which the continuation value is approximated by a regression to determine an optimal exercise policy, are today standard among practitioners in the context of CVA for two primary reasons. First, these methods can increase computational performance by avoiding nested Monte Carlo methods, i.e. Monte Carlo simulation within a Monte Carlo simulation, by using the same set of paths for pricing and for market simulation, as in De Prisco and Rosen (2005). Second, derivatives such as American and Bermudan swaptions, i.e. products that can be exercised at various dates prior to maturity, need to be priced using Monte Carlo methods. The benefit of having an American Monte Carlo CVA calculation is that all instruments will be handled the same manner within the CVA computation,

which makes it easy to aggregate trades and include netting, collateralisation, and others.

A general problem with regression functions as they are used in the least squares Monte Carlo method is that they do not necessarily provide accurate approximations of the MtM value of the derivative over all simulated paths, and can have significant bias and noise, resulting in an inaccurate CVA charge for such products. Additionally, schemes used to improve the approximation of the MtM value of such derivatives on the valuation date, such as using only in-the-money (ITM) paths for approximations by regression functions, cannot be used for CVA purposes because they are based on exposures along all paths and scenarios.

In this paper, we extend the SGBM as introduced by Jain and Oosterlee (2015) to compute the future exposure for Bermudan swaptions, where the one-factor Gaussian short rate model is used to simulate interest rates dynamics. We show through careful numerical experiments that the EE, PFE, and CVA computed using this approach have much smaller errors and noise when compared with using the standard LSM regression-based approach. One of the novel contributions of this paper is that, under the SGBM problem formulation, in terms of an inner and outer expectation, we can benefit from the flexibility to use different pricing measures within the same computation. Specifically in the case of Bermudan swaptions, doing so allows us to avoid simulation of the numeraire process, which helps to achieve significant variance reductions relative to the LSM approach.

A comprehensive overview of CVA methodologies can be found in Canabarro and Duffie (2003), Picoult (2005), Redon (2006), Pykhtin and Zhu (2007), Pykhtin and Rosen (2010), Gregory (2010) and Brigo, Morini and Pallavicini (2013). There is extensive literature on pricing Bermudan swaptions using Monte Carlo schemes, see for instance, Andersen (2000), Bender and Schoenmakers (2006), Kolodko and Schoenmakers (2006) and Piterbarg (2004).

The paper is organized as follows. Section 2 introduces notations, the general framework and formulates the Bermudan swaption pricing. Section 3 describes the SGBM algorithm for estimating EE, PFE and CVA charges. In Section 4 we present numerical examples to illustrate the method and its efficiency compared to the traditional LSM. And we conclude in Section 5.

2 Notation and General Framework

In this section, we introduce notation, the one-factor Gaussian short rate model (GSR) and define the pricing of Bermudan swaptions. Next, we introduce the methods for estimating counterparty risk using EE, PFE and CVA.

2.1   The One-Factor Gaussian Short Rate (GSR) Model

In the general one-factor GSR model the short rate $r(t)$ follows a mean-reverting process of the form,

$$dr(t) = \kappa(t)(\theta(t) - r(t))dt + \sigma(t)dW(t), \tag{1}$$

where parameter $\kappa(t)$ is the rate of mean-reversion, $\sigma(t)$ the volatility, and $W(t)$ a standard Brownian motion. The parameters $\kappa(t)$ and $\sigma(t)$ are usually obtained by calibrating the model to plain-vanilla option prices. The deterministic drift function $\theta(t)$ can be directly calculated from the yield curve and fits the curve for

$$\theta(t) = \frac{1}{\kappa(t)}\frac{\partial f(0,t)}{\partial t} + f(0,t) + \frac{1}{\kappa(t)}\int_0^t e^{-2\int_u^t \kappa(s)ds}\sigma^2(u)\,du.$$

A non-smooth initial forward curve can affect the calculation of $\partial f(0,t)/\partial t$, but by defining a new variable $x(t) = r(t) - f(0,t)$, computations are feasible. The dynamics are given by

$$dx(t) = \big(y(t) - \kappa(t)x(t)\big)dt + \sigma(t)dW(t), \tag{2}$$

where $x(0) = 0$ and

$$y(t) = \int_0^t e^{-2\int_u^t \kappa(s)ds}\sigma^2(u)\,du.$$

A benefit with the GSR model is that the risk-neutral expectation $\mathbb{E}_t^{\mathbb{Q}}[\cdot]$ of the discounted bond price $P(t,T)$ at time t with maturity T, that is,

$$P(t,T) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r(u)du}\right],$$

is known in closed-form and given by

$$P(t,T) = \frac{P(0,T)}{P(0,t)}\exp\left(-x(t)G(t,T) - \frac{1}{2}y(t)G^2(t,T)\right),$$

$$G(t,T) = \int_0^t e^{-\int_t^u \kappa(s)ds}du.$$

We use interchangeably the following notations $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot|\mathscr{F}_t]$, where $\mathscr{F}_t$ is the filtration at time t, generated by $W(t)$.

The analytic tractability of the GSR model makes it attractive for effective numerical implementations such as for calibration procedures and Monte Carlo simulation, e.g. for pricing and CVA calculations. Criticisms are that the model allows for negative short rates and that it has very limited flexibility for mod-

elling yield curve moments, since all points on the yield curve are perfectly correlated. However, many trading desks today value Bermudan swaptions by using a GSR model, e.g. by the one-factor Hull-White model (HW1F) by Hull and White (1990) due to its simplicity and tractability.

For practical reasons to be explained, we choose to work under the spot measure $\mathbb{Q}^B$. The numeraire induced by the spot measure is the discrete version of the continuous compounded money market account with rolling certificate of deposit $B(t)$, that is

$$B(t) = P(t, T_{i+1}) \prod_{n=0}^{i} P^{-1}(T_n, T_{n+1}), \ t \in (T_i, T_{i+1}],$$

with corresponding fixed discrete tenor structure, $0 = T_0 < T_1 < \ldots < T_N$. Let $\mathbb{E}_t^B = \mathbb{E}_t$ denote the conditional expectation with respect to the measure induced by $B(t)$. One benefit with the spot measure is that the numeraire asset $B(t)$ is *"alive"* throughout the tenor and therefore, allows for simulating paths irrespective of the tenor. This is practical for e.g. Bermudan swaptions and American Monte Carlo methods, since the contract can mature randomly at any of the dates in the discrete tenor structure.

Further details on the one-factor Gaussian short rate model, such as derivations of the bond equations, connection to HJM, is out of the scope of this paper, but may be found in Brigo and Mecurio (2001).

2.2 Bermudan swaptions

A vanilla interest rate swap is a contract that allows one to change payments between two different cashflows, often a floating leg against a fixed leg. The values of the *forward swap rate* $S_{n,m}(t)$ and *swap annuity* $A_{n,m}(t)$ at time t with payments $T_{n+1}, \ldots, T_m$ are given by.

$$S_{n,m}(t) = \frac{P(t, T_n) - P(t, T_m)}{A_{n,m}(t)}, \quad A_{n,m}(t) = \sum_{i=n}^{m-1} P(t, T_{i+1}) \tau_i,$$

where $\tau_i = T_{i+1} - T_i$.

Given a lockout, i.e., a no-call (no-exercise) period up to time T_1, the Bermudan swaption gives the holder the right - but not the obligation - at a set of fixing dates T_n, for $n \in \mathscr{I} = \{1, 2, \ldots, m-1\}$, i.e., for $T_n \in \mathscr{T} = \{T_1, T_2, \ldots, T_{m-1}\}$ to enter into a fixed-for-floating swap $S_{n,m}$ with fixing date T_n and last payment date T_m. The Bermudan swaption with the fixed coupon k, exercised at time T_n corresponds to the payout given by

$$U_n = \phi \mathcal{N} A_{n,m}(T_n) \left(S_{n,m}(T_n) - k \right),$$

where $\mathcal{N}$ denotes the notional, and $\phi \in \{-1, +1\}$ the payer or receiver factor (+1 for a payer and −1 for a receiver swaption). The holder of a payer Bermudan swaption will pay the fixed swap leg and receive the floating swap leg. The present value V_0 of a Bermudan swaption is the supremum taken over all discrete stopping times of all conditional expected discounted payoffs, that is,

$$V_0 \quad = \quad B(T_0) \sup_{n^* \in \mathcal{I}} \mathbb{E}_{T_0} \left[\frac{U_{n^*}}{B(T_{n^*})} \right]. \tag{3}$$

The option value at an arbitrary time T_n is the maximum of the intrinsic value U_n and the conditional continuation value H_n, i.e.,

$$V_n = \max(U_n, H_n), \tag{4}$$

where $H_m = 0$ at maturity T_m. The continuation value H_n is the conditional expected option value at time T_{n+1} and given by,

$$H_n = B(T_n) \mathbb{E}_{T_n} \left[\frac{V_{n+1}}{B(T_{n+1})} \right]. \tag{5}$$

The problem is solved via backward induction, starting from the terminal time T_m, and solved by recursively repeating (4) and (5) until we reach time T_0, where we get the value V_0 of the Bermudan swaption contract.

2.3   Counterparty Credit Risk

The exposure $\mathrm{E}(t)$ towards a counterparty C at time t is given by the positive side of a contract (or portfolio) value $V(t)$, that is,

$$\mathrm{E}(t) = \max\{V(t), 0\}. \tag{6}$$

This can be seen as the maximum loss if the counterparty defaults at time t. Let τ_C denote the counterparty's default time, and the $\mathbb{Q}$-probability that the counterparty C defaults before time t be given by $\mathrm{PD}(t) = \mathbb{Q}(\tau_C < t)$. A commonly used default probability approximation is

$$\mathrm{PD}(t) = 1 - \exp\left(-\int_0^t \gamma(t)\, dt\right), \tag{7}$$

where the probability factor $\gamma(t)$ is called the hazard rate or the instantaneous credit spread, see Gregory (2010). The probability that the counterparty defaults in dt years given that it has not defaulted so far is $\gamma(t)\, dt$. The default probability for a given counterparty is usually bootstrapped from quoted credit default

swaps (CDS)[2].

2.3.1   Credit Value Adjustment (CVA)

CVA is the market value of counterparty credit risk, i.e., the difference between the risk-free portfolio value and the value taking into account the counterparty's default probability. The charge is computed as the integral over all points in time of the discounted expected exposure given that the counterparty defaults at that time, multiplied with the default probability and the loss given default, i.e., one minus the recovery rate R. Following Gregory (2010) the CVA on an instrument (or portfolio) with maturity T is given by

$$\text{CVA} = (1 - R) B(0) \int_0^T \mathbb{E}\left[\frac{\text{E}(t)}{B(t)} \delta(t - \tau_C) \right] dt,$$

where δ is the Dirac delta function, which is one at the counterparty C's default at time τ_C, zero otherwise, and T is the maturity of the instrument. Assuming that there is no wrong-way risk (WWR),[3]. i.e., the default is independent of both the portfolio value and the numeraire, and application of Bayes' theorem, the CVA can be expressed as

$$\text{CVA} = (1 - R) B(0) \int_0^T \mathbb{E}\left[\frac{\text{E}(t)}{B(t)} \,\middle|\, t = \tau_C \right] \mathbb{E}[\delta(t - \tau_C)] \, dt.$$

The conditional expectation $\mathbb{E}[\cdot \,|\, t = \tau_C]$ is the current expected exposure at time t given that counterparty C defaulted at time t, i.e., $t = \tau_C$. The second expectation within the integrand is the counterparty C's default probability function, i.e., $PD(t)$ in (7). The CVA can therefore be written as,

$$\text{CVA} = (1 - R) B(0) \int_0^T \mathbb{E}\left[\frac{\text{E}(t)}{B(t)} \,\middle|\, t = \tau_C \right] dPD(t). \qquad (8)$$

---

[2]Basel III states that "Whenever such a CDS spread is not available, the bank must use a proxy spread that is appropriately based on the rating, industry and region of the counterparty". Calibration methods ranked from best to worst, first, from CDS spreads (if traded and quoted in the marked), second, from bond spreads (if traded and quoted in the marked), and third, from a rating transition matrix and last, from proxies such as stock prices or reported fundamental data.

[3]The International Swaps and Derivatives Association (ISDA) defines the wrong-way risk as "the risk that occurs when exposure to a counterparty is adversely correlated with the credit quality of that counterparty". If these two effects tend to happen together, then that co-dependence will increase the CVA on the contract and it will make the CVA larger than when the effects were independent. For details on WWR see for instance Hull and White (2012), Rosen and Saunders (2012), Redon (2006)

Let the expected exposure (EE) and the discounted expected exposure $\text{EE}^*(t)$ at time t be given by

$$
\begin{aligned}
\text{EE}(t) &= \mathbb{E}\left[\,\text{E}(t)\,|\,t = \tau_C\,\right], \\
\text{EE}^*(t) &= B(0)\,\mathbb{E}\left[\,\frac{\text{E}(t)}{B(t)}\,\bigg|\,t = \tau_C\,\right].
\end{aligned}
$$

Then, for a discrete time grid $0 = T_0 < T_1 < \cdots < T_m = T$ of observation dates Equation (8) can be approximated by

$$
\text{CVA} \approx (1 - R) \sum_{n=1}^{m-1} \text{EE}_n^* \left(PD_{n+1} - PD_n\right), \tag{9}
$$

where $\text{EE}_n^* = \text{EE}^*(T_n)$, to highlight that we work on a discrete time grid.

CVA can be seen as the weighted average of the discounted expected exposure with the weights given by the default probabilities. The complexity of CVA estimation lies within the evaluation of the exposure $\text{E}(t)$. Market practice is by American Monte Carlo methods where a large number of market scenarios of factors such as yield and inflation curves, FX rates, equity and commodity prices, credit spreads and others are simulated.

Next to EE and EE^*, trading desks are interested in additional exposure profiles such as the PFE. For a given date t, the α-percentile PFE_α is the maximum exposure of a portfolio with a high degree of statistical confidence α defined as,

$$
\text{PFE}_\alpha(t) = \inf\{x : \mathbb{P}\left(\text{EE}(t) \le x\right) \ge \alpha\}, \quad 0 \le t \le T.
$$

where $\mathbb{P}$ is the historical probability measure.

3 Monte Carlo Simulation of Counterparty Credit Risk

In this section we summarize the Least Squares Method (LSM) by Longstaff and Schwartz (2001) and present a version of the Stochastic Grid Bundling Method (SGBM) algorithm suitable for CVA calculation of Bermudan swaptions.

There are two choices for estimating the exposures on future scenarios, where the first approach includes all the payments including the one at the observation date, while the second approach only includes the future payments with respect to the observation date. We stay with the latter approach, which in case of cash settled early exercise options implies that the exposure of the option, if not exercised, is equal to its corresponding continuation value along the scenario at time T_n, i.e., $H_n = E_n$. If exercised at the observation date, we

assume no exposure for the option.

We let market state variable r_n represent the simulated market information at time T_n, and in our case, they are the short rates simulated using the one-factor GSR model in Equation (1).

3.1   The Least Squares Method (LSM)

The LSM, introduced by Carriere (1996) and popularized by Longstaff and Schwartz (2001), is a simulation-based method where one approximates the holding value H_n at each exercise time T_n of a Bermudan option using parametric functions. The parametric functions are approximated using least squares regression, giving the continuation value to have the form,

$$H_n = \sum_{i=0}^{q} \alpha_{i,n} \zeta_{i,n}, \tag{10}$$

for a set of q basis functions $\zeta_{i,n} : \mathbb{R}^d \rightarrow \mathbb{R}$, $i = 0, 1, \ldots, q$, and regression coefficients $\alpha_{i,n}$. The basis functions $\zeta_{i,n}$ are usually polynomials of the simulated state variables, in our case the short rates, e.g. $\zeta_{i,n} = r_n^i$ The regression coefficients are determined, when moving backwards in time, by minimizing

$$\sum_{\omega \in \Omega_{ITM}(T_n)} \left(\sum_{i=0}^{q} \alpha_{i,n} \zeta_{i,n}(\omega) - U_{\tau_n(\omega)}^*(\omega) \right),$$

where $\Omega_{ITM}(T_n)$ is the subset of paths on which the swaption is in-the-money (ITM) at time T_n, and

$$\tau_n(\omega) = \min\left(T_j : U_j(\omega) \geq H_j(\omega), \ j = n+1, \ldots, m \right).$$

$U_n^*(\omega)$ represents the corresponding future cashflows discounted along the path ω, given the observation date T_n. For the purpose of CVA, we do not make a restriction on the paths used for regression (based on whether or not they are ITM). Additionally, to avoid oversight bias, a second set of paths is generated, and regression functions from the initial simulation are used to approximate the continuation value and hence the exposures along the scenarios.

