The Journal of Computational Finance (1–35) Volume 15/Number 4, Summer 2012

Calibration and Monte Carlo pricing of the SABR–Hull–White model for long-maturity equity derivatives

Bin Chen
CWI – National Research Institute for Mathematics and Computer Science,
PO Box 94079, 1090 GB, Amsterdam, The Netherlands;
email: b.chen@cwi.nl
and
Rabobank International, Utrecht, The Netherlands

Lech A. Grzelak
CWI – National Research Institute for Mathematics and Computer Science,
PO Box 94079, 1090 GB, Amsterdam, The Netherlands;
email: l.a.grzelak@cwi.nl
and
Rabobank International, Utrecht, The Netherlands

Author: please check running head on odd pages – abbreviated form of title OK?

Please provide a full postal address for the Rabobank International affiliation.

Cornelis W. Oosterlee
CWI – National Research Institute for Mathematics and Computer Science,
PO Box 94079, 1090 GB, Amsterdam, The Netherlands;
email: c.w.oosterlee@cwi.nl
and
Delft Institute of Applied Mathematics, Delft University of Technology,
Mekelweg 4, 2628CD Delft, The Netherlands

We model the joint dynamics of stock prices and interest rates using a hybrid SABR–Hull–White model. The asset price dynamics are modeled by the SABR model of Hagan et al and the interest rate dynamics are modeled by the Hull–White short-rate model. We propose a projection formula, mapping the SABR–Hull–White model parameters onto the parameters of the nearest SABR model. Furthermore, a time-dependent parameter extension of this SABR–Hull–White model is introduced to make the calibration of the model consistent across maturities. The inverse of the projection formula enables a rapid calibration of the model. As the calibration quality is subject to the approximation errors of the projection formula, we subsequently apply a nonparametric numerical calibration technique based on the nonuniformly weighted Monte Carlo technique of Avellaneda et al to improve the calibration. In this step, the Monte Carlo weights are not uniform and are chosen in such a way that the calibration market instruments are perfectly replicated.

Direct citations removed from abstract according to journal style – OK? Also, changes to first sentence of text OK?

Changes to sentence OK?

1 INTRODUCTION

Equity derivative models and yield-curve models have been developed independently of each other for a long time. Whereas equity derivative models have been focused on the implied volatility skew/smile by local or stochastic volatility features (Gatheral (2006)), short-rate models have improved the accuracy of the yield-curve dynamics. With an increasing interest in long-maturity equity derivatives, as well as in equity–interest rate hybrid products from retail and long-term institutional investors (see Hunter and Picot (2005)), the industrial practice also demands models that are capable of describing the joint dynamics of interest rates and equity. Indeed, one would intuitively assert that the interest rate is stochastic and that there is nonzero correlation between the interest rate and the equity.

We propose a hybrid extension of the SABR ("stochastic $\alpha\beta\rho$") model (Hagan *et al* (2002)) to model the joint equity–interest rate dynamics. We construct a hybrid model, called the SABR–Hull–White (SABR–HW) model, in which the equity process is driven by the SABR model and the interest rate by the short-rate model of Hull and White (1996). In this framework, the equity process is assumed to be correlated with the interest rate process.

The SABR model, which is not often used in the equity derivative literature, has several attractive features for the modeling of long-term equity-linked products. It generalizes stochastic volatility models such as Heston's model by introducing an explicit stock price dependence in a power law local volatility term $f(x) = x^{\beta}$ (Mendoza *et al* (2010)). Furthermore, the SABR model admits a closed-form approximation formula ("Hagan's formula") for the Black implied volatilities, which greatly simplifies calibration. Third, the SABR process is, for certain parameter values, an absorbing process at the zero asset price boundary (Chen *et al* (2012)), which models the fact that companies may default in time (Mendoza *et al* (2010)). Last, but not least, the parameters in the SABR model have a direct connection to market instruments or market price features, in contrast with, for example, the speed-of-mean-reversion parameter in the Heston model.

One contribution made in this paper is an invertible projection formula of the constant-parameter SABR–HW model onto the plain SABR model. This formula enables a highly efficient calibration of the constant-parameter SABR–HW model based on the established calibration procedure for the SABR model. However, the resulting calibration parameters only remain valid for a single maturity time and cannot provide a consistent dynamic description of the underlying asset prices across multiple maturity times. We deal with this issue by adopting a dynamic SABR(–HW) extension in the spirit of Rebonato (2006).

'dynamical' changed to 'dynamic' throughout in keeping with much more common notation according to online research. OK?

Moreover, we will use the well-known weighted Monte Carlo technique proposed by Avellaneda *et al* (2000) as another stage of calibration of the SABR–HW model.

Using this, we can deal with the inconsistency between the true model dynamics and those implied by Hagan's asymptotic approximation formula (by which the calibration instruments are quoted).

For the Monte Carlo simulation, we adopt a low-bias discretization for the SABR–HW dynamics, which has some advantages over a basic Euler scheme as it gives a low bias (ie, stable and accurate) when large time steps are used (say, four time steps per year) (see Chen *et al* (2012)).