3.2   The Stochastic Grid Bundling Method (SGBM)

SGBM is a simulation-based dynamic programming method, which first generates Monte Carlo paths, forward in time (when the diffusion process appears in closed-form or in approximated closed-form the sample paths can be generated directly). This is then followed by finding the optimal early-exercise policy

by moving backwards in time. Although SGBM, like LSM, uses least squares regression to approximate parametric functions, the two approaches are significantly different. In loose terms, there are two key differences. First, in the case of LSM, regression is performed on the discounted future cashflows, while in the case of SGBM regression is directly performed on the value function. Second, in LSM, the regressed function is an approximation of the continuation value, and it is used for making the early exercise decisions. In the case of SGBM, the regressed function is an approximation of the option value function in a reduced space. The continuation value for a particular exposure date is determined as the conditional expectation of this regressed functional approximation on the next exposure date. A more detailed description of SGBM can be found in Jain and Oosterlee (2015).

In SGBM the exposure (continuation value), E_n at time T_n, is calculated using the law of iterated expectations, that is,

$$
\begin{aligned}
E_n &= B(T_n)\,\mathbb{E}\left[\left.\frac{V_{n+1}}{B(T_{n+1})}\right| r_n\right] \\
&= B(T_n)\,\mathbb{E}\left[\left.\mathbb{E}\left[\left.\frac{V_{n+1}}{B(T_{n+1})}\right|\zeta_{n+1}, r_n\right]\right| r_n\right],
\end{aligned}
\tag{11}
$$

where $\zeta_n = \left(\zeta_{0,n},\ldots,\zeta_{q,n}\right)^{\top}$ is a q-dimensional set of basis functions. For Bermudan swaptions in the one-factor Gaussian short rate model, we take a polynomial of the short rates as the basis functions. Writing the expected exposure as in Equation (11), decomposes the problem into two steps. The first step involves computing the inner conditional expectation,

$$
Z_{n+1} = \mathbb{E}\left[\left.\frac{V_{n+1}}{B(T_{n+1})}\right|\zeta_{n+1}, r_n\right],
\tag{12}
$$

which is followed by the computation of the outer expectation,

$$
E_n = B(T_n)\,\mathbb{E}\left[Z_{n+1}|r_n\right].
\tag{13}
$$

By carefully selecting the basis functions, Equation (13) can be computed in "closed-form". However, numerical approximations are required to calculate Z_{n+1} in Equation (12).

In order to compute Z_{n+1} in Equation (12), V_{n+1} needs to be conditioned on r_n. If computational costs were not a concern, this would imply simulating a new set of scenarios originating from each $r_n(\omega)$ and projecting the corresponding V_{n+1} for these sub-scenarios onto the basis functions ζ_{n+1}. This would result in a regressed function for each outer scenario $r_n(\omega)$ at T_n. However, nested Monte Carlo simulation is computationally inefficient as the number of paths grows exponentially with each time step.
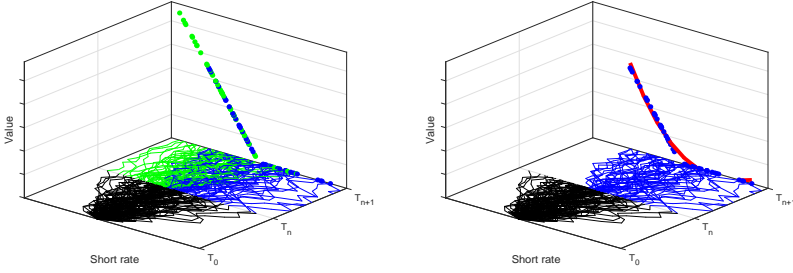
Figure 2: Illustrates SGBM. Left: Simulated paths at time T_n are clustered into two distinct bundles (green and blue). Right: The continuation values are approximated by a polynomial for the paths originating from the blue bundle. The procedure is repeated for each bundle.

A practical approach to condition V_{n+1} on r_n is to use bundling techniques. Bundling was introduced by Tilley (1993) and extended to higher dimensions in the State Space Partitioning Method (SSPM) in Jin, Tana and Sun (2007), and is a method to partition the state space into non-overlapping regions, so that each point in the space can be identified to lie in exactly one of the bundled regions. The intuitive idea behind bundling is that for $r_n(\omega)$, if the neighbouring paths are grouped together, the resulting distribution of paths at the next time step, in the limiting case of infinite scenarios and bundles, would be similar to the one obtained if new scenarios were generated starting from $r_n(\omega)$.

At each time T_n, the paths $r_n(\omega_k)$, for $k = 1,\dots,K$, are clustered into sets of non-overlapping bundles $\mathscr{B}(T_n)$. We bundle the grid points at each time step using the *recursive bifurcation algorithm*, explained in Appendix B. The number of bundles, after p iterations, equals 2^p. Figure 2 illustrates the bundling and regression procedure using two bundles. The computational complexity for the bundling is linear in the total number of grid points K, the dimensions d, and the number of iteration steps p. This makes the method of bundling practical and fast.

The inner expectation, Z_{n+1} given by Equation (12) is then approximated onto a polynomial subspace where the values are linear combinations of the basis functions. This is done by regressing locally, within each bundle, the option values, divided by the corresponding bank account process, at T_{n+1} for those paths that originate from the s-th bundle which contains $r_n(\omega_k)$, that is

$$\widehat{Z}_{n+1}^s = \sum_{i=0}^{q} \alpha_{i,n+1}^s r_{n+1}^i, \quad r_n(\omega_k) \in \mathscr{B}_n^s, \tag{14}$$

such that the following residual is minimized

$$\min_{\alpha^s} \sum_{r_n(\omega) \in \mathscr{B}_n^s} \left(\widehat{Z}_{n+1}^s(\omega) - \frac{V_{n+1}(\omega)}{B(T_{n+1},\omega)} \right)^2. \tag{15}$$

The exposure at a grid point $r_n(\omega_k)$ that belongs to bundle $\mathscr{B}_n^s$ is then approximated by,

$$
\begin{aligned}
\widehat{E}_n(\omega_k) &= B(T_n,\omega_k)\mathbb{E}\left[\widehat{Z}_{n+1}^s \big| r_n(\omega_k)\right] \\
&\approx B(T_n,\omega_k)\sum_{i=0}^{q} \alpha_{i,n+1}^s \mathbb{E}\left[\zeta_{i,n+1}\big| r_n(\omega_k)\right].
\end{aligned} \tag{16}
$$

Equation (16) converges to the true expected exposure, when the number of asset paths K and ithe number of bundles tend to infinity, see Jain and Oosterlee (2015) for details.

Once we have calculated exposures at each time step T_n using Equation (16) we can approximate the expected exposure as

$$\widehat{EE}_n \approx \frac{1}{K}\sum_{k=1}^{K} \widehat{E}_n(\omega_k),$$

and the discounted expected exposure as

$$\widehat{EE}_n^* \approx \frac{1}{K}\sum_{k=1}^{K} B(T_0)\frac{\widehat{E}_n(\omega_k)}{B(T_n,\omega_k)},$$

for $k = 1,\ldots,K$. Then, by using Equation (9), the CVA charges follow.

Regression-based American Monte Carlo methods depend on the choice of the regression variables. To avoid over-fitting one should not use too many regression variables since they are easily affected by outliers in the simulation. For Bermudan swaptions it is common to use a second-order polynomial (of the underlying swap value or the short rate) for the regression, see for instance, Glasserman and Yu (2004).

Remark 10 *As the regression approximation depends on a rather arbitrary choice of the basis functions, one should ideally have an estimate of both the upper and lower bound values for the true price. A lower bound for the option price can be computed using the so-called path-estimator approach, where the option value is computed as an expectation of the discounted payoffs from a sub-optimal exercise policy, see for example, Broadie and Glasserman (2004). The policy is sub-optimal, because of the numerical errors in its computation. One should*

use a fresh set of paths for the path-estimator, and not the same ones used to obtain the early-exercise policy, to avoid a foresight bias. An upper bound is found using the duality approach, based upon the work of Rogers (2002) and Haugh and Kogan (2004). This approach has moreover been extensively studied and described in Andersen and Broadie (2004) and Kolodko and Schoenmakers (2004). Belomestny, Bender and Schoenmakers (2009) present an efficient method for obtaining the upper bound using the duality approach, which can be used for Bermudan swaptions. The quality of the upper bound produced by the duality approach depends on the quality of the estimated exercise policy in the first pass, a more accurate policy gives tighter upper and lower bounds. In this paper, however, we focus only on the lower bounds and show that the ones obtained using SGBM, with significantly fewer paths, converge to the same lower bound value as those obtained using LSM. One can, in a relatively straightforward way, use the exercise policy obtained from SGBM in the duality approach to obtain a corresponding upper bound.

3.2.1   Hybrid Measure Monte Carlo

CVA calculations are done at netting set level, where the netting set can have several different types of deals and underlying driving risk-factors. Additionally, a CVA quote for a new deal, added or removed from an existing netting set, should be ideally priced in real-time. As the computational time for Monte Carlo simulations scales with number of scenarios, it is important that the standard error and bias of the results from the simulation are as small as possible. Variance reduction then is a highly desired feature for calculations related to CVA pricing.

An advantage of using SGBM is that it allows adapting the problem, to break the expectation, which would otherwise be solely computed using the Monte Carlo approach, to sub-problems where part of expectation is known in closed form. This feature helps in significantly bringing down the variance of the solution. In particular for the Bermudan swaptions, we employ hybrid measures to achieve variance reduction.

The T-forward measure, with corresponding expectation $\mathbb{E}^T$ and the T-maturity zero coupon bond $P(t, T)$ as the numeraire has the advantage that it allows for decoupling the payoff $V(T)$ from the numeraire and take out the discount factor from the expectation, i.e,

$$
\begin{aligned}
V(t) &= B(t)\mathbb{E}_t\left[\frac{V(T)}{B(T)}\right] \\
&= P(t, T)\mathbb{E}_t^T[V(T)].
\end{aligned}
$$

One benefit however of the spot measure compared to the T-forward measure is that the numeraire asset $B(t)$ is alive throughout the tenor and therefore allows for simulating paths irrespective of tenor. In SGBM, we will employ *hybrid measures* to obtain an efficient Monte Carlo simulation. In order to apply the hybrid measure we modify the inner-expectation, as given in Equation (12), to the following,

$$Z_{n+1} = \mathbb{E}[V_{n+1}|\zeta_{n+1}, r_n], \tag{17}$$

which is followed by the computation of the following outer expectation (as opposed to Equation (13)),

$$E_n = B(T_n)\,\mathbb{E}\left[\left.\frac{Z_{n+1}}{B(T_{n+1})}\right| r_n\right]. \tag{18}$$

The inner expectation is approximated by regression on short-rates simulated under the spot measure. Note that the minimization problem for regression problem changes from Equation (15) to:

$$\min_{\alpha^s} \sum_{r_n(\omega)\in\mathscr{B}_n^s} \left(\widehat{Z}_{n+1}^s(\omega) - V_{n+1}(\omega)\right)^2. \tag{19}$$

The outer-expectation in Equation (18) can be computed under the T-forward measure, rather than the spot measure. This would allow computing the expectation, without explicitly simulating the bank account process B. The exposure at grid point $r_n(\omega)$ that belongs to bundle $\mathscr{B}_n^s$ is therefore computed as,

$$
\begin{aligned}
\widehat{E}_n(\omega) &= B(T_n,\omega)\,\mathbb{E}\left[\left.\frac{\widehat{Z}_{n+1}^s}{B(T_{n+1})}\right| r_n(\omega)\right] \\
&= P(T_n, T_{n+1}, \omega)\,\mathbb{E}^{T_{n+1}}\left[\widehat{Z}_{n+1}^s \,\middle|\, r_n(\omega)\right] \\
&\approx P(T_n, T_{n+1}, \omega)\sum_{i=0}^{q} \alpha_{i,n+1}^s \mathbb{E}^{T_{n+1}}\left[\zeta_{i,n+1}\,\middle|\, r_n(\omega)\right].
\end{aligned}
\tag{20}
$$

As $\zeta_{i,n+1}$ is a polynomial function of the short-rates (simulated in the Gaussian factor model), its conditional moments are known in closed form under the T forward measure. An outcome of formulating the problem as above is that we only need to simulate the future option price, and not additionally the corresponding future bank account process, to obtain the option price on a given exposure date. As a result we achieve significant variance reduction in the exposure calculation when compared to plain LSM.

3.2.2   The SGBM-CVA Algorithm

As explained, SGBM computes the continuation value in two steps: we first compute the expected option value, conditioned on a finer information set, given by Equation (12), which is followed by the computation of the outer expectation, given by Equation (18).

The SGBM-CVA algorithm is therefore divided into two parts, a first and second pass. In the first pass, we perform a forward phase where K_1 Monte Carlo paths are simulated, future cash flows are calculated and a regression basis is constructed. This is subsequently followed by a backward phase, where we estimate the payoffs and the polynomials by regression. In order to get unbiased values and lower bound values, we perform a second pass with a new forward phase where we simulate K_2 Monte Carlo paths, evaluating the payoffs using the regression functions estimated in the first pass but with the new set of paths.

For clarity, we summarize the steps for the SGBM-CVA algorithm.

I. FIRST PASS: Estimate Regression Functions.

---

1. Generate K_1 paths $\omega_1, \dots, \omega_{K_1}$, using Equation (2).

2. For each path ω_k and time T_n, for $k = 1, \dots, K_1$ and $n = 1, \dots, N-1$, compute the state variable $r_n(\omega_k)$ and values $V_n(\omega_k)$, where V_N, is known and for $n = 1, \dots, N-1$, it is solved recursively as below.

3. For each $n = N-1 \dots, 1$,

   (a) Bundle the grid points at T_{n-1}, into a distinct bundles (except at T_0, where there is only one point) using the algorithm described in Appendix B.

   (b) Compute the regression functions, Z_n^s, $s = 1, \dots, a$, given by Equation (14), using the option values V_n at T_n for the paths originating from the s-th bundle, $\mathscr{B}_{n-1}^s$, at T_{n-1}.

   (c) Compute the E_n for the grid points in the s-th bundle at T_{n-1}, using Equation (20) for those paths for which $r_n(\omega_k)$ belongs to the bundle $\mathscr{B}_{n-1}^s$, for $s = 1, \dots, a$.

II. SECOND PASS: Estimate CVA.

---

1. In order to compute an unbiased CVA, generate a fresh set of K_2 paths $\omega_1', \dots, \omega_{K_2}'$, and compute new state variables $r\left(\omega_k'\right)$ and values $V_n\left(\omega_k'\right)$

2. For each $n = N - 1 \ldots, 1$,
   (a) Bundle the grid points at T_{n-1} using the same algorithm as in the first pass and described in Appendix B.
   (b) Compute the exposures for the grid points in bundle s, at time step T_{n-1}, using the regression function Z_n^s $s = 1, \ldots, a$, obtained in the first pass.
   (c) Compute the EE_n, EE_n^* and PFE_α (T_n) for the grid points in the s-th bundle at T_{n-1}, for those paths for which $r_n \left(\omega_k' \right)$ belongs to the bundle $\mathcal{B}_{n-1}^s$, for $s = 1, \ldots, a$.

3. The CVA charge is then calculated as,

$$\text{CVA} \approx (1 - R_C) \sum_{n=0}^{N-1} \text{EE}_n^* (PD_{n+1} - PD_n).$$

Remark 11 *Valuation of Bermudan swaptions with American Monte Carlo methods requires an estimate of the early exercise boundary. Exposure can then be seen as a barrier option (knock-in) with the estimated exercise boundary as the barrier. Once the option has been exercised (knocked) along a path at time T_n the exposure E_m at T_m for $T_n < T_m$ for that path becomes zero.*

Remark 12 *The market standards for swaptions are cash-settled contracts, i.e., contracts that settle into a cash payment when exercised. The benefit is that one avoids credit exposure (and the obligation of collateral posting due to the Credit Support Annex, or CSA) and therefore have zero exposure after the exercise date. For physically settled contracts, i.e., contracts entered into an interest rate swap when the contract is exercised, one would have to calculate the exposure of the swap from the exercise date throughout the swap tenor. The Bermudan swaption formulas in Section 2.2 describe physical-settled contracts. The standard pricing formulas for cash-settled agreements are not properly justified, since one would have to calculate the annuity $A_{n,m}$ by discounting at the fixed swap rate $S_{n,m} (T_0)$. Since the SGBM-CVA algorithm presented here works irrespectively of settlement type, we assume for simplicity that the annuity for cash- and physical-settled Bermudan swaptions are the same.*

4   Numerical Results

In this section we study the performance of SGBM-CVA by means of numerical experiments. The numerical examples presented below demonstrate the efficiency of calculating CVA using SGBM.

4.1   Setup

We use a one-dimensional market state variable $r(t)$ to represent the market information, and we let the short rates be simulated using the HW1F model by Hull and White (1990), which is commonly used for pricing Bermudan swaptions. Under HW1F the short rate dynamics are given by Equation (1) with κ and σ constant. We calibrate the HW1F model parameters to the initial zero coupon bond prices observed in the market 2 January 2014. For the default probability function in equation (7), we set the hazard rate $\gamma(t) = 0.05$, and the recovery rate $R_C = 0.40$.