This paper is organized as follows. In Section 2 we define the dynamic SABR–HW model and discuss the building blocks. We show how to project the constant-parameter SABR–HW dynamics onto a plain SABR model in Section 3. In Section 4 we utilize this projection formula to calibrate the constant-parameter version of the model. We also show how to determine the time-dependent functions in the dynamic SABR–HW model and use the weighted Monte Carlo technique (Avellaneda *et al* (2000)). In Section 5 the low-bias Monte Carlo simulation for the dynamic SABR–HW model is presented. Numerical experiments for validation and calibration are discussed throughout the paper.

2 THE DYNAMIC SABR–HULL–WHITE MODEL

This section describes the construction of the dynamic SABR–HW equity–interest rate model.

We assume efficient markets and the existence of an equivalent martingale measure $\mathbb{Q}$ when appropriate numeraires need to be chosen.

2.1 Model definition

We define the dynamic SABR–HW model in a similar fashion to Rebonato's SABR–LIBOR market model (Rebonato (2006)) for forward rates. The full-scale dynamic SABR–HW model for equity–interest rate products, under the $\mathbb{Q}$-measure associated with $B(t)$ (a money-saving account), is given by:

$$\left.\begin{aligned} \mathrm{d}S(t)/S(t) &= r(t)\,\mathrm{d}t + \Sigma(t)S^{\beta-1}(t)\,\mathrm{d}W_x(t), \quad S(0) > 0 \\ \mathrm{d}r(t) &= \lambda(\theta(t) - r(t))\,\mathrm{d}t + \eta\,\mathrm{d}W_r(t), \quad r(0) > 0 \\ \Sigma(t) &= g(t)k(t) \\ \mathrm{d}k(t) &= h(t)k(t)\,\mathrm{d}W_\Sigma(t), \quad k(0) = 1 \\ \mathrm{d}W_x(t)\,\mathrm{d}W_\sigma(t) &= \rho_{x,\sigma}\,\mathrm{d}t, \quad -1 \leqslant \rho_{x,\sigma} \leqslant 1 \\ \mathrm{d}W_x(t)\,\mathrm{d}W_r(t) &= \rho_{x,r}\,\mathrm{d}t, \quad -1 \leqslant \rho_{x,r} \leqslant 1 \end{aligned}\right\} \tag{2.1}$$

with constant parameters β, λ, η, $\rho_{x,r}$ and $\rho_{x,\sigma}$, and appropriately chosen time-dependent functions $\theta(t)$, $g(t)$ and $h(t)$. For simplicity, we assume here that the

interest rates are independent of the stochastic volatility, $\mathrm{d}W_r(t)\,\mathrm{d}W_\sigma(t) = 0$. The parameters will be discussed in the sections to follow, and details of the functional form of $g(t)$ and $h(t)$ are given in Section 4.2.

REMARK 2.1 For some additional insight into the functions $g(t)$ and $h(t)$, we derive the dynamics of the time-dependent volatility $\Sigma(t)$ by applying Ito's lemma:

$$\begin{aligned}\mathrm{d}\Sigma(t) \equiv \mathrm{d}(g(t)k(t)) &= k(t)\,\mathrm{d}g(t) + g(t)\,\mathrm{d}k(t) + \mathrm{d}k(t)\,\mathrm{d}g(t)\\ &= \left(\frac{1}{g(t)}\frac{\mathrm{d}g(t)}{\mathrm{d}t}\right)\Sigma(t)\,\mathrm{d}t + h(t)\Sigma(t)\,\mathrm{d}W_\Sigma(t)\\ &= \frac{\mathrm{d}\log g(t)}{\mathrm{d}t}\Sigma(t)\,\mathrm{d}t + h(t)\Sigma(t)\,\mathrm{d}W_\Sigma(t)\end{aligned}$$

Overall, this implies that the dynamics for the volatility can be defined as:

$$\mathrm{d}\Sigma(t) = \hat{g}(t)\Sigma(t)\,\mathrm{d}t + h(t)\Sigma(t)\,\mathrm{d}W_\Sigma(t)$$

with:

$$\hat{g}(t) = \frac{\mathrm{d}\log g(t)}{\mathrm{d}t}$$

which means that we deal with a lognormal process with time-dependent drift and volatility terms. It is obvious that the function $h(t)$ plays the role of volatility coefficient for this volatility process and the function $g(t)$, appearing in the drift term, shifts the volatility up and down deterministically.

2.1.1 The HW model

One of the building blocks of hybrid model (2.1) is the HW single-factor no-arbitrage yield-curve model in which the short-term interest rate, $r(t)$, is driven by an Ornstein–Uhlenbeck mean-reverting process, with $\theta(t) > 0$, $t \in \mathbb{R}^+$, a time-dependent drift term, to fit theoretical bond prices to the yield curve observed in the market. Parameter η determines the overall level of volatility and the reversion rate parameter, λ, determines the relative volatilities.