For the LSM and SGBM regression, we use a third-order polynomial with the short rate as the basis and $\zeta_{i,n} = r_n^i$. The moments for the short rates under the HW1F dynamics in Equation (1) are given in Appendix C.

We consider Bermudan swaptions exercisable once a year with Moneyness (MN) i.e, the spot vs. strike ratio of 80%, 100% and 120%, and with realistic HW1F parameters $\kappa = 0.01, 0.02$ and $\sigma = 0.01, 0.02$.[4]

We simulate the first pass with $K_1 = 4096$ seeds using the Mersenne twister pseudo random number generator to estimate the regression functions. Subsequently, we simulate the second pass with $K_2 = 8192$ quasi-Monte Carlo Sobol paths using the regression function estimated in the first pass to estimate the unbiased Bermudan swaptions values, EE, PFE and CVA charges. Each test is repeated 100 times with different seeds in the first pass, to remove the overall influence of the Mersenne twister pseudo random number generator.

We use the bundling scheme described in Appendix B with 8 bundles and with the same number of bundles at each time step, except at time T_0, where there is only one point, r_0. We report the values obtained from the second pass. The prices are reported in basis points, with the notional $\mathcal{N} = 10,000$.

The variance reduction is defined as the ratio between the variance from LSM and the variance from SGBM, where both estimates are obtained from 100 simulations.

4.2   EE and PFE values

Figure 3 illustrates the $\text{PFE}_{5\%}$, $\text{PFE}_{95\%}$ and EE values generated by LSM and SGBM for 5Y, 10Y, 15Y and 20Y Bermudan swaptions. We observe that both methods generate the same values and the characteristic shapes, i.e., the exposure tends to increase first, since there is an increased probability that the Bermudan swaptions will be deeper in-the-money at a future exercise date.

The efficiency of SGBM compared to LSM for estimating $\text{PFE}_{5\%}$, $\text{PFE}_{95\%}$ and EE is illustrated in Figure 4. Clearly, LSM is affected by outliers for the high quan-

---

[4]For instance, at the beginning of 2015, the HW1F, with value of κ and σ calibrated to USD, co-terminal swaptions were both around 0.01.

tile PFE estimation. For the EE by SGBM, we obtain on average a variance re-
duction of a factor 100.

4.3 CVA

For the CVA computations, we consider Bermudan swaptions with maturities
of 5Y and 10Y. Tables 1 and 2 report the lower bound values for the Bermudan
swaptions and CVA charges via LSM and SGBM. The numbers in parentheses
are sample standard deviations and the values from LSM and SGBM differ at
most 5 bps. As a first observation, the standard deviation for the SGBM lower
bounds is much smaller than the ones obtained from LSM. The efficiency of
SGBM compared to LSM for pricing and CVA calculation is illustrated in Figure
5. For the lower volatility scenarios, i.e., $\sigma = 0.01$ we obtain for the 5Y Bermudan
swaption CVA a variance reduction of a factor 200 and for the 10Y a factor of 400.
For the high volatility case, i.e., with $\sigma = 0.02$ we observe a variance reduction of
a factor 100 for the 5Y and 200 for the 10Y test case. The interpretation here is
that for a 10Y Bermudan swaption under HW1F with $\sigma = 0.01$ we will on average
need 400 times more Monte Carlo seeds for LSM compared to SGBM in order to
obtain equally "accurate" CVA values.

4.4 Approximation Error

For the approximation error we study the convergence by increasing the num-
ber of paths in the first and the second pass. As the *"true"* value, we select the
mean of the LSM computations with $K_1 = 131,072$ and $K_2 = 2K_1$, repeated 100
times. Then, for different values of K_1, with $K_2 = 2K_1$, we repeat the simula-
tion 100 times, and estimate the relative error with respect to the *"true"* value,
for LSM and SGBM with 1, 2, 4, 8 and 16 bundles. In Figure 6, we illustrate the
mean and the standard deviation of the error for a 5Y Bermudan swaption with
$\kappa = 0.01$, $\sigma = 0.01$ and an MN value of 100%. One can observe that LSM re-
quires a large number of paths to converge to the true value. For SGBM-1 (i.e.
SGBM with 1 bundle) we see an upward-biased value, but we observe a signif-
icant improvement in convergence by SGBM-2 which converges at $K_1 = 4096$
demonstrating essentially the same accuracy as LSM in the case of $K_1 = 131,072$
paths. The error is further reduced by increasing the number of bundles and the
computations with 4, 8 and 16 bundles converge at $K_1 = 16384$ paths. SGBM-16
is slightly upward-biased for small numbers of paths, most likely because some
bundles will then contain too few paths to allow a feasible regression without
too large error. It can be seen that the SGBM-16 is slightly upward-bias for low
number of paths compared to SGBM-4 and SGBM-8. We observe similar pat-
terns and convergence for different MN values, maturities and parameter setup,
when the number of paths per bundle is too small.

MN	κ	σ	SGBM		LSM		CVA SGBM		CVA LSM	
0.8	0.01	0.01	477.30	(0.17)	477.39	(3.90)	51.70	(0.014)	51.70	(0.21)
1.0	0.01	0.01	548.12	(0.18)	548.26	(4.13)	58.58	(0.014)	58.53	(0.22)
1.2	0.01	0.01	599.25	(0.18)	599.65	(4.32)	63.41	(0.014)	63.35	(0.21)
0.8	0.01	0.02	736.43	(0.46)	737.74	(7.09)	76.31	(0.035)	75.41	(0.36)
1.0	0.01	0.02	801.25	(0.46)	803.21	(7.14)	82.29	(0.033)	83.04	(0.38)
1.2	0.01	0.02	846.78	(0.50)	847.85	(7.32)	86.45	(0.034)	86.16	(0.35)
0.8	0.02	0.01	471.07	(0.18)	471.89	(4.18)	51.06	(0.02)	51.32	(0.22)
1.0	0.02	0.01	542.15	(0.18)	542.76	(4.09)	57.99	(0.01)	58.01	(0.22)
1.2	0.02	0.01	593.49	(0.18)	594.28	(4.52)	62.86	(0.01)	63.06	(0.21)
0.8	0.02	0.02	723.13	(0.45)	726.45	(6.90)	75.03	(0.03)	74.86	(0.37)
1.0	0.02	0.02	788.16	(0.46)	791.75	(6.85)	80.96	(0.03)	80.88	(0.35)
1.2	0.02	0.02	833.95	(0.46)	837.23	(6.92)	85.13	(0.03)	85.63	(0.39)

Table 1: Lower bound 5Y Bermudan swaption risk-free prices and CVA charges under HW1F using LSM and SGBM. Prices are in basis points and standard deviations within parentheses.

MN	κ	σ	SGBM		LSM		CVA SGBM		CVA LSM	
0.8	0.01	0.01	947.30	(0.23)	946.74	(7.87)	175.04	(0.03)	175.41	(0.66)
1.0	0.01	0.01	1187.0	(0.24)	1186.7	(8.77)	215.55	(0.03)	214.56	(0.75)
1.2	0.01	0.01	1367.7	(0.23)	1368.2	(9.04)	245.44	(0.03)	244.85	(0.70)
0.8	0.01	0.02	1584.1	(0.54)	1586.4	(13.25)	283.59	(0.07)	282.52	(1.15)
1.0	0.01	0.02	1805.5	(0.55)	1809.4	(13.59)	319.16	(0.08)	319.46	(1.14)
1.2	0.01	0.02	1966.0	(0.53)	1968.5	(14.52)	344.73	(0.07)	343.11	(1.15)
0.8	0.02	0.01	921.21	(0.23)	920.38	(8.21)	170.12	(0.03)	170.34	(0.72)
1.0	0.02	0.01	1162.7	(0.22)	1161.2	(8.59)	211.08	(0.03)	210.45	(0.72)
1.2	0.02	0.01	1345.0	(0.24)	1344.3	(8.83)	241.23	(0.03)	241.07	(0.71)
0.8	0.02	0.02	1529.5	(0.51)	1536.7	(12.96)	273.61	(0.07)	274.5	(1.13)
1.0	0.02	0.02	1752.3	(0.54)	1758.7	(14.18)	309.44	(0.07)	311.73	(1.16)
1.2	0.02	0.02	1914.2	(0.56)	1919.6	(13.73)	335.13	(0.07)	334.72	(1.20)

Table 2: Lower bound 10Y Bermudan swaption risk-free prices and CVA charges under HW1F using LSM and SGBM. Prices are in basis points and standard deviations within parentheses.
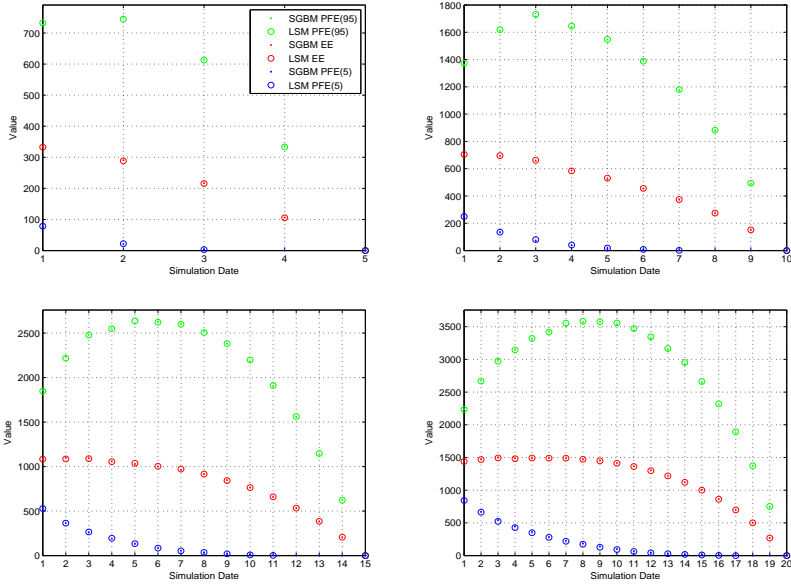
Figure 3: Bermudan swaption EE, PFE$_{5\%}$ and PFE$_{95\%}$ under HW1F with $\kappa = 0.01$ and $\sigma = 0.01$ with notional $\mathcal{N} = 10,000$. Upper Left: 5Y Maturity. Upper Right: 10Y Maturity. Lower Left: 15Y Maturity. Lower Right: 20Y Maturity.

SGBM demonstrate a faster convergence and produces more stable values with significant lower variances. The reason is that LSM uses the regressed continuation values directly to make the early-exercise decision. A large number of paths and basis functions are required to reduce the noise in this regressed function. Therefore, the quality of the LSM early-exercise policy may not be accurate for a small number of paths and basis functions. In SGBM, however, the regressed function is just the inner expectation, which is not used for decision-making. The outer expectation, which can be analytically computed, gives the continuation value and is used to make the early exercise decision. As the noise, or the error due to regression, is normally distributed with a zero mean, the outer expectation of the noise would be zero.

5   Conclusion

Usually banks have a large number of trades in a portfolio and it would be computationally inefficient to require several runs for the trades in the port-
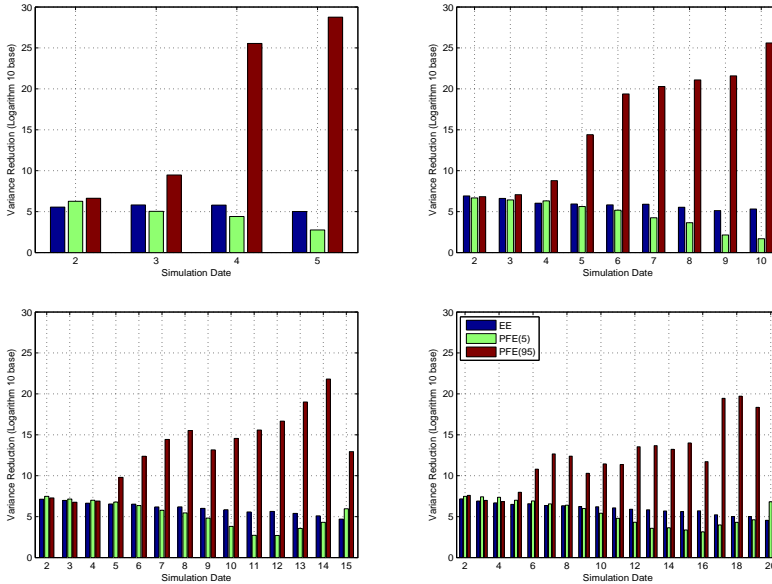
Figure 4: EE, PFE$_{5\%}$ and PFE$_{95\%}$ variance reduction for a Bermudan swaption under HW1F with $\kappa = 0.01$, $\sigma = 0.01$ and notional $\mathcal{N} = 10,000$. Upper Left: 5Y Maturity. Upper Right: 10Y Maturity. Lower Left: 15Y Maturity. Lower Right: 20Y Maturity.

folio to get a CVA which we can be confident about if there is high variance. This paper presented the application of the Stochastic Grid Bundling Method (SGBM) for calculating exposures, potential future exposure and approximating CVA charges for Bermudan swaptions in an American Monte Carlo simulation framework. SGBM is a regression-based Monte Carlo method which is accurate and easy to implement. Variance reduction, based on iterated conditioning, in combination with the bundling technique form the necessary ingredients for accurate CVA valuation with a relative small number of paths and basis functions. The computational time for the method is comparable to the least squares method in Longstaff and Schwartz (2001), but a higher accuracy is achieved. Our numerical examples demonstrate the efficiency of calculating CVA using SGBM, making it a very suitable candidate with a potential to calculate *"real-time"* CVA charges and easy extension to other charges within the XVA family.
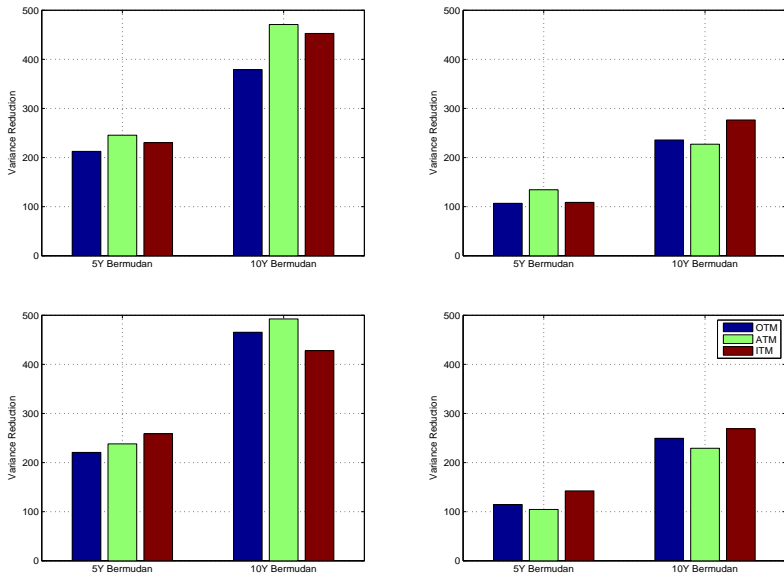
Figure 5: CVA variance reduction for 5Y and 10Y Bermudan swaptions under HW1F with notional $\mathcal{N} = 10,000$. Upper Left: $\kappa = 0.01$, $\sigma = 0.01$. Upper Right: $\kappa = 0.01$ $\sigma = 0.0.2$. Lower Left: $\kappa = 0.02$ $\sigma = 0.01$. Lower Right: $\kappa = 0.02$, $\sigma = 0.02$.
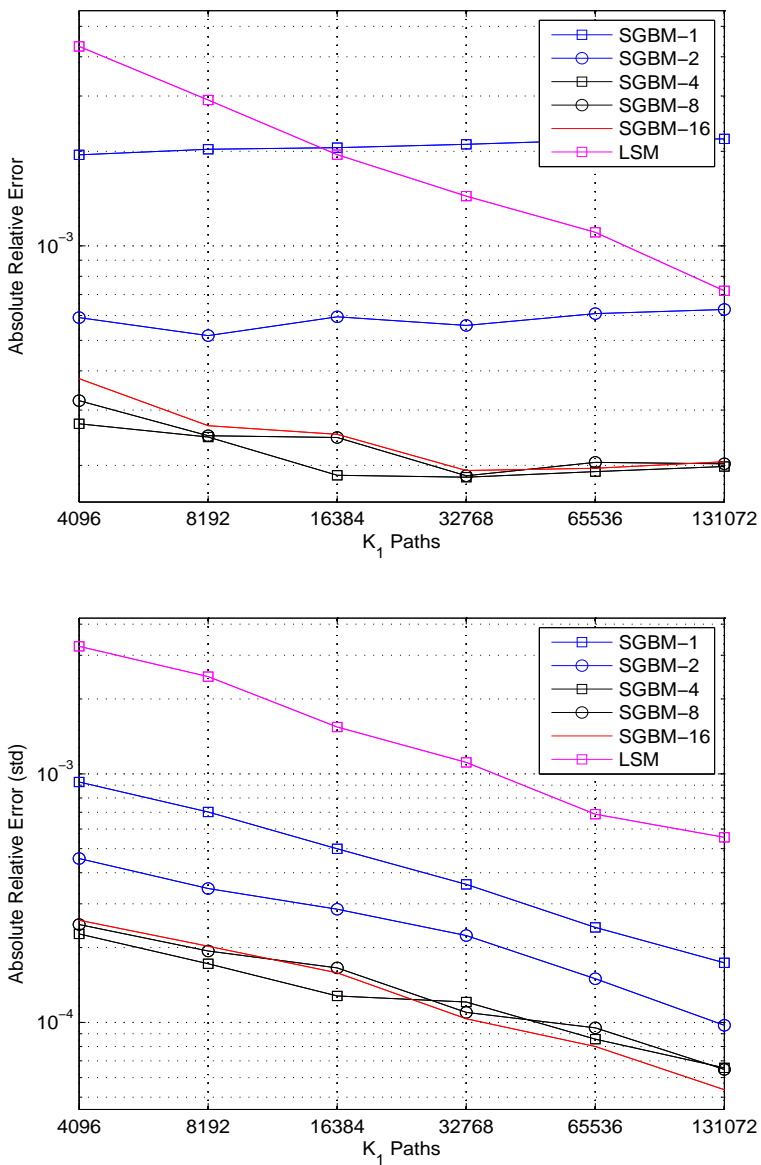
Figure 6: CVA error. Mean (upper figure) and standard deviation (lower figure) for 5Y Bermudan swaption.

B Bundling

Suppose we need to bundle K_s grid points at epoch T_n, given by $S(T_n, \omega_k)$, where $k = 1, \ldots, K_s$. The following steps are performed recursively.

1. Compute the *mean* of the given set of grid points,

$$\mu_n^s = \frac{1}{K_s} \sum_{k=1}^{K_s} S(T_n, \omega_k).$$

2. Bundling the grid points is performed by dividing the grid points into two groups, depending on whether the asset price for the grid point is greater or less than the mean of the asset prices for the given set of grid points:

$$\begin{aligned} \mathscr{B}^1(T_n, \omega_k) &= \mathbf{1}\left(S(T_n, \omega_k) > \mu_n^s\right), \\ \mathscr{B}^2(T_n, \omega_k) &= \mathbf{1}\left(S(T_n, \omega_k) \leq \mu_n^s\right), \end{aligned}$$

for $k = 1, \ldots, K_s$. $\mathscr{B}^1(T_n, \omega_k)$ returns 'true', when the asset price $S(T_n, \omega_k)$ is greater than the mean, μ_n^s and belongs to bundle 1. $\mathscr{B}^2(T_n, \omega_k)$ returns 'true', if it less than the mean and belongs to bundle 2. Formally, $\mathscr{B}^s(T_n, \omega_k)$ returns 'true', if the grid point $S(T_n, \omega_k)$ belongs to bundle s.

3. Bundles $\mathscr{B}^1(T_n)$ and $\mathscr{B}^2(T_n,)$ can be split again, returning to step 1.

C HW1F Moments

Let $M_k(s, t) = \mathbb{E}\left[r(t)^k \mid s\right]$ be the k-th moment. The three first moments for the HW1F are given by,

$$\begin{aligned} M_1(s, t) &= e^{-\kappa(t-s)} r(s) + \frac{\theta(t)}{\kappa}\left(1 - e^{-\kappa(t-s)}\right), \\ M_2(s, t) &= M_1^2(s, t) + \frac{\sigma^2}{2\kappa}\left(1 - e^{-2\kappa(t-s)}\right), \\ M_3(s, t) &= M_1^3(s, t) + 3M_1(s, t)\left(M_2(s, t) - M_1^2(s, t)\right). \end{aligned}$$

PAPER 4

Finite Element Based Monte Carlo Simulation of Option Prices on Lévy Driven Assets

Abstract

This paper extends the simulation algorithm by Andreasen and Huge (2011) to the simulation of option prices and deltas on Lévy driven assets. The simulation is performed and relies on the inverse transition matrix of a discretised partial differential equation (PDE). We demonstrate how one can obtain accurate prices and deltas of European options on the variance gamma (VG) and CGMY model through finite element-based Monte Carlo simulations.

1   Introduction

The traditional method of pricing and hedging over-the-counter (OTC) deriva-
tives originates from a partial differential equation (PDE). The PDE describes
option prices over time given certain underlying dynamics of the prices process,
such as a geometric Brownian motion in Black and Scholes (1973), a stochastic
volatility model in Heston (1993), a local volatility model in Dupire (1994), a
jump model in Merton (1976), Carr, Chang, and Madan (1998) and Carr et al.
(2002), or a combination of these processes. The model is calibrated to a set
of market data such that it can regenerate significant market features, including
the volatility smile (and skew). The calibration needs to be fast and is performed
in the frequency domain most of the time, such as by using the methods by Carr
and Madan (1999) and Fang and Oosterlee (2008). After calibrating the model,
one turns to the main task of pricing and hedging exotic derivatives. This is ac-
complished by simulating a discretisation of the stochastic differential equation
(SDE) using Monte Carlo methods, or by a discretisation of the PDE to iteratively
solve for the evolution of option prices and sensitivities with respect to bound-
ary conditions.

A drawback with the traditional method is the inconsistency between cal-
ibration and the pricing step because they rely on two different types of dis-
cretisations and, therefore, generate two different types of discretisation er-
rors. Andreasen and Huge (2011) demonstrate a calibration and pricing method
through which discretisation errors from the two steps are fully consistent with
each other. They consider a PDE driven by an underlying local stochastic volatil-
ity model. They also describe a method to calibrate the model to plain vanilla
options and an algorithm for pricing exotic derivatives such that both steps rely
on the same discretisation and for which the Monte Carlo simulation of op-
tion prices is performed using the discrete formulation of an equivalent Black-
Scholes PDE. More specifically, each row of the inverse transition matrix de-
scribes the evolution of option prices and, therefore, can be seen as the under-
lying's transition probability distribution and from where the samples are gen-
erated.

The Finite Difference Method (FDM) and the Finite Element Method (FEM)
are two widely used methods for numerically solving PDEs. FDM was first ap-
plied to options in Brennan and Schwartz (1978) and consists of approximating
the solution on a grid by replacing the derivatives in the PDE with difference
quotients. Wilmott, Howison and Dewynne (1993) introduced FEM to option
pricing. FEM is a two-step procedure in which one first performs a discretisa-
tion in the price domain. In other words, a variational formulation of the PDE
is found and then the solution is approximated using piecewise polynomials to
obtain a system of coupled ordinary differential equations (ODEs). Second, one
performs a time discretisation to solve the ODEs.

FDM is the most widely used method for solving PDEs and has gained in

popularity because of its simplicity to implement. However, this advantage is lost with exotic derivatives. FDM requires that the payoff, i.e. terminal and boundary conditions, must be sufficiently smooth to guarantee the existence and uniqueness of a solution. For simple products, such as vanilla European and digital options, a problem already exists because of the discontinuous payoff at the strike (e.g. the slope of a European call option is discontinuous at the strike). For crude discretisation, this discontinuity will give rise to the *odd–even effect*, where the solution will jump up and down as the discretisation grid is increased. These shortcomings can be handled by FEM, and represent one reason why option pricing with FEM has increased within computational finance. As mentioned, the advantages of FEM over FDM is the low smoothness assumptions on terminal conditions (e.g. the option payoff) and the faster obtained convergence rate, as shown in, for example, Hilber, Kehtari, Schwab and Winter (2010).

This paper focuses on the simulation of jump diffusion models, which is a process characterised by a Lévy measure and for which the evolution of option prices are characterised by a partial integro-differential equation (PIDE). Jumps in option pricing were introduced by Merton (1976) and assume that the non-normality in returns could be captured by a jump-diffusion process of the finite activity type. The direction of the jumps has further been extended with infinite activity jumps by the variance gamma (VG) model in Carr, Chang, and Madan (1998) and its extension, the Carr-Geman-Madan-Yor (CGMY) model in Carr et al. (2002). These models allow for jump components of finite and infinite activity type and with finite or infinite variation. Jumps have been proven very useful in capturing the extreme implied smile and skew typically observed for short-dated options that cannot be explained by a normal or log-normal model; see, for instance, Tompkins (2001). FEM and Lévy processes were previously studied in Matache, Petersdorff and Schwab (2004), Hilber, Kehtari, Schwab and Winter (2010), and Achdou and Pironneau (2005). A limited number of papers describe the simulation of the Lévy process. Examples include Asmussen and Rosinski (2001), who splits the Lévy process into two parts: one part consists of small jumps that are simulated using a diffusion process and the second part consists of large jumps that are simulated as a compound Poisson process (will not capture infinite activity). Ribeiro and Webber (2004) demonstrates an accurate and fast simulation algorithm for the VG using Gamma bridges and the difference between two Gamma processes. Monte Carlo simulation schemes for some general Lévy processes can also be found in Schoutens (2003).

This paper contributes to the simulation of Lévy processes by extending the simulation algorithm in Andreasen and Huge (2011) to the simulation of option values and sensitivities for which the underlying dynamics are specified by a Lévy process. As reference models, we consider VG and CGMY. In this paper, we discretise the VG and CGMY PIDE using FEM and analytically solve the inte-

gral arising in the PIDE attributable to the jumps. With the discrete formulation of the PIDE, we thereafter demonstrate an algorithm that allows for efficient and accurate simulation of option values and sensitivities on Lévy-driven assets. The main advantage is that the MC schemes are fully consistent with the discretised PIDE. We follow Hilber, Kehtari, Schwab and Winter (2010) and discretise the PIDE using the Galerkin method.

This paper is organised as follows. Section 2 provides a quick introduction of Lévy processes. Section 3 describes numerical methods for pricing options on Lévy-driven assets. Section 4 describes how the random grids algorithm in Andreasen and Huge (2011) can be applied to simulate option prices and Greeks on Lévy-driven assets. In Section 5, we present various numerical examples to illustrate the method and finally conclude in Section 6.

2 Lévy Processes

This section describes the fundamental theory of Lévy processes and provides a brief overview of two important Lévy processes: the VG and the CGMY model.

In general, any Lévy process $X = \{X_t : t \geq 0\}$ is completely identified by its characteristic triplet (γ, σ^2, ν) and can be written as

$$X_t = \gamma t + \sigma W_t + Z_t, \tag{1}$$

where γ is called the mean correcting martingale parameter, W_t is a Brownian motion, σ is the volatility, and Z_t is a jump process described using a Lévy measure, $d\nu(x) = k(x) dx$, for some Lévy density $k(x)$. The characteristic exponent ψ is defined by the Lévy-Khintchine representation $\mathbb{E}\left[e^{iuX_t}\right] = \exp\left(t\psi(u)\right)$, and given by

$$\psi(u) = i\gamma u - \frac{\sigma^2 u^2}{2} + \int_{\mathbb{R}} \left(e^{iuz} - 1 - iuz1_{|z|\leq1}\right) d\nu(z). \tag{2}$$

For the price process, we assume underlying risk-neutral spot dynamics with zero dividends, defined as an exponential Lévy process

$$S_t \quad = \quad S_0 \exp(rt + X_t), \tag{3}$$

where r is the risk free rate, and γ is chosen such that $\exp(X_t)$ is a martingale in the risk-neutral measure, which holds for

$$\gamma = -\frac{\sigma^2}{2} - \int_{\mathbb{R}} \left(e^z - 1 - z\right) d\nu(z). \tag{4}$$

The option value $V = V(t, x)$ at time $t \in (0, T)$ with payoff $g(\cdot)$ is given by the

conditional expectation

$$V(t, x) = \mathbb{E}\left[e^{-r(T-t)} g\left(e^{r(T-t)+X_T}\right) \big| X_t = x \right].$$ (5)

Where $\mathbb{E}[\cdot | X_t] = \mathbb{E}[\cdot | \mathscr{F}_t]$ is the conditional expectation with respect to the risk-neutral distribution and $\mathscr{F}_t$ is the filtration generated by X_t. By the Feynman-Kaĉ representation theorem for Lévy processes X, V given by Equation (5) is a solution to the boundary value problem as indicated in

$$\frac{\partial V}{\partial t} - \mathscr{A}V - rV \quad = \quad 0, \quad \text{in } (0, T) \times \mathbb{R},$$ (6)

where the integro-differential operator is given by

$$
\begin{aligned}
\mathscr{A}V(x) \quad = \quad & \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial x^2} V(x) + \gamma \frac{\partial}{\partial x} V(x) \\
& + \int_{\mathbb{R}} \left(V(x+z) - V(x) - z\frac{\partial}{\partial x} V(x) \right) dv(z),
\end{aligned}
$$ (7)

for functions $V \in C^2(\mathbb{R})$ with bounded derivatives, and where the boundary condition satisfies the terminal condition $V(T, x) = g(e^x)$. One obtains the Black-Scholes equation for $v(z) = 0$.

Next, we provide a brief overview of two important examples of Lévy processes, the VG and CGMY models.

2.1 The VG Model

The VG model in Carr, Chang, and Madan (1998) is a continuous pure-jump Lévy process $X^{VG} = \{X_t^{VG}, t \geq 0\}$ of infinite activity type with independent and stationary VG distributed increments.[1] The VG is a popular model for option pricing because of its analytical tractability and its allowance for flexible parameterisations of the skewness and kurtosis increments. The characteristic function of the VG process is given by

$$\psi_{VG}(u) \quad = \quad \left(\frac{GM}{GM + (M-G)\,iu + u^2} \right)^C,$$ (8)

---

[1]A process is said to be of finite (infinite) activity type if the process have a finite (infinite) number of jumps along any finite time interval. A process with finite (infinite) variation have a finite (infinite) variance along any finite time interval

for $C > 0$ and $G, M \geq 0$; see, for instance, Carr et al. (2002) for further details. The Lévy density is given by

$$k_{\text{VG}}(z) = C \frac{\exp(G|z|)}{|z|} 1_{\{z<0\}} + C \frac{\exp(M|z|)}{|z|} 1_{\{z>0\}}. \tag{9}$$

The mean correcting martingale parameter γ_{VG} for the VG process is given by

$$\gamma_{\text{VG}} = \frac{-C\left(G\left(\exp(-M)-1\right) - M\left(\exp(-G)-1\right)\right)}{MG}. \tag{10}$$

The parameters G and M control the skewness and exponential decay, or the right and left tail behaviours. For $G = M$, we have a symmetric Lévy measure, for $G > M$ the right tail is heavier than the left one, and for $G < M$ the left tail is heavier than the right one.

2.2   The CGMY Model

Carr et al. (2002) introduces the CGMY model, which is a pure-jump process that allows for both finite and infinite variations. The model extends the VG process by introducing an additional parameter Y, which defines the fine structure of the process and determines whether the process is of finite or infinite activity type.

The CGMY process has finite activity with finite variation for $Y < 0$, infinite activity with finite variation for $0 \leq Y < 1$, and infinite activity and variation for $1 \leq Y < 2$.

The characteristic function of the CGMY process is given by

$$\psi_{\text{CGMY}}(u) = \exp\left(tC\Gamma(-Y)\left[(M - iu)^Y - M^Y + (G + iu)^Y - G^Y\right]\right), \tag{11}$$

where Γ is the upper incomplete gamma function and is given by

$$\Gamma(-Y) = \int_0^\infty w^{-Y-1} e^{-t} dw. \tag{12}$$

The Lévy density is given by

$$k_{\text{CGMY}}(z) = C \frac{\exp(G|z|)}{|z|^{1+Y}} 1_{\{z<0\}} + C \frac{\exp(M|z|)}{|z|^{1+Y}} 1_{\{z>0\}}. \tag{13}$$

The mean correcting martingale parameter γ_{CGMY} for the CGMY process is

given by

$$\gamma_{\text{CGMY}} = -C\Gamma(-Y)\left((M-1)^Y - M^Y + (G+1)^Y - G^Y\right). \tag{14}$$

The parameters are restricted to $C > 0$, $G, M \geq 0$, and $Y < 2$; see, for instance, Carr et al. (2002) for further details. It is possible to allow both C and Y to have different values dependent on the sign of x. For $Y = 0$, we have the VG process [2].

Providing additional details on Lévy processes is out of the scope of this paper. For more details and the application of the processes to option pricing, we refer to Cont and Tankov (2005), Schoutens (2003), and references therein.

3   Finite Element Method (FEM) for Lévy Driven Assets

In this section, we provide a brief overview of FEM for option pricing on general Lévy models. We localise the PIDE, i.e. define it on a bounded domain. We then write it in variational form and, finally, define a discrete version.