Under the HW model, the dynamics of the zero-coupon bond, paying €1 at maturity T, are given by:

$$\frac{\mathrm{d}P(t,T)}{P(t,T)} = r(t)\,\mathrm{d}t + \frac{\eta}{\lambda}(\mathrm{e}^{-\lambda(T-t)} - 1)\,\mathrm{d}W_r(t) \tag{2.2}$$

Since the HW model belongs to the class of affine diffusion processes, the solution of (2.2) is known analytically and reads:

$$P(t,T) = \exp(A(t,T) + B_r(t,T)r(t)) \tag{2.3}$$

with:

$$B_r(t,T) = \frac{1}{\lambda}(\mathrm{e}^{-\lambda(T-t)} - 1) \tag{2.4}$$

$$A(t,T) = \exp\left(\log\left(\frac{P(0,T)}{P(0,t)}\right) - B_r(t,T)f(0,t) - \frac{\eta^2}{4\lambda}(1-\mathrm{e}^{-2\lambda t})(B_r(t,T))^2\right) \tag{2.5}$$

where $f(0,t) := -\partial P(0,t)/\partial t$, with $P(0,t)$ the market discount factor for maturity t.

By the Radon–Nikodým derivative (Geman *et al* (1996)):

$$\frac{\mathrm{d}\mathbb{Q}^T}{\mathrm{d}\mathbb{Q}} = \frac{P(t,T)}{P(0,T)B(t)} \tag{2.6}$$

we find the following change of measure: $\mathrm{d}W_r^T(t) = \mathrm{d}W_r(t) - \eta B_r(t,T)\,\mathrm{d}t$.

The short rate $r(t)$ under the T-forward measure is governed by the following dynamics:

$$\mathrm{d}r(t) = (\lambda(\theta(t) - r(t)) + \eta^2 B_r(t,T))\,\mathrm{d}t + \eta\,\mathrm{d}W_r^T(t)$$

which can be written as:

$$\mathrm{d}r(t) = \lambda(\hat{\theta}(t) - r(t))\,\mathrm{d}t + \eta\,\mathrm{d}W_r^T(t) \tag{2.7}$$

with:

$$\hat{\theta}(t) = \theta(t) + \frac{\eta^2}{\lambda}B_r(t,T)$$

and $B_r(t,T)$ in (2.4). Since the process under the T-forward measure in (2.7) is of "HW form", it is normally distributed (Brigo and Mercurio (2007)) with expectation and variance given by:

$$\mathbb{E}^T(r(t)) = r_0\mathrm{e}^{-\lambda t} + \lambda\int_0^t \hat{\theta}(s)\mathrm{e}^{-\lambda(t-s)}\,\mathrm{d}s \tag{2.8}$$

$$\mathrm{var}^T(r(t)) = \frac{\eta^2}{2\lambda}(1-\mathrm{e}^{-2\lambda t}) \tag{2.9}$$

A disadvantage of the HW model is that it may give rise to negative interest rates. The negative interest rate, however, may be present in the real market.[1] An alternative to the HW model is the Cox–Ingersoll–Ross (CIR) model. A hybrid SABR–CIR model is, however, not tractable if there is a nonzero correlation between the interest rate and the SABR equity process. The choice between a CIR or an HW model within the hybrid process is a trade-off between nonzero correlation and nonnegative rates.

For hybrid structured products, a nonzero correlation is a crucial feature that should be incorporated into a model (see Grzelak and Oosterlee (2011) for analysis and further

Change OK?

[1] See http://en.wikipedia.org/wiki/Interest_rate.

arguments), whereas the appearance of negative interest rates in an HW process is an inherent feature of the model and has been known by practitioners for quite some time (Brigo and Mercurio (2007)). There are practical solutions to this problem, eg, choosing parameters that give rise to lower probabilities for negative rates. We therefore prefer the HW process over the CIR process as part of our equity–interest rate hybrid model.

'equity' and 'interest rate' switched around to match earlier formulation – OK?

2.1.2 The constant-parameter SABR–HW model

The second building block of model (2.1) is the SABR stochastic volatility model of Hagan *et al* (2002).

The SABR stochastic differential equation (SDE) system with constant parameters was originally defined under the T-forward measure as:

$$\left.\begin{aligned} \mathrm{d}S(t) &= \sigma(t)S(t)^{\beta}\,\mathrm{d}W_x^T(t) \\ \mathrm{d}\sigma(t) &= \gamma\sigma(t)\,\mathrm{d}W_\sigma^T(t) \end{aligned}\right\} \tag{2.10}$$

with $\mathrm{d}W_x^T(t)\,\mathrm{d}W_\sigma^T(t) = \rho_{x,\sigma}\,\mathrm{d}t$.

One of the reasons why the original SABR model is not applied to equity derivatives is that a drift term is lacking. Risk-neutral equity price processes are defined with a drift term, and are assumed to be arbitrage-free under the risk-neutral measure associated with the money-savings account. For long-maturity equity options and equity–interest rate hybrids, however, industrial practice is to treat the interest rate as a stochastic process as well. As shown below, when combining the HW interest rate model with the SABR equity model, the drift term appears naturally in the SABR equity dynamics under the risk neutral $\mathbb{Q}$-measure:

$$\left.\begin{aligned} \mathrm{d}S(t)/S(t) &= r(t)\,\mathrm{d}t + \sigma(t)S^{\beta-1}(t)\,\mathrm{d}W_x(t), & S(0) &> 0 \\ \mathrm{d}\sigma(t) &= \gamma\sigma(t)\,\mathrm{d}W_\sigma(t), & \sigma(0) &> 0 \\ \mathrm{d}r(t) &= \lambda(\theta(t) - r(t))\,\mathrm{d}t + \eta\,\mathrm{d}W_r(t), & r(0) &> 0 \end{aligned}\right\} \tag{2.11}$$

with constant model parameters $0 < \beta < 1$, $\gamma > 0$, $\lambda > 0$ and $\eta > 0$. As in system (2.1), we assume nonzero correlations:

$$\begin{aligned} \mathrm{d}W_x(t)\,\mathrm{d}W_\sigma(t) &= \rho_{x,\sigma}\,\mathrm{d}t \\ \mathrm{d}W_x(t)\,\mathrm{d}W_r(t) &= \rho_{x,r}\,\mathrm{d}t \\ \mathrm{d}W_r(t)\,\mathrm{d}W_\sigma(t) &= 0 \end{aligned}$$

Since the interest rate diffusion coefficient in (2.11) is not explicitly dependent on the interest rate, it is convenient to move from the spot measure, generated by the

money-savings account, $B(t)$, to the forward measure, for which the numeraire is the zero-coupon bond, $P(t,T)$:

$$F(t) := \frac{S(t)}{P(t,T)} \tag{2.12}$$

(details of $P(t,T)$ are given in (2.3)).

By Ito's lemma, the dynamics of forward price $F(t)$ in (2.12) are given by:

$$\begin{aligned}\mathrm{d}F(t)/F(t) = (\eta^2 B_r^2(t,T) - \rho_{x,r}\eta B_r(t,T)\sigma(t)S^{\beta-1}(t))\,\mathrm{d}t \\ + \sigma(t)S^{\beta-1}(t)\,\mathrm{d}W_x(t) - \eta B_r(t,T)\,\mathrm{d}W_r^T(t)\end{aligned}$$

combined with the volatility process for $\sigma(t)$ in system (2.11). Since the forward $F(t)$ is a martingale under the T-forward measure, the forward dynamics should not contain a drift term. This implies that "dt" terms will not appear in the (reformulated) dynamics of $\mathrm{d}F(t)$, ie:

$$\left.\begin{aligned}\mathrm{d}F(t) &= \sigma(t)F^{\beta}(t)\left(P^{\beta-1}(t,T)\,\mathrm{d}W_x^T(t) - \frac{\eta B_r(t,T)F(t)}{\sigma(t)F^{\beta}(t)}\,\mathrm{d}W_r^T(t)\right)\\ \mathrm{d}\sigma(t) &= \gamma\sigma(t)\,\mathrm{d}W_\sigma^T(t)\end{aligned}\right\} \tag{2.13}$$

We assume that the interest rate is independent of the volatility process, so a change of measure will not affect the dynamics of the variance process $\sigma(t)$.

By factorization, model (2.13) can be expressed as:

$$\left.\begin{aligned}\mathrm{d}F(t) &= \sigma(t)v(t)F^{\beta}(t)\,\mathrm{d}W_F^T(t)\\ \mathrm{d}\sigma(t) &= \gamma\sigma(t)\,\mathrm{d}W_\sigma^T(t)\end{aligned}\right\} \tag{2.14}$$

with:

$$v^2(t) := P^{2(\beta-1)}(t,T) + \left(\frac{\eta B_r(t,T)}{\sigma(t)F^{\beta-1}(t)}\right)^2 - 2\rho_{x,r}\frac{\eta B_r(t,T)P^{\beta-1}(t,T)}{\sigma(t)F^{\beta-1}(t)} \tag{2.15}$$

Now, the instantaneous correlation coefficient $\rho_{F,\sigma}$ must be determined, which is defined as:

$$\begin{aligned}\rho_{F,\sigma} &= \frac{\mathrm{cov}(\mathrm{d}F(t),\mathrm{d}\sigma(t))}{\sqrt{\mathrm{var}(\mathrm{d}F(t))\,\mathrm{var}(\mathrm{d}\sigma(t))}}\\ &= \rho_{x,\sigma}\Psi(t,\sigma(t),F(t),P(t,T))\end{aligned} \tag{2.16}$$

with:[2]

$$\begin{aligned}&\Psi(t,\sigma(t),F(t),P(t,T))\\ &\qquad := \frac{F^{\beta}P^{\beta-1}\sigma}{\sqrt{\sigma^2F^{2\beta}P^{2(\beta-1)} + \eta^2B_r^2F^2 - \rho_{x,r}\sigma F^{\beta+1}P^{\beta-1}\eta B_r}}\end{aligned} \tag{2.17}$$

[2] To simplify notation we suppress the arguments t and T here.

Model (2.14) with (2.15) and (2.16) is not in the well-known plain SABR model form because the local volatility is not expressed only by $\sigma(t)F^{\beta}(t)$ but also contains additional terms like $v(t)$. Moreover, the instantaneous correlation between forward and volatility processes $\rho_{F,\sigma}$ is a state-dependent function of time. In order to make use of Hagan's asymptotic formulas (Hagan *et al* (2002)) for the plain SABR model in the current setting, we propose a projection formula in the next section.

Changes to sentence OK?

3 PROJECTION FORMULA FOR THE CONSTANT-PARAMETER SABR–HULL–WHITE MODEL

In this section we describe the model approximations that bring the SABR–HW model into the desired SABR model form. The approximations enable us to carry out an efficient calibration based on the analytic implied volatility formulas for the SABR model.