3.1   Localization

To simplify the numeric, we define the PIDE as a forward parabolic problem by changing to time-to-maturity $t := T - t$. To obtain more stability, we remove the drift by defining

$$u(t,x) = e^{rt}V\left(T - t, x - (\gamma + r)t\right). \tag{15}$$

The boundary value problem is then given by

$$\frac{\partial u}{\partial t} - \mathcal{A}u = 0, \quad \text{in } (0,T) \times \mathbb{R}, \tag{16}$$

with an initial condition satisfying $u(0,x) = g(e^x)$. The jump-diffusion Merton (1976) model is an example of a model that contains both a Brownian motion and a jump component, and has an integro-differential operator that is fully described by Equation (7). However, as mentioned in Section 2, because we only focus on pure jump models, as the VG and CGMY models, we remove the diffusion operator in Equation (7) by setting $\sigma = 0$. The operator in Equation

---

[2]See the proof of Theorem 1 in Carr et al. (2002) for the relationship between ψ_{VG} and ψ_{CGMY} for $Y = 0$.

(16) then becomes

$$\mathscr{A} u(x) \;=\; \int_{\mathbb{R}} \left(u(x+z) - u(x) - z\frac{\partial}{\partial x} u(x) \right) d\nu(z). \tag{17}$$

The PIDE is discretised on a bounded rectangular domain $(t, x) \in [0, T] \times D$, where D is the state space of admissible logarithmic stock prices x. The space variable is truncated into a bounded domain $D := [-R, R]$ with boundary conditions on ∂D.

The bounded domain is chosen such that the risk-neutral probability for the process in Equation (3) to jump outside the domain is one basis point.

3.2   Variational Formulation

The variational formulation consists of multiplying the truncated PIDE by a smooth test function $v \in C_0^\infty(D)$ that satisfies $v(-R) = v(R) = 0$ and then applies integration by parts. The goal is then to find a continuous function u defined in $[0, T]$ with values in the Hilbert space H equipped with the norm $\|\cdot\|$ and inner product $(u, v) = \int_D u(x) v(x) \, dx$, such that

$$\frac{\partial}{\partial t}(u, v) + a(u, v) = 0, \quad \forall v \in H(\mathbb{R}), \tag{18}$$

where $u(0) = g(e^x)$. The bi-linear form $a(\cdot, \cdot) : H(\mathbb{R}) \times H(\mathbb{R}) \to \mathbb{R}$, associated with operator $\mathscr{A}$ in Equation (17), is given by

$$a(u, v) = \int_D \int_D \left(u(x+z) - u(x) - zu'(x) \right) v(x) \, d\nu(z) \, dx. \tag{19}$$

The main problem is the singularity of the Lévy measure at $z = 0$; but for the VG and CGMY model, it is possible to analytically solve the integro-differential operator $\mathscr{A}$. Integrate the jump generator by twice applying integration by parts,

$$\int_D \left(u(x+z) - u(x) - zu'(x) \right) k(z) \, dz \tag{20}$$

$$= \left(u(x+z) - u(x) - zu'(x) \right) k^{(-1)}(z) \big|_D - \int_D \left(u'(x+z) - u'(x) \right) k^{(-1)}(z) \, dz \tag{21}$$

$$= -\left(u'(x+z) - u'(x) \right) k^{(-2)}(z) \big|_D + \int_D u''(x+z) \, k^{(-2)}(z) \, dz \tag{22}$$

$$= \int_D u''(x+z) \, k^{(-2)}(z) \, dz, \tag{23}$$

where the first term in the second and third equality vanishes at the boundary and where the i-th antiderivative of $k = k^{(0)}$ is defined as,

$$k^{(-i)}(z) = \begin{cases} \int_{-\infty}^{z} k^{(-i+1)}(x)\, dx, & \text{if } z < 0, \\ -\int_{z}^{\infty} k^{(-i+1)}(x)\, dx, & \text{if } z > 0. \end{cases} \tag{24}$$

Applying integration by parts, we can then write the operator in Equation (19) as

$$a(u, v) = \int_{D} \int_{D} u'(y)\, v'(x)\, k^{(-2)}(y - x)\, dy\, dx. \tag{25}$$

3.3   Discretization

We define the discrete logarithmic stock mesh as $-R = x_0 < x_1 < \cdots < x_N < x_{N+1} = R$, with equidistant points $x_n = -R + n \cdot h$, for $n = 0, \ldots, N+1$ and width $h = 2R/(N+1)$. For discretisation, the Galerkin method is applied and the basic concept is that, for each $t \in [0, T]$, the solution $u(t, x)$ is approximated by an element $u_N(t, x) \in V_N = \text{span}\{b_i(x) : i = 1, \ldots, N\}$. For the basis $\{b_i\}_{i=1}^{N}$ of V_N, we chose the linear hat-functions b_i given by

$$b_i(x) = \max\{0, 1 - h^{-1}|x - x_i|\}, \quad i = 1, \ldots, N. \tag{26}$$

One advantage to using linear hat-functions as a basis is that solving integrals with inner products becomes much easier, as is later demonstrated in, for instance, Appendix D. Further details can be found in Braess (2007). We approximate the solution u_N by a linear combination of the basis function, that is,

$$u_N(t, x) = \sum_{j=1}^{N} u_{N,j}(t)\, b_j(x), \tag{27}$$

Equation (18) can then be written as

$$\frac{\partial}{\partial t}\left(\sum_{j=1}^{N} u_{N,j}(t)\, b_j(x), b_i(x)\right) + a\left(\sum_{j=1}^{N} u_{N,j}(t)\, b_j(x), b_i(x)\right) = 0, \quad \forall i = 1, \ldots, N. \tag{28}$$

Numerous ways exist to solve the ODEs that arise. However, to be able to simulate forward in time, we apply the implicit Euler scheme and define the time grid as $t_m = m \cdot k$, for $m = 0, \ldots, N_M$, and with step size $k = T/N_M$. With the linear hat-functions, the matrix formulation is given by finding $u_h^m \in \mathbb{R}^{N_m}$ such

that for $m = 0, \ldots, N_M - 1$,

$$\frac{1}{k}\mathbf{M}\left(u_h^{m+1} - u_h^m\right) + \mathbf{A}u_h^{m+1} = 0, \tag{29}$$

where $\mathbf{M}$ is the mass matrix and $\mathbf{A}$ is the stiffness matrix. Ern and Guermond (2004) show that Equation (29) converges to the true value as $k \to 0$ and $h \to 0$.

The entries $\mathbf{A}_{j,i}$ of the stiffness matrix is given by the following proposition,

Proposition 2 (*Stiffness matrix entries*) *The entries of the stiffness matrix $\mathbf{A}_{j,i}$ are, for $i = j$ given by*

$$\mathbf{A}_{i,i} = \frac{1}{h^2}\left(-4hk^{(-3)}(0) - 6k^{(-4)}(0) + 8k^{(-4)}(h) - 2k^{(-4)}(2h)\right). \tag{30}$$

For $i = j + 1$, given by

$$\mathbf{A}_{i,i+1} = \frac{1}{h^2}\left(2hk^{(-3)}(0) + 4k^{(-4)}(0) - 7k^{(-4)}(h) + 4k^{(-4)}(2h) - k^{(-4)}(3h)\right). \tag{31}$$

For $i \geq j + 2$, and $d = i - j$, given by

$$\mathbf{A}_{j,i} = \frac{1}{h^2}\sum_{l=-2}^{2} \alpha_j k^{(-4)}\left((d+l)h\right), \tag{32}$$

where $\alpha = (-1, 4, -6, 4, -1)$.

Proof. See Appendix D. ∎

One benefit of VG and CGMY is that the first four anti-derivatives can be calculated analytically. In order to keep the discussion simple we will only consider Lévy processes with symmetric density, that is, for $G = M$. The first four anti-derivative for the VG and CGMY model given by the following two propositions.

Proposition 3 (*Symmetric VG Anti-derivatives*) *For $z > 0$ and $Y = 0$ the anti-derivatives $k_{VG}^{(-i)}$, for $i = 1, 2, 3, 4$, are for the symmetric VG process, given by*

$$k_{VG}^{(-1)}(z) \quad = \quad -C Ei(\beta z), \tag{33}$$

$$k_{VG}^{(-2)}(z) \quad = \quad z k_{VG}^{(-1)}(z) + \frac{C}{\beta} e^{-\beta z}, \tag{34}$$

$$k_{VG}^{(-3)}(z) \quad = \quad z k_{VG}^{(-2)}(z) - \frac{1}{2} z^2 k_{VG}^{(-1)}(z) + \frac{1}{2} \frac{C}{\beta} \left(\frac{1}{\beta} - z \right) e^{-\beta z}, \tag{35}$$

$$k_{VG}^{(-4)}(z) \quad = \quad z k_{VG}^{(-3)}(z) - \frac{1}{2} z^2 k_{VG}^{(-2)}(z) + \frac{1}{6} z^3 k_{VG}^{(-1)}(z) \tag{36}$$

$$+ \frac{C}{6} \left(\frac{1}{\beta} z^2 - \frac{1}{\beta} z + \frac{1}{\beta} \right) e^{-\beta z},$$

where $\beta = G = M$ and $Ei(x) = \int_x^\infty e^{-t} t^{-1} dt$, is the exponential integral.

   Proof. See Appendix E. ∎

Proposition 4 (*Symmetric CGMY Anti-derivatives*) *For $z > 0$ and $0 < Y < 1$ the anti-derivatives $k_{CGMY}^{(-i)}$, for $i = 1, 2, 3, 4$, are for the symmetric CGMY process given by*

$$k_{CGMY}^{(-1)}(z) \quad = \quad -C L^Y \Gamma(Lz, -Y), \tag{37}$$

$$k_{CGMY}^{(-2)}(z) \quad = \quad z k_{CGMY}^{(-1)}(z) + C L^{Y-1} \Gamma(Lz, 1 - Y), \tag{38}$$

$$k_{CGMY}^{(-3)}(z) \quad = \quad z k_{CGMY}^{(-2)}(z) - \frac{1}{2} z_{CGMY}^2 k^{(-1)}(z) - \frac{1}{2} C L^{Y-2} \Gamma(Lz, 2 - Y), \tag{39}$$

$$k_{CGMY}^{(-4)}(z) \quad = \quad z k_{CGMY}^{(-3)}(z) - \frac{1}{2} z^2 k_{CGMY}^{(-2)}(z) + \frac{1}{6} z^3 k_{CGMY}^{(-1)}(z) \tag{40}$$

$$+ \frac{1}{6} C L^{Y-3} \Gamma(Lz, 3 - Y),$$

where $\Gamma(x, a) = \int_x^\infty t^{a-1} e^{-t} dt$, is the upper incomplete gamma function.

   Proof. See Appendix F. ∎

Remark 13 *For the piecewise linear finite elements, we obtain a convergence rate of $\mathcal{O}(h^2)$ in space. Increasing the space dimension, such as for multi-asset instruments, would require the multidimensional hat functions and the introduction of tensor products. This will have a significant impact on the convergence rate. However, using the random grids algorithm by Andreasen and Huge (2011), we can still obtain one-dimensional convergence (multiplied by the number of assets) through simulation using copulas.*

4   Simulation

In this section, we provide a brief summary of the Random Grids simulation algorithm in Andreasen and Huge (2011) for simulating underlying assets through its discrete PDE matrix formulation. We also present a method for estimating Greeks through Monte Carlo simulations.

From Equation (29), we have that

$$u(t_{h+1}) \quad = \quad (\mathbf{M} + k\mathbf{A})^{-1}\mathbf{M}u(t_h) \tag{41}$$

$$\equiv \quad \mathbf{A}_x^{-1}u(t_h), \tag{42}$$

which is the matrix formulation we use for the grid simulation. Each row of $\mathbf{A}_x^{-1}$ represents a transition probability distribution, and each element represents the conditional transition probability, that is,

$$P\left(x(t_{h+1}) = x_j \,|\, x(t_h) = x_i\right) = \left(\mathbf{A}_x^{-1}\right)_{ij}. \tag{43}$$

We can then define the cumulative distribution function as

$$Q_{ij} \quad = \quad P\left(x(t_{h+1}) \le x_j \,|\, x(t_h) = x_i\right) \tag{44}$$

$$= \quad \sum_{k \le j} \left(\mathbf{A}_x^{-1}\right)_{ik}. \tag{45}$$

For the simulation of the underlying asset, we rely on the discrete inverse transform method together with a simple table lookup. We define the discrete inverse transform method as

$$F_i^{-1}(u) \equiv \inf\{x : F_i(x) \ge u\}, \tag{46}$$

where F_i^{-1} denotes the inverse of the distribution function and where $Q_{ij} = F_i(x_j)$ for $i = 1, \dots, n$ and where $x(t_h) = x_i$.

Figure 1 illustrates the Laplace density $\mathbf{A}_x^{-1}$ and the cumulative distribution function for the CGMY process. Figure 2 illustrates a slice of the Laplace densities $\mathbf{A}_x^{-1}$ and the cumulative distribution function for the CGMY process for an arbitrarily grid point x_j.

Remark 14 *The discretisation of a pure diffusion problem, such as the Black-Scholes equation, is a $N \times N$ tri-diagonal sparse matrix with a matrix inversion requiring $\mathcal{O}(N)$ operations. The matrix formulation for PIDEs of option prices on by Lévy-driven assets is dense because of its non-localness. Calculating the inverse of a dense matrix requires $\mathcal{O}(N^2)$ operations, which are more computer intense compared with the tri-diagonal case. One could overcome this problem by*
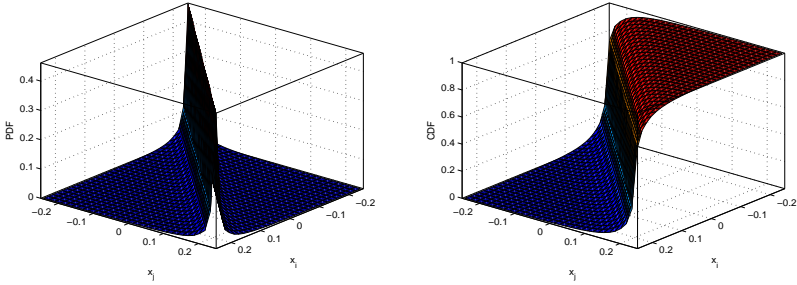
Figure 1: Left: The transition matrix $\mathbf{A}_x^{-1}$. Right: The cumulative distribution function.
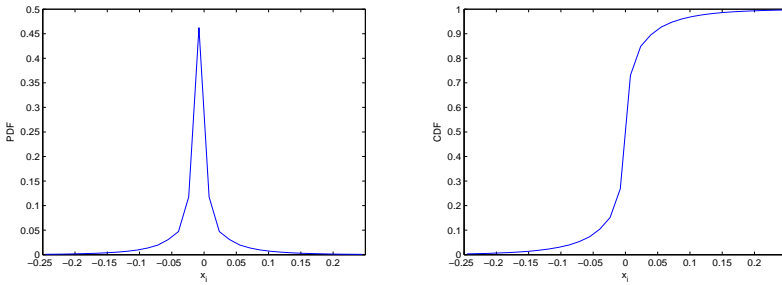
Figure 2: Left: Slice of the transition matrix $\mathbf{A}_x^{-1}$. Right: Slice of the cumulative distribution function.

considering a wavelet Galerkin scheme; see, for instance, Hilber, Kehtari, Schwab and Winter (2010). However, because we only consider a one-dimensional PIDE and only calculate one matrix inversion prior to the Monte Carlo simulation, this is not an issue.

4.1   Sensitivities

There are two different classes of Greeks. The first class is the solution sensitivity with respect to a bump of one of the input parameters, such as the Vega, which is the solution sensitivity with respect to the volatility parameter, $\partial_\sigma C$. The second class is a change in the price with respect to a bump in the input arguments, such as the delta $\partial_S C$. In FEM, the Greeks of the second class are easily obtained by post processing the solution. However, for Monte Carlo Greeks, one needs to

adjust the pricing simulation.

We focus on the second class of sensitivities: the deltas. To compute the delta using Monte Carlo simulation, one could for instance use a finite difference approximation, such as the bumping method; however, the method produces quite biased values and might sometimes be very unstable. The pathwise differentiation method, such as in Broadie and Glasserman (1996), is a better method that produces a direct and unbiased estimate of the true derivative because,

$$\frac{\partial}{\partial \theta} \mathbb{E}[g(\theta)] = \mathbb{E}\left[\frac{\partial}{\partial \theta} g(\theta)\right]. \tag{47}$$

For a European call option with maturity T the option value at time $t = 0$ is given by

$$V(0) = e^{-rT} \mathbb{E}[(S_T - K)^+]. \tag{48}$$

The delta is obtained by applying the chain rule along each path, that is,

$$\frac{dV(T)}{dS(T)} = e^{-rT} \mathbb{E}\left[\frac{S(T)}{S(0)} \mathbf{1}\{S(T) > K\}\right]. \tag{49}$$

The delta from Equation (49) is easily computed from a simulated path of the underlying by evaluating the indicator function $\mathbf{1}\{S(T) > K\}$. The indicator function takes the value of one for in-the-money options and zero otherwise.