3.1 Projection step for the constant-parameter SABR–HW model

In order to present model (2.14) in SABR form, we need to approximate the additional terms from the local volatility for the forward process $F(t)$ and simplify the associated correlation structure. In a plain SABR model, the volatility process $\sigma(t)$ is lognormal, which suggests that a projection of the volatility term $\sigma(t)v(t)$ in (2.14) on a lognormal distribution may give the desired SABR form, which is:

$$\left.\begin{aligned} \mathrm{d}F(t) &= \hat{\sigma}(t)F^{\beta}(t)\,\mathrm{d}W_F^T(t), \quad F(0) > 0 \\ \mathrm{d}\hat{\sigma}(t) &= \hat{\gamma}\hat{\sigma}(t)\,\mathrm{d}W_{\sigma}^T(t), \qquad \hat{\sigma}(0) > 0 \end{aligned}\right\} \tag{3.1}$$

with constant parameters $\hat{\sigma}(0)$ and $\hat{\gamma}$, and constant correlation $\hat{\rho}_{F,\sigma}$.

The term $v(t)$ in (2.15) depends on forward $F^{\beta-1}(t)$, volatility $\sigma(t)$ and on zero-coupon bond $P(t,T)$. With a function $v(t)$, which is independent of these state variables, the expression simplifies. This can be achieved by freezing the forward and variance, $F(t)$ and $\sigma(t)$, at their initial values, ie, $F(t) \approx F(0)$ and $\sigma(t) \approx \sigma(0)$, respectively, and by projecting $P(t,T)$ on its expectation, ie:

$$P(t,T) \approx \mathbb{E}^T[P(t,T) \mid \mathcal{F}_t] =: \xi(t)$$

Function $v^2(t)$ is then approximated by:

$$v^2(t) \approx \xi^{2(\beta-1)}(t) + \left(\frac{\eta B_r(t,T)}{\sigma(0)F^{\beta-1}(0)}\right)^2 - 2\rho_{x,r}\frac{\eta B_r(t,T)\xi^{\beta-1}(t)}{\sigma(0)F^{\beta-1}(0)} \tag{3.2}$$

With the help of the well-known formulas (2.2) and (2.7), we obtain the following closed-form solution for $\xi(t)$:

$$\xi(t) = \exp(A(t,T) + B_r(t,T)\mathbb{E}^T[r(t)] + \tfrac{1}{2}B_r^2(t,T)\,\mathrm{var}^T(r(t))) \tag{3.3}$$

with $\mathbb{E}^T[r(t)]$ and $\text{var}^T(r(t))$ given by (2.8) and (2.9), respectively, and $A(t,T)$ defined in (2.5).

Note that $F(t)$ and $\sigma(t)$ are both martingales, due to (2.14), which implies that the values of their expectations oscillate around their initial values, $F^{\beta-1}(0)$ and $\sigma(0)$. Function $v(t)$ has become deterministic from the approximations made.

We then determine the dynamics for the linearized volatility structure, $\bar{\sigma}(t) := \sigma(t)v(t)$. By applying the Ito product rule, we find:

$$\mathrm{d}\bar{\sigma}(t)/\bar{\sigma}(t) = v'(t)\,\mathrm{d}t + \gamma v(t)\,\mathrm{d}W_\sigma^T(t)$$

The $\bar{\sigma}(t)$-dynamics are thus governed by a state-dependent drift term. Therefore, they are not yet in standard SABR volatility form, which does not contain any drift term.

However, since $v(t)$ is approximated by a deterministic time-dependent function, the process $\bar{\sigma}(t) := v(t)\sigma(t)$ remains lognormal. The idea is now to determine the first two moments of process $\bar{\sigma}(t)$ and to project them onto the moments of the SABR volatility process in (3.1), defined as:

$$\mathrm{d}\hat{\sigma}(t) = \hat{\gamma}\hat{\sigma}(t)\,\mathrm{d}W_\sigma^T(t), \quad \hat{\sigma}(0) > 0$$

with parameters $\hat{\gamma}$ and $\hat{\sigma}(0)$.

The expectation and variance of process $\hat{\sigma}(t)$ in (3.1) are given by:

$$\mathbb{E}^T[\hat{\sigma}(t)] = \hat{\sigma}(0), \qquad \text{var}^T(\hat{\sigma}(t)) = \hat{\sigma}^2(0)(\mathrm{e}^{\hat{\gamma}^2 t} - 1) \tag{3.4}$$

On the other hand, the expectation and the variance of $\bar{\sigma}(t) = v(t)\sigma(t)$ are given by:

$$\mathbb{E}^T[\bar{\sigma}(t)] = v(t)\sigma(0), \qquad \text{var}^T(\bar{\sigma}(t)) = v^2(t)\sigma^2(0)(\mathrm{e}^{\gamma^2 t} - 1) \tag{3.5}$$

The main objective is to find the effective parameters $\hat{\sigma}(0)$ and $\hat{\gamma}$, so that the expectations and variances in (3.4) and (3.5) match.

By matching the expectations and variances, we arrive at the following optimization problem:

$$\underset{\hat{\gamma},\hat{\sigma}(0)}{\text{argmin}} \begin{cases} \displaystyle\int_0^T (\mathbb{E}^T[\hat{\sigma}(t)] - \mathbb{E}^T[\bar{\sigma}(t)])\,\mathrm{d}t \\ \displaystyle\int_0^T (\text{var}^T(\hat{\sigma}(t)) - \text{var}^T(\bar{\sigma}(t)))\,\mathrm{d}t \end{cases} \tag{3.6}$$

Typically, the optimization problem in (3.6) is easy since the expectations and variances are analytic deterministic functions. In Result 3.1, a straightforward approach for parameter estimation is presented.

RESULT 3.1 A simple approximation for $\hat{\sigma}(0)$ is the averaged parameter estimate, given by:

$$\int_0^T \mathbb{E}^T[\hat{\sigma}(t)]\,\mathrm{d}t = \int_0^T \mathbb{E}^T[\bar{\sigma}(t)]\,\mathrm{d}t \implies \hat{\sigma}(0) = \frac{\sigma(0)}{T}\int_0^T v(s)\,\mathrm{d}s \tag{3.7}$$

By matching the variances we obtain:

$$\int_0^T \mathrm{var}^T(\hat{\sigma}(t))\,\mathrm{d}t = \int_0^T \mathrm{var}^T(\bar{\sigma}(t))\,\mathrm{d}t$$

$$\Longrightarrow\ \hat{\gamma} = \frac{1}{T}\int_0^T \sqrt{\frac{1}{s}\log\left(\frac{\mathrm{var}^T(\bar{\sigma}(s))}{\hat{\sigma}^2(0)}+1\right)}\,\mathrm{d}s \tag{3.8}$$

with $\mathrm{var}^T(\bar{\sigma}(t))$ given in (3.5).

In the search for the optimal parameters $\hat{\sigma}(0)$ and $\hat{\gamma}$, so that the constant SABR–HW model in (2.11) is connected to the SABR dynamics given by (3.1), the correlation $\rho_{F,\sigma}$ has not yet been included. This allows us to determine the effective correlation, $\hat{\rho}_{F,\sigma}$, independent of the other approximations. Since the equation for correlation $\rho_{F,\sigma}$ is involved and state-dependent, we look for a simplification here as well. By freezing the volatility and forward to their initial values, and by projection of zero-coupon bonds on their expectations, the correlation in Equation (2.16) can be approximated by:

$$\rho_{F,\sigma} \approx \rho_{x,\sigma}\Psi(t,\sigma(0),F(0),\mathbb{E}^T[P(t,T)\mid\mathcal{F}_t] =: \rho_{x,\sigma}\psi(t) \tag{3.9}$$

Unmatched opening parenthesis in this display – remove, or add a closing parenthesis for balance? Please advise a clear course of action.

In order to use Hagan's implied volatility SABR formula, the correlation must be constant, so we need to determine an averaged correlation, which is defined as:

$$\hat{\rho}_{F,\sigma} = \frac{\rho_{x,\sigma}}{T}\int_0^T \psi(s)\,\mathrm{d}s \tag{3.10}$$

with $\psi(t)$ given by (3.9).

The estimates obtained for $\hat{\gamma}$, $\hat{\sigma}(0)$ and $\hat{\rho}_{F,\sigma}$ allow us to use the Hagan implied volatility formula for the plain SABR model as a first approximation in the calibration procedure.

REMARK 3.2 Our approximations in (3.7) and (3.8) perform well for relatively short maturity times such as $T \leqslant 10$ years. In the case of larger maturity times, such as $T \geqslant 10$ years, we prefer to solve problem (3.6) using an optimization procedure (for example, the Nelder–Mead simplex algorithm). Furthermore, the weighted Monte Carlo method, to be discussed in Section 4, will be used to improve the calibration in those cases.

Changes to this sentence and the next OK?

In the next section we check the accuracy of the approximations developed for a few parameter sets. The SABR model in (3.1) with the parameters $\hat{\gamma}$, $\hat{\sigma}(0)$ and $\hat{\rho}_{F,\sigma}$ will be called the SABR–HW$_1$ model here.

3.2 Numerical validation of the SABR–HW projection method

We check the performance of our approximation model, SABR–HW$_1$, in comparison with the constant-parameter SABR–HW model.

TABLE 1 Sets of parameters used in the simulations.

Parameter	β (%)	γ (%)	λ (%)	η (%)	$\sigma(0)$ (%)	$\rho_{x,\sigma}$ (%)	$\rho_{x,r}$ (%)
Set 1	30	30	20	1	20	−30	20
Set 2	50	40	1	0.5	20	−10	40
Set 3	40	10	60	0.1	30	−30	−30

The numerical experiment is set up as follows. First we prescribe a set of parameters for the constant-parameter SABR–HW model in (2.11) for which, by means of an Euler-based Monte Carlo scheme, the European option prices are simulated. Secondly, we compute the effective parameters $\hat{\gamma}$, $\hat{\sigma}(0)$ and $\hat{\rho}_{F,\sigma}$ by solving (3.6) and calculating (3.10). These parameters are inserted into the plain SABR model (3.1). For the resulting SABR–HW$_1$ model, we then calculate the corresponding implied volatilities using Hagan's asymptotic formula. We compare these results and, in addition, we determine the error for the case where γ, $\sigma(0)$ and $\rho_{x,\sigma}$ were used instead of $\hat{\gamma}$, $\hat{\sigma}(0)$ and $\hat{\rho}_{x,\sigma}$.