4.2   The FEM-MC Simulation Algorithm

For clarity we summarise the steps of the simulation:

I. Initialization

1. Define time grid $t_0 = 0 < t_1 < \cdots < t_m = T$ and set $x_0 = \log(S_0)$.
2. Generate the mass and stiffness matrix $\mathbf{M}$ and $\mathbf{A}$

II. Simulation. For each path $k = 1, \ldots, N_k$.

1. **For each time point t_h $(1 \le h \le m)$,**
   (a) Update the transition matrix $\mathbf{A}_x^{-1} v(t_h)$ in Equation (41)
   (b) For $x(t_h) = x_j$, set $j = i$.
   (c) Draw a uniform random number $\tilde{u} \sim U(0, 1)$.

(d) Find the quantile:

    i. If $\tilde{u} \leq Q_{ii}$: While $\tilde{u} \leq Q_{i,j-1}$ set $j = j - 1$.

    ii. If If $\tilde{u} > Q_{ii}$: While $\tilde{u} > Q_{i,j-1}$ set $j = j + 1$.

(e) Set $x(t_{h+1}) = x_j$

2. Calculate the option payoff $g(x)$ and Greeks

5 Numerical Examples

In this section, we present the numerical results of European style call options with payoff at maturity T given by $V = (S_T - K)^+$, where K is the strike. The underlying S is given by Equation (3) and follows an exponential Lévy process of VG in Section 2.1 and CGMY type in Section 2.2. The corresponding European put options can be obtained by applying the put-call parity.

The VG and CGMY model parameters are arbitrarily chosen to $C = 5$, $G = M = 10$ and $Y = 0.5$ and the interest rate is assumed to be zero, $r = 0$. We choose to work with the moneyness $MN := S/K = [0.8, 1.0, 1.2]$ and maturities $T = [0.25, 0.50, 1, 10]$ years.

The simulation is performed using 20,000 scrambled (Sobol generated) quasi Monte Carlo (QMC) seeds on a FEM Galerkin PIDE discretisation spanned by the simple hat functions consisting of 100×100 (space×time) grid points, as described in Section 3. As common praxis, the truncation $[-R, R]$ is chosen such that for a given spot level S_0 at time $t = 0$, the likelihood of falling outside the truncation for some maturity T should be less than one basis point.

To remove the initial choice of the Monte Carlo seed in the overall comparison, we repeat the Monte Carlo simulation 100 times with different pseudo-random number seeds in each iteration. The reported mean and standard deviation values for the FEM-MC are the means and standard deviations of the 100 outer Monte Carlo simulations. For the simulation, the underlying initial spot price was chosen to be the value on the discretisation grid closest to the current MN value but also could be adjusted to any spot level using an interpolation technique.

To benchmark our values from FEM and a FEM Monte Carlo (FEM-MC) simulation, we use the COS method by Fang and Oosterlee (2008), which is based on the Fourier-cosine series expansion and can price options on the Lévy process at high precision.

In our first example, we simulate the prices of European call options and the corresponding implied Black-Scholes volatilities generated by the COS method, FEM, and FEM Monte Carlo (FEM-MC) simulation. The values are reported in Table 1 and Table 2, and for which the standard deviations are reported within parentheses. For short-dated options, the FEM-MC performs very well, and is

where jumps are needed to generate the extreme smile that one typically observes for such options, where Black-Scholes- and Heston-like models typically fails.

We observe that for all test cases, the simulated values generate a tight confidence interval that covers both FEM and COS prices. The 95% confidence interval generated by a FEM-MC simulation covers the COS and FEM values for all test cases, indicating that the simulating technique presented in this paper is suitable for simulating processes with jumps.

The small differences in prices between FEM and FEM-MC are the result of Monte Carlo noise and interpolation using table lock-up. Neglecting the interpolation error, these values will converge towards each other as the number of simulations N increases. The convergence rate of a pure Monte Carlo is on the order of $\mathcal{O}\left(N^{-1/2}\right)$. Numerical experiments with many types of integrands show that the convergence rate of the QMC method frequently leads to $\mathcal{O}\left(N^{-1}\right)$; for further details, see for instance Glasserman (2003). The small discrepancy between FEM and COS arise given different numerical techniques. Ignoring the discretisation errors, increasing the number of integration steps in the Fourier-cosine integration in the COS method (see Fang and Oosterlee (2008)), together with decreasing the step size in space and time, i.e. $k \to 0$ and $h \to 0$ for FEM, converges the two numerical techniques' values towards each other.

In our second example, we simulate the deltas for the same setup as in example one using the three different methods. The simulated deltas are reported in Table 3. Regarding the prices and implied volatilities, we observe accurate FEM-MC deltas with tight confidence intervals and where the 95% confidence interval generated by the FEM-MC simulation covers the COS and FEM values for all test cases.

In closing, we did not notice different results by using an alternative simulation grid, i.e. increasing or decreasing the (space×time) grid.

6 Conclusion

In this paper, we extended the simulation method by Andreasen and Huge (2011) and demonstrated how one could simulate options prices and sensitivities on Lévy processes, where each simulated path of the underlying is sampled from the matrix formulation of the numerical discretised PIDE. Reported Monte Carlo values demonstrate that the simulated option prices and deltas are consistent with the values coming from both FEM and Fourier integration.

TTM	MN	VG			CGMY		
		COS	**FEM**	**FEM-MC**	**COS**	**FEM**	**FEM-MC**
0.25	1.2	0,7758	0,7758	0,7767 (0,0194)	3,0046	3,0046	3,0097 (0,0401)
0.25	1.0	7,1709	7,1709	7,1726 (0,0337)	11,146	11,146	11,152 (0,0588)
0.25	0.8	21,125	21,125	21,124 (0,0325)	25,628	25,628	25,622 (0,0649)
0.5	1.2	1,8940	1,8940	1,9020 (0,0352)	6,1929	6,1929	6,2160 (0,0793)
0.5	1.0	8,9103	8,9103	8,9111 (0,0429)	17,344	17,344	17,345 (0,1045)
0.5	0.8	24,500	24,500	24,497 (0,0500)	30,092	30,092	30,077 (0,0952)
1.0	1.2	4,6064	4,6064	4,6146 (0,0643)	11,725	11,725	11,745 (0,1289)
1.0	1.0	13,810	13,810	13,811 (0,0723)	21,184	21,184	21,171 (0,1424)
1.0	0.8	28,165	28,165	28,162 (0,0778)	38,745	38,745	38,743 (0,1788)
10	1.2	27,086	27,086	27,168 (0,4327)	45,880	45,880	46,403 (1,5711)
10	1.0	38,407	38,407	38,412 (0,5709)	65,724	65,724	66,309 (2,3308)
10	0.8	55,837	55,837	55,868 (0,6393)	84,684	84,684	85,297 (2,4876)

Table 1: European call option prices on VG and CGMY with $C = 5$, $G = M = 10$ and $Y = 0.5$. The Monte Carlo standard errors are reported within the parentheses.

TTM	MN	VG				CGMY			
		COS	FEM	FEM-MC		COS	FEM	FEM-MC	
0.25	1.2	0,3451	0,3451	0,3452	(0,0024)	0,5312	0,5312	0,5316	(0,0030)
0.25	1.0	0,2875	0,2875	0,2876	(0,0017)	0,5193	0,5193	0,5196	(0,0030)
0.25	0.8	0,3190	0,3190	0,3188	(0,0029)	0,5217	0,5217	0,5214	(0,0040)
0.5	1.2	0,3237	0,3237	0,3242	(0,0022)	0,5292	0,5292	0,5303	(0,0037)
0.5	1.0	0,3021	0,3021	0,3021	(0,0015)	0,5247	0,5247	0,5247	(0,0037)
0.5	0.8	0,3101	0,3101	0,3099	(0,0024)	0,5245	0,5245	0,5239	(0,0036)
1.0	1.2	0,3164	0,3164	0,3166	(0,0022)	0,5294	0,5294	0,5300	(0,0039)
1.0	1.0	0,3095	0,3095	0,3095	(0,0018)	0,5279	0,5279	0,5276	(0,0037)
1.0	0.8	0,3102	0,3102	0,3101	(0,0022)	0,5271	0,5271	0,5271	(0,0046)
10	1.2	0,3172	0,3172	0,3180	(0,0044)	0,5307	0,5307	0,5375	(0,0200)
10	1.0	0,3168	0,3168	0,3169	(0,0051)	0,5305	0,5305	0,5373	(0,0257)
10	0.8	0,3165	0,3165	0,3168	(0,0053)	0,5304	0,5304	0,5370	(0,0253)

Table 2: Black Scholes implied volatility from European call options on VG and CGMY with $C = 5$, $G = M = 10$ and $Y = 0.5$. The Monte Carlo standard errors are reported within the parentheses.

		VG				CGMY			
TTM	**MN**	**COS**	**FEM**	**FEM-MC**		**COS**	**FEM**	**FEM-MC**	
0.25	1.2	0,0910	0,0913	0,0911	(0,0015)	0,2471	0,2465	0,2464	(0,0022)
0.25	1.0	0,6028	0,6022	0,6046	(0,0022)	0,5699	0,5705	0,5698	(0,0021)
0.25	0.8	0,9147	0,9137	0,9153	(0,0011)	0,8192	0,8193	0,8201	(0,0015)
0.5	1.2	0,1764	0,1762	0,1763	(0,0021)	0,3535	0,3529	0,3531	(0,0026)
0.5	1.0	0,5475	0,5481	0,5472	(0,0024)	0,6164	0,6171	0,6163	(0,0020)
0.5	0.8	0,8623	0,8618	0,8629	(0,0012)	0,7704	0,7710	0,7711	(0,0016)
1.0	1.2	0,3036	0,3031	0,3028	(0,0025)	0,4610	0,4603	0,4604	(0,0029)
1.0	1.0	0,5894	0,5901	0,5891	(0,0021)	0,6063	0,6066	0,6058	(0,0025)
1.0	0.8	0,8002	0,8005	0,8010	(0,0017)	0,7589	0,7597	0,7590	(0,0020)
10	1.2	0,6209	0,6171	0,6204	(0,0057)	0,7642	0,7501	0,7699	(0,0191)
10	1.0	0,6903	0,6871	0,6894	(0,0057)	0,8103	0,7968	0,8152	(0,0214)
10	0.8	0,7616	0,7591	0,7617	(0,0049)	0,8400	0,8270	0,8444	(0,0189)

Table 3: European call option deltas on VG and CGMY with $C = 5$, $G = M = 10$ and $Y = 0.5$. The Monte Carlo standard errors are reported within the parentheses.

D Finite Element for Lévy Models

With basis function given by Equation (26) the stiffness matrix (25) is then computed by

$$
\begin{aligned}
a\left(b_i, b_j\right) &= \int_{x_{i-1}}^{x_{i+1}} \int_{x_{j-1}}^{x_{j+1}} b_i^{'}(x)\, b_j^{'}(y)\, k^{(-2)}(y-x)\, dx\, dy \\
&= \frac{1}{h^2}\left(\int_0^h \int_{(j-i)h}^{(j-i+1)h} k^{(-2)}(y-x)\, dx\, dy - \int_0^h \int_{(j-i+1)h}^{(j-i+2)h} k^{(-2)}(y-x)\, dx\, dy \right. \\
&\quad \left. - \int_0^h \int_{(j-i-1)h}^{(j-i)h} k^{(-2)}(y-x)\, dx\, dy + \int_0^h \int_{(j-i)h}^{(j-i+1)h} k^{(-2)}(y-x)\, dx\, dy \right) \\
&= \frac{1}{h^2}\left(2\int_0^h \int_{(j-i)h}^{(j-i+1)h} k^{(-2)}(y-x)\, dx\, dy - \int_0^h \int_{(j-i+1)h}^{(j-i+2)h} k^{(-2)}(y-x)\, dx\, dy \right. \\
&\quad \left. - \int_0^h \int_{(j-i-1)h}^{(j-i)h} k^{(-2)}(y-x)\, dx\, dy \right).
\end{aligned}
$$

For $i = j$

$$
\begin{aligned}
\mathbf{A}_{i,i} &= \frac{1}{h^2}\left(2\int_0^h \int_0^h k^{(-2)}(y-x)\, dx\, dy - \int_0^h \int_h^{2h} k^{(-2)}(y-x)\, dx\, dy \right. \\
&\quad \left. - \int_0^h \int_{-h}^0 k^{(-2)}(y-x)\, dx\, dy \right) \\
&= \frac{2}{h^2}\left(\int_0^h \int_0^h k^{(-2)}(y-x)\, dx\, dy - \int_0^h \int_h^{2h} k^{(-2)}(y-x)\, dy\, dx \right).
\end{aligned}
$$

Where

$$
\begin{aligned}
\int_0^h \int_0^h k^{(-2)}(y-x)\, dy\, dx &= \int_0^h \left(\int_0^x k^{(-2)}(y-x)\, dy + \int_x^h k^{(-2)}(y-x)\, dy \right) dx \\
&= \int_0^h \left(k^{(-3)}(0) + k^{(-3)}(x) + k^{(-3)}(h-x) - k^{(-3)}(0) \right) dx \\
&= -2h k^{(-3)}(0) + 2k^{(-4)}(h) - 2k^{(-4)}(0).
\end{aligned}
$$

and

$$
\begin{aligned}
\int_0^h \int_h^{2h} k^{(-2)}(y-x)\, dy\, dx &= \int_0^h \left(k^{(-2)}(2h-x) - k^{(-2)}(h-x) \right) dx \\
&= 2k^{(-2)}(h) - k^{(-2)}(0) - k^{(-2)}(2h).
\end{aligned}
$$

The entries for $i = j$ is then given by,

$$\mathbf{A}_{i,i} = \frac{1}{h^2} \left(-4hk^{(-3)}(0) - 6k^{(-4)}(0) + 8k^{(-4)}(h) - 2k^{(-4)}(2h) \right).$$

For $i = j + 1$

$$\mathbf{A}_{i,i+1} = \frac{1}{h^2} \left(2 \int_0^h \int_h^{2h} k^{(-2)}(y-x)\,dxdy - \int_0^h \int_0^h k^{(-2)}(y-x)\,dxdy \right.$$
$$\left. - \int_0^h \int_{2h}^{3h} k^{(-2)}(y-x)\,dxdy \right).$$

and where,

$$\int_0^h \int_{2h}^{3h} k^{(-2)}(y-x)\,dxdy = \int_0^h \left(k^{(-3)}(3h-x) - k^{(-3)}(2h-x) \right) dy$$
$$= 2k^{(-4)}(2h) - k^{(-4)}(h) - k^{(-4)}(3h).$$

The entries for $i = j + 1$ is then given by,

$$\mathbf{A}_{i,i+1} = \frac{1}{h^2} \left(2hk^{(-3)}(0) + 4k^{(-4)}(0) - 7k^{(-4)}(h) + 4k^{(-4)}(2h) - k^{(-4)}(3h) \right).$$

For $d = i - j$, $i \geq j + 2$

$$\mathbf{A}_{j,i} = h^{-2} \left(2 \int_0^h \int_{dh}^{(d+1)h} k^{(-2)}(y-x)\,dxdy - \int_0^h \int_{(d+1)h}^{(d+2)h} k^{(-2)}(y-x)\,dxdy \right.$$
$$\left. - \int_0^h \int_{(d-1)h}^{dh} k^{(-2)}(y-x)\,dxdy \right).$$

Consider the case when $d = 2$

$$\mathbf{A}_{j,i} = h^{-2} \left(2 \int_0^h \int_{2h}^{3h} k^{(-2)}(y-x)\,dxdy - \int_0^h \int_{3h}^{4h} k^{(-2)}(y-x)\,dxdy \right.$$
$$\left. - \int_0^h \int_h^{2h} k^{(-2)}(y-x)\,dxdy \right),$$

where

$$\int_0^h \int_{3h}^{4h} k^{(-2)}(y-x)\,dxdy = \int_0^h \left(k^{(-3)}(4h-x) - k^{(-3)}(3h-x) \right) dx$$
$$= k^{(-4)}(4h-h) - k^{(-4)}(3h-h) - \left(k^{(-4)}(4h) - k^{(-4)}(3h) \right).$$

The entries for $d = i - j$, $i \geq j + 2$ is then given by,

$$\mathbf{A}_{i,i+1} = \frac{1}{h^2} \left(-k^{(-4)}(0) + 4k^{(-4)}(h) - 6k^{(-4)}(2h) + 4k^{(-4)}(3h) - k^{(-4)}(4h) \right),$$

the pattern appears and hold for all $d = i - j$, $i \geq j + 2$.

E   VG Anti-derivatives

Let $\mathrm{Ei}(x) = \int_x^\infty e^{-t} t^{-1} dt$ be the exponential integral, $k^{(-i)}$ the i-th anti-derivative, and the VG Lévy measure k_{VG} by Equation (9).

For $i = 1$,

$$\begin{aligned}
k_{\mathrm{VG}}^{(-1)}(z) &= -\int_z^\infty k_{\mathrm{VG}}(x)\,dx \\
&= -\int_z^\infty \frac{Ce^{-\beta|x|}}{|x|}\,dx \\
&= -C\mathrm{Ei}(\beta z).
\end{aligned}$$

For $i = 2$,

$$\begin{aligned}
k_{\mathrm{VG}}^{(-2)}(z) &= -\int_z^\infty k_{\mathrm{VG}}^{(-1)}(x)\,dx \\
&= zk_{\mathrm{VG}}^{(-1)}(z) + C\int_z^\infty x\frac{e^{-\beta x}}{x}\,dx \\
&= zk_{\mathrm{VG}}^{(-1)}(z) + \frac{C}{\beta}e^{-\beta z},
\end{aligned}$$

For $i = 3$,

$$\begin{aligned}
k_{\mathrm{VG}}^{(-3)}(z) &= -\int_z^\infty k_{\mathrm{VG}}^{(-2)}(x)\,dx \\
&= zk_{\mathrm{VG}}^{(-2)}(z) - \frac{1}{2}z^2 k_{\mathrm{VG}}^{(-1)}(z) - \frac{C}{2}\int_z^\infty xe^{-\beta x}\,dx \\
&= zk_{\mathrm{VG}}^{(-2)}(z) - \frac{1}{2}z^2 k_{\mathrm{VG}}^{(-1)}(z) + \frac{1}{2}\frac{C}{\beta}\left(\frac{1}{\beta} - z\right)e^{-\beta z},
\end{aligned}$$

where the second equality is obtained by applying integration by parts twice.

For $i = 4$,

$$
\begin{aligned}
k_{\mathrm{VG}}^{(-4)}(z) &= -\int_z^\infty k_{\mathrm{VG}}^{(-3)}(x)\,dx \\
&= z k_{\mathrm{VG}}^{(-3)}(z) - \frac{1}{2}z^2 k_{\mathrm{VG}}^{(-2)}(z) + \frac{1}{6}z^3 k_{\mathrm{VG}}^{(-1)}(z) + \frac{C}{6}\int_z^\infty x^2 e^{-\beta x}\,dx,
\end{aligned}
$$

where the second equality is obtained by applying integration by parts three times and where

$$
\int_z^\infty x^2 e^{-\beta x}\,dx = e^{-\beta z}\left(\frac{1}{\beta}z^2 - \frac{1}{\beta}z + \frac{1}{\beta}\right).
$$

F   CGMY Anti-derivatives

Let $\Gamma(x,a) = \int_x^\infty t^{a-1}e^{-t}\,dt$ be the upper incomplete gamma function, the i-th anti-derivative $k^{(-i)}$ be given by Equation (24) and the CGMY Lévy measure by Equation (13).

For $i = 1$,

$$
\begin{aligned}
k_{\mathrm{CGMY}}^{(-1)}(z) &= -C\int_z^\infty \frac{e^{-L|x|}}{|x|^{1+Y}}\,dx \\
&= -CL^Y\Gamma(Lz,-Y).
\end{aligned}
$$

For $i = 2$

$$
\begin{aligned}
k_{\mathrm{CGMY}}^{(-2)}(z) &= -\int_z^\infty k_{\mathrm{CGMY}}^{(-1)}(x)\,dx \\
&= -z k_{\mathrm{CGMY}}^{(-1)}(z) + c\int_{Lz}^\infty \frac{e^{-t}}{t^Y}L^{Y-1}\,dt \\
&= C\left(-z k_{\mathrm{CGMY}}^{(-1)}(z) + L^{Y-1}\Gamma(Lz,1-Y)\right).
\end{aligned}
$$

For $i = 3$

$$
\begin{aligned}
k_{\mathrm{CGMY}}^{(-3)}(z) &= -\int_z^\infty k_{\mathrm{CGMY}}^{(-2)}(x)\,dx \\
&= z k_{\mathrm{CGMY}}^{(-2)}(z) - \frac{1}{2}z^2 k_{\mathrm{CGMY}}^{(-1)}(z) - \frac{1}{2}\int_z^\infty x^2 k_{\mathrm{CGMY}}(x)\,dx,
\end{aligned}
$$

where the second equality is obtained from applying integration by parts twice

and where

$$\int_z^\infty x^2 k_{\text{CGMY}}(x)\,dx = \int_z^\infty x^2 C \frac{e^{-Lx}}{x^{Y+1}}\,dx$$

$$= C \int_{Lz}^\infty e^{-t} t^{1-Y} L^{Y-2}\,dt$$

$$= C L^{Y-2} \Gamma(Lz, 2-Y).$$

For $i = 4$, and applying integration by parts three times

$$k_{\text{CGMY}}^{(-4)}(z) = -\int_z^\infty k_{\text{CGMY}}^{(-3)}(x)\,dx$$

$$= z k_{\text{CGMY}}^{(-3)}(z) - \frac{1}{2} z^2 k_{\text{CGMY}}^{(-2)}(z) + \frac{1}{6} z^3 k_{\text{CGMY}}^{(-1)}(z) + \frac{1}{6} \int_z^\infty x^3 k_{\text{CGMY}}(x)\,dx,$$

and where

$$\int_z^\infty x^3 k_{\text{CGMY}}(x)\,dx = \int_z^\infty x^3 C \frac{e^{-Lx}}{x^{Y+1}}\,dx$$

$$= C L^{Y-3} \int_{Lz}^\infty \frac{e^{-t}}{t^{Y-2}}\,dt$$

$$= C L^{Y-3} \Gamma(Lz, 3-Y).$$

EPILOGUE

The Future of Quantitative Finance

Abstract

In the Golden Age of Quants, pricing exotic derivatives was the hottest topic. However, when the financial crisis occurred and the aftermath emerged, everything changed. Quant activities largely seem to have had to take a back seat given the drying up of exotics. However, as valuation adjustments (XVA) had to be priced in, vanilla derivatives became the new exotics and new challenges emerged.

Today, quants have opportunities to face greater challenges outside the classic area of pricing. Hedging is more complex and needs to be tackled from different angles, as XVA needs to be incorporated. To allow for efficient XVA hedging, a more liquid credit default swaps market is required. Moreover, with fewer exotics and more standardised products, the flow desks will continue to grow. Together with increasing derivatives, clearing will require liquidity optimisation.

Collateralisation has recently been priced in, and without standardisation it is open for enhancement. Smart derivatives with block-chain and automated systems for peer-to-peer collateral payments might be the game changer. Relying on technology behind cryptocurrencies, such as Bitcoin, has the potential to retire the XVAs for good.

Today, the demand for quants possessing high levels of IT skills is larger. Quants have moved from analytics to more development roles, where skills to develop production libraries are essential. As trading has become more automated, quants have already explored algorithmic trading, replacing old-school traders and moving from the sell-side to the buy-side and Fintech companies.

We have gone from a less model-driven period to a more (big) data driven one. Moreover, with big data, new challenges emerge and one needs to be able to use and process these data quickly and take advantage of having access to large data sets. We need to explore artificial intelligence (AI) from a broad business scope to make sense of the data to improve trading and to develop new business opportunities. Banks need to bring these advantages to their clients by implementing Facebook-like services, such as being able to tailor-make investments on the basis of client behaviour. Moreover, with the Internet of Things (IoT), we can take it further as the removal of humans and better real-time big data techniques evolve. Allowing for fast data access and being able to analyse large sets of data in real-time will further increase the pressure on the IT infrastructure. The industry spent the last years exploring GPUs and FPGAs to accelerate trading activities. The next big thing – quantum computing – still in its infant stage, will take computing power to a new level; in the 2030s, it is expected to be as common as CPU-GPU computing today.

However, as Moore's law is diminishing, we also have to focus on writing smarter algorithms rather than adding new CPU-GPUs each time performance issues arise. Recently, Google's AlphaGo became the first algorithm to beat a professional human Go player.[1] Instead of relying on brute force methods, it relies on neural networks inspired by the human brain to develop intuition and strategies rather than pattern recognition. However, as the area of AI increases, one needs to be aware of its potential operational risk. It only takes one mistake to wipe out an entire business; therefore, sophisticated risk-management

---

[1] http://www.nature.com/news/google-ai-algorithm-masters-ancient-game-of-go-1.19234

systems are essential.

As the definition of the traditional quant is emerging, pricing will still be involved. Over the last years, quants have explored and are still exploring innovative modelling, such as methodologies to efficiently estimate XVA. Something that previously worked well on a desk level does not certainly hold on an aggregated level. Moreover, given all of the new regulations waiting around the corner, the story is different. We need to have a generalised pricing approach that can be applied uniformly across asset classes and that can adapt as regulations and policies develop.

Although the *Golden Age of Quants* during which quants had exotic trading desk as their clients is over, quants have moved to a situation in which everybody needs them. More than ever, quants need to be aware of new regulations and policies and how they affect daily business. The modern quant needs to be predictive, understand the big picture, and seek business opportunities, and will have more client-facing roles in the future. Instead of being notorious rocket scientists, quants face greater demand to possess a broader range of skills that cover mathematics, finance, computer science, and business administration. Therefore, quants will evolve into universal soldiers and become more valuable as we enter the *New Golden Age of Quants*.

BIBLIOGRAPHY

Achdou, Y. and Pironneau, O (2005). *Computational Methods for Option Pricing*, SIAM Frontiers in Applied Mathematics.

Andersen, L. (2000). A Simple Approach to the Pricing of Bermudan Swaptions in the Multifactor LIBOR Market Model. *The Journal of Computational Finance* 3: 5–32.

Andersen, L. and Broadie, M. (2004). A Primal-Dual Simulation Algorithm for Pricing Multi-Dimensional American Options, *Management Science* 50: 1222-1234.

Andersen, L. and Brotherton-Ratcliffe, R. (2005). Extended Libor Market Models with Stochastic Volatility. *Journal of Computational Finance* 9: 1–40.

Andersen, L and Piterbarg, V. (2010). *Interest Rate Modeling: Models, Products, Risk Management*. In three volumes. Atlantic Financial Press.

Anderson, G., Goldberg, L., Kercheval, A.N., Miller, G and Sorge, K. (2005). On the Aggregation of Local Risk Models for Global Risk Management. *Journal of Risk* 8: 25–40.

Andreasen, J. and Huge, B. (2011). Random Grids, *Risk Magazine* 3: 66–71.

Asmussen, S. and Rosinski, J. (2001). Approximations of small jumps of Lévy processes with a view towards simulation, *Journal of Applied Probability* 38: 482–493.

Bachelier, L. (1990). *Theorie de la Speculation*, Ph.D. Dissertation, LEcole Normale Superieure.

Belomestny, D., Bender, C., and Schoenmakers, J. (2009). True Upper Bounds for Bermudan Products via Non-Nested Monte Carlo, *Mathematical Finance* 19: 53–71.

Belomestny, D., Kolodko, A. and Schoenmakers, J. (2009). Pricing CMS spreads in the LIBOR market model. *International Journal of Theoretical and Applied Finance* 13: 45-62.

Bender, C., and Schoenmakers, J. (2006). An Iterative Method for Multiple Stopping: Convergence and Stability. *Advances in Applied Probability* 38: 729–749.

Black, F. (1976). The Pricing of Commodity Contracts. *Journal of Financial Economics* 3: 167–179.

Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81: 637–654.

Brace, A., Gatarek, D. and Musiela, M. (1997). The Market Model of Interest Rate Dynamics. *Mathematical Finance* 7: 127–154.

Braess, D. (2007). *Finite elements*, 3rd Ed, Cambridge University Press, Cambridge.

Brennan, M.J. and Schwartz, E.S. (1978). Finite difference methods and jump processes arising in the pricing of contingents claims: a synthesis, *Journal of Financial Quantitative Analysis* 13: 462–474.

Brigo, D. and Mercurio, F. (2006). *Interest Rate Models: Theory and Practice - with Smile, Inflation and Credit*, 2nd Ed, Heidelberg: Springer Verlag.

Brigo, D., Morini M. and Pallavicini, A. (2013). *Counterparty Credit Risk, Collateral and Funding: With Pricing Cases For All Asset Classes*, Chichester: Wiley.

Broadie, M. and Cao, M. (2008). Improved lower and upper bound algorithms for pricing American options by simulation. *Quantitative Finance* 8: 845-861.

Broadie, M. and Glasserman, P (1996). Estimating security price derivatives using simulation, *Management Science* 52: 269–285.

Broadie, M. and Glasserman, P. (2004). A Stochastic Mesh Method for Pricing High-Dimensional American Options. *The Journal of Computational Finance* 7: 35–72.

Canabarro, E. and Duffie, D. (2003). Measuring and marking counterparty risk. Chap. 9 in *Asset/Liability Management of Financial Institutions*, edited by Tilman, L., 122-134. New York: Institutional Investor Books.

Carr, P., Chang, E. and Madan, D.B (1998). The variance gamma process and option pricing, *European Finance Review* 2: 79–105.

Carr, P., Geman, H., Madan, D.B. and Yor, M (2002). The fine strucutre of asset returns: An empirical investigation, *Journal of Business* 75: 305–332.

Carr, P. and Madan, D.B (1999). Option Valuation Using the Fast Fourier Transform, *Journal of Computational Finance* 2: 61–73.

Carriere. J. (1996). Valuation of Early-Exercise Price of Options Using Simulations and Nonparamnetric Regression, *Insurance Mathematics and Economics* 19: 19-30.

Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016a). Pricing of Long–Dated Commodity Derivatives with Stochastic Volatility and Stochastic Interest Rates. Available at SSRN: `http://ssrn.com/abstract=2712025`.

Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016b). Empirical Pricing Performance on Long–Dated Crude Oil Derivatives: Do Models with Stochastic Interest Rates Matter?. Available at SSRN: `http://ssrn.com/abstract=2721716`.

Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016c). 'Hedging Futures Options with Stochastic Interest Rates', Available at SSRN: `http://ssrn.com/abstract=2840635`.

Cheng, B. Nikitopoulos Sklibosios, C. and Schlögl, E. (2016d). 'Empirical Hedging Performance on Long-Dated Crude Oil Derivatives', Available at SSRN: `http://ssrn.com/abstract=2840622`.

Choy, B., Dun, T. and Schlögl, E. (2004). Correlating Market Models. *Risk* 17: 124–129.

Clement, E., Lamberton, D. and Protter, P. (2002). An analysis of a least squares regression method for American option pricing. *Finance and Stochastics* 6: 449-471.

Cont, R. and Tankov, P (2005). *Financial Modelling with Jump Processes*, Chapman & Hall, CRC Press.

Cortazar, G. and Schwartz, E. S. (1994). The Valuation of Commodity Contingent Claims. *Journal of Derivatives* 1: 27–39.

Cox, J.C., Ingersoll, J.E. and Ross, S.A. (1981). The Relation between Forward Prices and Futures Prices. *Journal of Financial Econometrics* 9: 321–346.

Cox, J., Ingersoll, J. and Ross, S. (1985). A theory of the term structure of interest rates, *Econometrica* 53, 385-407.

De Prisco, B. and Rosen, D. (2005). *Modeling Stochastic Counterparty Credit Exposures for Derivatives Portfolios, In Counterparty Credit Risk Modeling*, edited by M. Pykhtin, Risk Books, London.

Derman, E. (2003). Laughter in the Dark — The Problem of the Volatility Smile. May 26, 2003.

Dupire, B. (1994). Pricing with a Smile, *Risk* 7: 18–20.

Ern E. and Guermond, J.L. (2004). Theory and practice of finite elements, volume 159 of Applied Mathematical Sciences. Springer, New York.

Fang, F. and Oosterlee, C.W (2008). A novel pricing method for European options based on Fourier-cosine series expansions, *SIAM Journal of Scientific Computing* 31: 826–848.

Gibson, R. and Schwartz, E. S. (1990). Stochastic Convenience Yield and the Pricing of Oil Contingent Claims. *Journal of Finance* 45: 959–976.

Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*, Springer.

Glasserman, P and Yu, B. (2004). Number of paths versus number of basis functions in American option pricing, *Annals of Applied Probability* 14: 2090-119.

Golub, G. and Van Loan, C. (1996). *Matrix Computations*, 3rd Ed., Johns Hopkins.

Gower, J.C. and Dijksterhuis, G.B. (2004). *Procrustes Problems*, Oxford University Press.

Gregory, J. (2010). *Counterparty Credit Risk and Credit Value Adjustment: A Continuing Challenge for Global Financial Markets.* Chichester: Wiley.

Grzelak, L. and Oosterlee, C. (2011a). On Cross-Currency Models with Stochastic Volatility and Correlated Interest Rates. *Applied Mathematical Finance* 19: 1–35.

Grzelak, L. and Oosterlee, C. (2011b). 'On the Heston Model with Stochastic Interest Rates', *SIAM J. Financial Math.* 2: 255–286.

Hagan, P., Kumar, D., Lesniewski, A. and D. Woodward. (2002). Managing Smile Risk. *Wilmott Magazine*, September, 84–108.

Haugh, M. and Kogan, L. (2004). Pricing American options: a duality approach, *Operations Research* 52: 258–270.

Heath, D., Jarrow, R. and A. Morton. (1992). Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation. *Econometrica* 60: 77–105.

Heston, S. (1993). A closed-form solution for options with stochastic volatility with application to bond and currency options, *Review of Financial Studies* 6: 327–343.

Hilber, N., Kehtari, S., Schwab, C. and Winter, C (2010). Wavelet finite element method for option pricing highdimensional diffusion market models, SAM Report, 2010–01.