The simulations were performed with 100 000 paths and $20T$ steps. The initial stock price is set to $S(0) = 0.8$, and the zero-coupon bonds, $P(0,T)$, were generated by the HW model with constant long-term mean, $\theta = 0.03$. We also define the strikes, as in Piterbarg (2006), with expiry times given by $T \in \{1, 5, 10, 15\}$ years. The strikes are computed by the formula:

$$K_n(T) = F(0)\exp(0.1\delta_n\sqrt{T}) \quad \text{with } \delta_n = \{-1.5, -1.0, -0.5, 0, 0.5, 1.0, 1.5\} \tag{3.11}$$

and $F(0)$ is as in (2.12). This formula for the strikes is convenient, since, for $n = 4$, the strikes $K_4(\cdot)$ are equal to the forward prices.

In Table 1 we present three different sets of parameters. For those sets we determine the estimators $\hat{\gamma}$, $\hat{\sigma}(0)$ and $\hat{\rho}_{F,\sigma}$. They are shown in Table 2 on the next page.

We measure the maximum absolute difference in the implied volatilities for model (3.1) with the estimates in (3.7), (3.8) and the constant-parameter SABR–HW model (2.11).

Two errors are defined: error 1 is the error when the naive approach is used, ie, $\hat{\gamma} = \gamma$ and $\hat{\sigma}(0) = \sigma(0)$; error 2 corresponds to the bias obtained using the adjusted parameters $\hat{\gamma}$ and $\hat{\sigma}(0)$ (the SABR–HW$_1$ model). Table 3 on the next page presents these results.

Our approach for $\hat{\gamma}$, $\hat{\sigma}(0)$ and $\hat{\rho}_{F,\sigma}$ provides a significantly better fit to the constant-parameter SABR–HW model than the model with the naively chosen parameters. For

TABLE 2 Effective constant parameters $\hat{\gamma}$, $\hat{\sigma}(0)$ and $\hat{\rho}_{F,\sigma}$, defined in (3.6), and determined by solving the nonlinear least-squares problem (by MATLAB function "lsqcurvefit").

	T			
Estimators	**1 year (%)**	**5 years (%)**	**10 years (%)**	**15 years (%)**
Set 1				
$\hat{\gamma}$	29.76	29.02	28.34	27.85
$\hat{\sigma}(0)$	20.30	21.51	23.01	24.56
$\hat{\rho}_{F,\sigma}$	−30.08	−30.25	−30.30	−30.29
Set 2				
$\hat{\gamma}$	39.75	38.93	38.16	37.50
$\hat{\sigma}(0)$	20.24	21.28	22.73	24.52
$\hat{\rho}_{F,\sigma}$	−10.00	−10.01	−9.97	−9.90
Set 3				
$\hat{\gamma}$	9.97	9.78	9.56	9.35
$\hat{\sigma}(0)$	30.26	31.36	32.83	34.39
$\hat{\rho}_{F,\sigma}$	−30.01	−30.02	−30.02	−30.02

TABLE 3 The absolute maximum percentage difference between implied volatilities from two different models.

	$T = 1$ year		$T = 5$ years		$T = 10$ years		$T = 15$ years	
	Error 1 (%)	**Error 2 (%)**	**Error 1 (%)**	**Error 2 (%)**	**Error 1 (%)**	**Error 2 (%)**	**Error 1 (%)**	**Error 2 (%)**
Set 1	0.36	0.02	1.63	0.19	2.93	0.43	3.94	0.61
Set 2	0.27	0.01	1.38	0.08	2.65	0.32	3.74	0.51
Set 3	0.31	0.03	1.63	0.04	3.12	0.07	4.31	0.11

Both errors relate to the constant-parameter SABR–HW model and respective approximations.

the maturity times of one, five and fifteen years, Figure 1 on the facing page presents the corresponding implied volatilities.

The performance of the SABR–HW_1 approximation is most accurate when the volatility for the short rate, determined by η, is not too large, ie, $\eta < 1.5\%$. Fortunately, this is very often the case in the calibration of the HW model to market data. In the experiments we have chosen parameter $\beta \leqslant 50\%$ (in the case $\beta > 50\%$ an even better accuracy is expected, because then the model behavior is closer to that of a lognormal model (Chen *et al* (2010))). We also see that the correlation approximation, $\hat{\rho}_{F,\sigma}$, is close to the initial correlation $\rho_{F,\sigma}$. This is because the func-

Word added – OK?

FIGURE 1 Comparison of implied Black–Scholes volatilities for European equity options and parameter set 2 in Table 1 on page 11.