Hull, J. and White, A. (1987). The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 271-301.

Hull, J. and White, A. (1990). Pricing interest-rate derivative securities, *The Review of Financial Studies* 3: 573–592.

Hull, J. and White, A. (2012). CVA and wrong way risk, *Financial Analysts Journal* 68: 58–69.

Jain, S. and Oosterlee, C.W. (2012). Pricing high-dimensional Bermudan options using stochastic grid method. *International Journal of Computer Mathematics* 89: 1186-1211.

Jain, S. and Oosterlee, C.W. (2015). The stochastic grid bundling method: efficient pricing of Bermudan options and their Greeks. *Applied Mathematics and Computation* 269: 412–431.

Jamshidian, F. (1997). LIBOR and Swap Market Models and Measures. *Finance and Stochastics* 1: 293–330.

Jarrow, R., (2010). Option Pricing Theory: Historical Perspectives, in Encyclopedia of Quantitative Finance, eds Rama Cont, John Wiley & Sons.

Jarrow, R. and Turnbull, S. (1992). Credit risk: drawing the analogy, *Risk Magazine* 5.

Jarrow, R. and Turnbull, S. (1995). Pricing derivatives on financial securities subject to credit risk, *Journal of Finance* 50: 53-85.

Jin, X., Tana, H.H. and Sun, J. (2007). A State-Space Partitioning Method For Pricing High-Dimensional American-Style Options, *Mathematical Finance* 17: 399–426.

Joshi, M. and R. Rebonato (2003). 'A Displaced–Diffusion Stochastic Volatility LIBOR Market Model: Motivation, Definition and Implementation', *Quantitative Finance* 3: 458–469.

Kercheval, A.N. (2006). Optimal Covariances in Risk Model Aggregation. *Proceedings of the Third IASTED International Conference on Financial Engineering and Applications*, ACTA Press, 30-35.

Kolodko, A., and Schoenmakers, J. (2004). Upper bounds for Bermudan style derivatives, *Monte Carlo Methods and Applications* 10: 331–343.

Kolodko, A., and Schoenmakers, J. (2006). Iterative construction of the optimal Bermudan stopping time, *Finance and Stochastics* 10: 27–49.

Koschat, M.A. andD.F. Swayne (1991). A weighted Procrustes criterion. *Psychometrika* 56: 229-223.

Lando, D. (1998). On Cox processes and credit risky securities, *Review of Derivatives Research* 2: 99-120.

Longstaff, F.A. and Schwartz, E.S. (2001). Valuing American Options by Simulation: A Simple Least-Squares Approach, *Review of Financial Studies* 14: 113–147.

Matache, A. M., von Petersdorff, T. and Schwab, C. (2004). Fast deterministic pricing of options on Lévy driven assets, *M2AN Math. Model. Numer. Analysis* 38: 37–71.

Merton, R.C. (1973). The theory of rational option pricing, *Bell Journal of Economics and Management Science* 4: 141-183.

Merton, R.C. (1974). On the pricing of corporate debt: the risk structure of interest rates, *Journal of Finance* 29: 449-470.

Merton, R.C. (1976). Option Pricing when underlying returns are discontinuous, *Journal of Financial Economics* 3: 125–144.

Milstein, G.N. and Tretyakov, M. V. (2009). Practical variance reduction via regression for simulating diffusions. *SIAM Journal of Numerical Analysis* 47: 887-910.

Miltersen, K. (2003). Commodity Price Modelling that Matches Current Observables: A New Approach. *Quantitative Finance* 3: 51–58.

Miltersen, K., Sandmann, K. and Sondermann, D. (1997). Closed Form Solutions for Term Structure Derivatives with Log-Normal Interest Rates. *Journal of Finance* 52: 409–430.

Miltersen, K. and Schwartz, E. (1998). Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yield and Interest Rates. *Journal of Financial and Quantitative Analysis* 33: 33–59.

Musiela, M. and Rutkowski, M. (1997a). 'Continuous–time Term Structure Models: Forward Measure Approach', *Finance and Stochastics* 1: 261–291.

Musiela, M. and Rutkowski, M. (1997b). *Martingale Methods in Financial Modelling*, Springer Verlag.

Nocedal, J. and Wright, S.J. (2006). *Numerical Optimization*, Springer Verlag.

Pedersen, M.B. (1998). Calibrating Libor Market Models. *SimCorp Financial Research Working Paper*.

Picoult, E. (2005). *Calculating and Hedging Exposure, Credit Value Adjustment and Economic Capital for Counterparty Credit Risk in Counterparty Credit Risk Modelling*. Risk Books

Pilz, K.F. and Schlögl, E. (2013). A hybrid commodity and interest rate market model. *Quantitative Finance* 13: 543–560.

Piterbarg V. (2004). Pricing and hedging callable Libor exotics in forward Libor models, *The Journal of Computational Finance* 8: 65–117.

Piterbarg, V. (2005a). Time to smile. *Risk Magazine* 18: 71–75.

Piterbarg, V. (2005b). Stochastic volatility model with time dependent skew. *Applied Mathematical Finance* 12: 147–185.

Pykhtin, M. and Rosen, D. (2010). Pricing Counterparty Risk At The Trade Level and Credit Valuation Adjustment Allocations, *The Journal of Credit Risk* 6: 3–38.

Pykhtin, M. and Zhu, S. (2007). A Guide to Modeling Counterparty Credit Risk, *GARP Risk Review* 37: 16–22.

Rebonato, R. (1999). 'Calibrating the BGM Model', *Risk*, March, 88–94.

Rebonato, R. (2004). *Volatility and Correlation: The Perfect Hedger and the Fox*, 2nd edition, Wiley, Chichester.

Rebonato, R. and P. Jäckel (2000). 'The Most General Methodology to Create a Valid Correlation Matrix for Risk Management and Option Pricing Purposes', *Journal of Risk* 2: 17–27.

Rebonato, R., McKay, K. and R. White (2009). *The SABR/LIBOR Market Model*, Wiley.

Redon, C. (2006). Wrong Way Risk Modelling, *Risk Magazine*, April, 90–95.

Ribeiro, C. and Webber, N. (2004). Valuing path-dependent options in the variance-gamma model by Monte Carlo with a gamma bridge, *Journal of Computational Finance* 7: 81–100.

Rogers, L. (2002). Monte Carlo valuation of American options, *Mathematical Finance* 12: 271–286.

Rosen, D. and Saunders, D. (2012). CVA the Wrong Way, *Journal of Risk Management in Financial Institutions* 5: 252–272.

Samuelson, P. (1965). Rational theory of warrant pricing, *Industrial Management Review* 6: 13-39.

Schlögl, E. (2002a). Arbitrage-Free Interpolation in Models of Market Observable Interest Rates. In K. Sandmann and P. Schonbucher (Eds.), *Advances in Finance and Stochastics*, Springer, Heidelberg.

Schlögl, E. (2002b). A Multicurrency Extension of the Lognormal Interest Rate Market Models. *Finance and Stochastics* 6: 173–196.

Schoutens, W. (2003). *Lévy Processes in Finance: Pricing Financial Derivatives*, Wiley.

Schwartz, E. S. (1982). The Pricing of Commodity–Linked Bonds. *Journal of Finance* 37: 525–539.

Schwartz, E. S. (1997). The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *Journal of Finance* 52: 923–973.

Tilley, J.A. (1993). Valuing American Options in a Path Simulation Model, *Transactions of the Society of Actuaries* 45: 83–104.

Tompkins, R.G. (2001). Implied volatility surfaces: uncovering regularities for options of financial futures, *The European Journal of Finance* 7: 198–230.

Vasicek, O. (1977). An equilibrium characterization of the term structure, *Journal of Financial Economics* 5: 177-188.

Wilmott, P., Howison, S. and Dewynne, J. (1993). *The mathematics of financial derivatives: a student introduction*, Cambridge University Press, Cambridge.

Lund Economic Studies

1. Guy Arvidsson   Bidrag till teorin för verkningarna av räntevariationer, 1962

2. Björn Thalberg   A Trade Cycle Analysis. Extensions of the Goodwin Model, 1966

3. Bengt Höglund   Modell och observationer. En studie av empirisk anknytning och aggregation för en linjär produktionsmodell, 1968

4. Alf Carling   Industrins struktur och konkurrensförhållanden, 1968

5. Tony Hagström   Kreditmarknadens struktur och funktionssätt, 1968

6. Göran Skogh   Straffrätt och samhällsekonomi, 1973

7. Ulf Jakobsson och Göran Norman   Inkomstbeskattningen i den ekonomiska politiken. En kvantitativ analys av systemet för personlig inkomstbeskattning 1952-71, 1974

8. Eskil Wadensjö   Immigration och samhällsekonomi. Immigrationens ekonomiska orsaker och effekter, 1973

9. Rögnvaldur Hannesson   Economics of Fisheries. Some Problems of Efficiency, 1974

10. Charles Stuart   Search and the Organization of Marketplaces, 1975

11. S Enone Metuge   An Input-Output Study of the Structure and Resource Use in the Cameroon Economy, 1976

12. Bengt Jönsson   Cost-Benefit Analysis in Public Health and Medical Care, 1976

13. Agneta Kruse och Ann-Charlotte Ståhlberg   Effekter av ATP - en samhällsekonomisk studie, 1977

14. Krister Hjalte   Sjörestaureringens ekonomi, 1977

15. Lars-Gunnar Svensson   Social Justice and Fair Distributions, 1977

16. Curt Wells   Optimal Fiscal and Monetary Policy - Experiments with an Econometric Model of Sweden, 1978

17. Karl Lidgren   Dryckesförpackningar och miljöpolitik - En studie av styrmedel, 1978

18. Mats Lundahl   Peasants and Poverty. A Study of Haiti, London, 1979

19. Inga Persson-Tanimura   Studier kring arbetsmarknad och information, 1980

20. Bengt Turner   Hyressättning på bostadsmarknaden - Från hyresreglering till bruksvärdesprövning, Stockholm 1979

21. Ingemar Hansson   Market Adjustment and Investment Determination. A Theoretical Analysis of the Firm and the Industry, Stockholm 1981

22. Daniel Boda Ndlela   Dualism in the Rhodesian Colonial Economy, 1981

23. Tom Alberts — Agrarian Reform and Rural Poverty: A Case Study of Peru, 1981

24. Björn Lindgren — Costs of Illness in Sweden 1964-75, 1981

25. Göte Hansson — Social Clauses and International Trade. An Economic Analysis of Labour Standards in Trade Policy, 1981

26. Noman Kanafani — Oil and Development. A Case Study of Iraq, 1982

27. Jan Ekberg — Inkomsteffekter av invandring, 1983

28. Stefan Hedlund — Crisis in Soviet Agriculture?, 1983

29. Ann-Marie Pålsson — Hushållen och kreditpolitiken. En studie av kreditrestriktioners effekt på hushållens konsumtion, sparande och konsumtionsmönster, 1983

30. Lennart Petersson — Svensk utrikeshandel, 1871-1980. En studie i den intraindustriella handelns framväxt, 1984

31. Bengt Assarsson — Inflation and Relative Prices in an Open Economy, 1984

32. Claudio Vedovato — Politics, Foreign Trade and Economic Development in the Dominican Republic, 1985

33. Knut Ödegaard — Cash Crop versus Food Crop Production in Tanzania: An Assessment of the Major Post-Colonial Trends, 1985

34. Vassilios Vlachos — Temporära lönesubventioner. En studie av ett arbetsmarknadspolitiskt medel, 1985

35. Stig Tegle — Part-Time Employment. An Economic Analysis of Weekly Working Hours in Sweden 1963-1982, 1985

36. Peter Stenkula — Tre studier över resursanvändningen i högskolan, 1985

37. Carl Hampus Lyttkens — Swedish Work Environment Policy. An Economic Analysis, 1985

38. Per-Olof Bjuggren — A Transaction Cost Approach to Vertical Integration: The Case of Swedish Pulp and Paper Industry, 1985

39. Jan Petersson — Erik Lindahl och Stockholmsskolans dynamiska metod, 1987

40. Yves Bourdet — International Integration, Market Structure and Prices. A Case Study of the West-European Passenger Car Industry, 1987

41. Krister Andersson and Erik Norrman — Capital Taxation and Neutrality. A study of tax wedges with special reference to Sweden, 1987

42. Tohmas Karlsson — A Macroeconomic Disequilibrium Model. An Econometric Study of the Swedish Business Sector 1970-84, 1987

43. Rosemary Vargas-Lundius — Peasants in Distress. Poverty and Unemployment in the Dominican Republic, 1989

Method, 1997

67. Lars Pålsson Syll   Utility Theory and Structural Analysis, 1997

68. Richard Henricsson   Time Varying Parameters in Exchange Rate Models, 1997

69. Peter Hördahl   Financial Volatility and Time-Varying Risk Premia, 1997

70. Lars Nilsson   Essays on North-South Trade, 1997

71. Fredrik Berggren   Essays on the Demand for Alcohol in Sweden - Review and Applied Demand Studies, 1998

72. Henrik Braconier   Essays on R&D, Technology and Growth, 1998

73. Jerker Lundbäck   Essays on Trade, Growth and Exchange Rates, 1998

74. Dan Anderberg   Essays on Pensions and Information, 1998

75. P. Göran T. Hägg   An Institutional Analysis of Insurance Regulation – The Case of Sweden, 1998

76. Hans-Peter Bermin   Essays on Lookback and Barrier Options - A Malliavin Calculus Approach, 1998

77. Kristian Nilsson   Essays on Exchange Rates, Exports and Growth in Developing Countries, 1998

78. Peter Jochumzen   Essays on Econometric Theory, 1998

79. Lars Behrenz   Essays on the Employment Service and Employers' Recruitment Behaviour, 1998

80. Paul Nystedt   Economic Aspects of Ageing, 1998

81. Rasha M. Torstensson   Empirical Studies in Trade, Integration and Growth, 1999

82. Mattias Ganslandt   Games and Markets – Essays on Communication, Coordination and Multi-Market Competition, 1999

83. Carl-Johan Belfrage   Essays on Interest Groups and Trade Policy, 1999

84. Dan-Olof Rooth   Refugee Immigrants in Sweden - Educational Investments and Labour Market Integration, 1999

85. Karin Olofsdotter   Market Structure and Integration: Essays on Trade, Specialisation and Foreign Direct Investment, 1999

86. Katarina Steen Carlsson   Equality of Access in Health Care, 1999

87. Peter Martinsson   Stated preference methods and empirical analyses of equity in health, 2000

88. Klas Bergenheim   Essays on Pharmaceutical R&D, 2000

89. Hanna Norberg   Empirical Essays on Regional Specialization and Trade in Sweden, 2000

90. Åsa Hansson   Limits of Tax Policy, 2000

91. Hans Byström   Essays on Financial Markets, 2000

92. Henrik Amilon	Essays on Financial Models, 2000
93. Mattias Lundbäck	Asymmetric Information and The Production of Health, 2000
94. Jesper Hansson	Macroeconometric Studies of Private Consumption, Government Debt and Real Exchange Rates, 2001
95. Jonas Månsson	Essays on: Application of Cross Sectional Efficiency Analysis, 2001
96. Mattias Persson	Portfolio Selection and the Analysis of Risk and Time Diversification, 2001
97. Pontus Hansson	Economic Growth and Fiscal Policy, 2002
98. Joakim Gullstrand	Splitting and Measuring Intra-Industry Trade, 2002
99. Birger Nilsson	International Asset Pricing, Diversification and Links between National Stock Markets, 2002
100. Andreas Graflund	Financial Applications of Markov Chain Monte Carlo Methods, 2002
101. Therése Hindman Persson	Economic Analyses of Drinking Water and Sanitation in Developing Countries, 2002
102. Göran Hjelm	Macroeconomic Studies on Fiscal Policy and Real Exchange Rates, 2002
103. Klas Rikner	Sickness Insurance: Design and Behavior, 2002
104. Thomas Ericson	Essays on the Acquisition of Skills in Teams, 2002
105. Thomas Elger	Empirical Studies on the Demand for Monetary Services in the UK, 2002
106. Helena Johansson	International Competition, Productivity and Regional Spillovers, 2003
107. Fredrik Gallo	Explorations in the New Economic Geography, 2003
108. Susanna Thede	Essays on Endogenous Trade Policies, 2003
109. Fredrik CA Andersson	Interest Groups and Government Policy, A Political Economy Analysis, 2003
110. Petter Lundborg	Risky Health Behaviour among Adolescents, 2003
111. Martin W Johansson	Essays on Empirical Macroeconomics, 2003
112. Joakim Ekstrand	Currency Markets - Equilibrium and Expectations, 2003
113. Ingemar Bengtsson	Central bank power: a matter of coordination rather than money supply, 2003
114. Lars Pira	Staples, Institutions and Growth: Competitiveness of Guatemalan Exports 1524-1945, 2003
115. Andreas Bergh	Distributive Justice and the Welfare State, 2003

LUND UNIVERSITY

School of Economics and Management