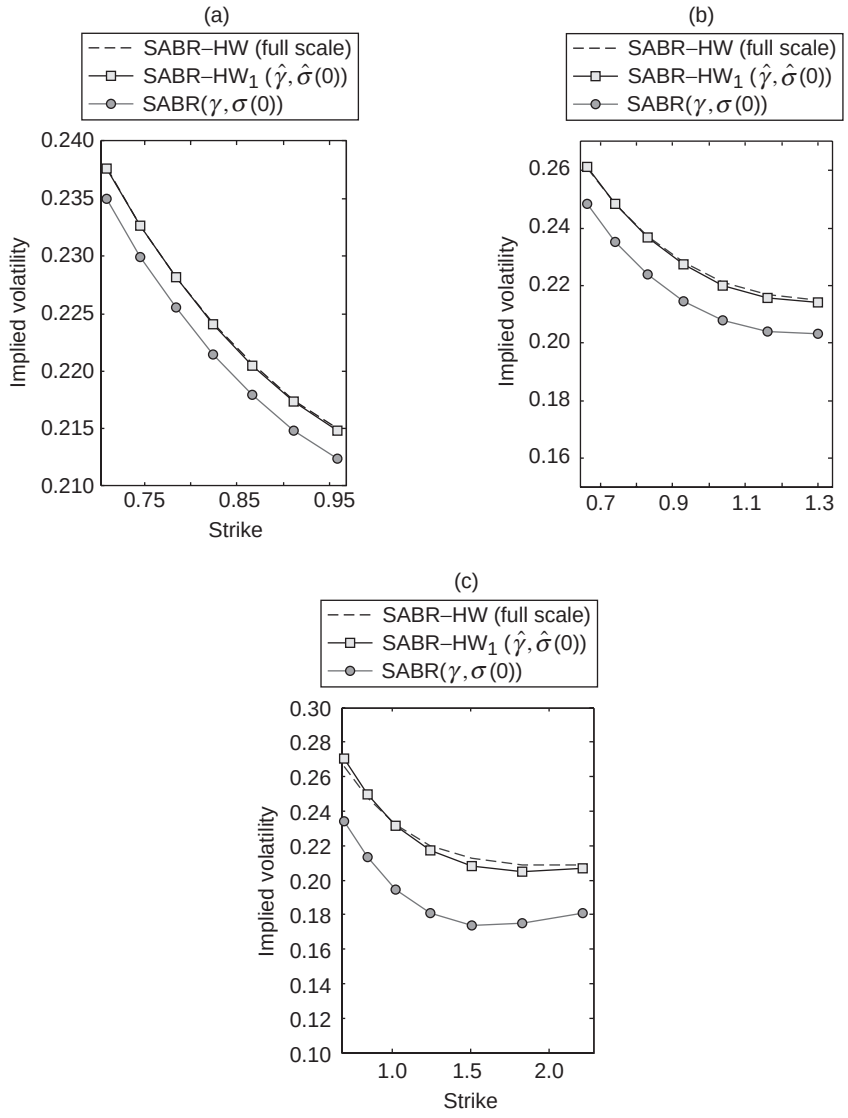

Implied volatilities for (a) $T = 1$, (b) $T = 5$ and (c) $T = 15$. For the SABR–HW model, Euler Monte Carlo was used with 100 000 paths and $20T$ intermediate steps.

tion $\Psi(t, \sigma(t), F(t), P(t,T))$ in (2.17) converges to 1 as $t \to T$, implying that $\rho_{F,\sigma} \to \hat{\rho}_{F,\sigma}$.

4 THE CALIBRATION PROCEDURE

We present a calibration procedure for the SABR–HW model in three stages, and start by applying the inverse projection formulas from the previous section to calibrate the constant-parameter SABR–HW model for every single maturity. In the second stage, we determine the parameters of the time-dependent functions in the dynamic SABR model in order to produce coherent model dynamics across the different maturities. In the final stage, the calibration is refined by means of a weighted Monte Carlo simulation. These stages are discussed in subsequent subsections.

4.1 Stage I: parameter projection for the SABR–HW model

In the calibration of the SABR–HW model, the HW part, which is connected to the function $\theta(t)$, is calibrated to the yield curve, whereas the parameters λ and η are calibrated to swaption prices separately. This is well-known (we refer the interested reader to Brigo and Mercurio (2007) for further information on this topic). The asset–interest rate correlation will be prescribed *a priori* based on historical data.

After the calibration of the HW model, we consider the determination of the parameters of the stochastic volatility SABR part.

One of the consequences of the projection of the constant-parameter SABR–HW model onto a plain SABR model is the rapid calibration by means of Hagan's formula (West (2005)). The projection formula described in Section 3 can also be inverted numerically to retain the constant-parameter SABR–HW parameters, $\sigma(0)$, γ, $\rho_{x,\sigma}$, $\rho_{x,r}$, λ, $\theta(t)$ and η, from those of a plain SABR model, $\hat{\sigma}(0)$, $\hat{\gamma}$ and $\hat{\rho}_{F,\sigma}$. Since two parameters β and $\rho_{x,\sigma}$ control the skewness of the implied volatility curve, one of them (parameter β in our case) is fixed *a priori*, as in Rebonato (2009).

We briefly recall the calibration of the plain SABR model, in which different values are prescribed for β, such as $\beta \in \{0.25, 0.5, 0.75, 1\}$ (see, for example, West (2005) and Rebonato (2009)). By numerical experiments we observe that different combinations of β and ρ give rise to parameter fits of very similar quality. This is especially true for short-maturity implied volatilities (see Figure 2 on page 17). The specific β that gives the best fit for both short and long maturities will be determined in the second calibration stage.

Parameter $\hat{\sigma}(0)$ is determined with the help of the at-the money (ATM) implied volatility. West (2005) shows that when the forward in the plain SABR model is equal to the strike price, $F = K$, the ATM implied volatility in Hagan's formula simplifies

to:

$$\sigma_{\text{ATM}} = (F^{1-\beta})^{-1} \times \hat{\sigma}(0)\left(1 + \left(\frac{(1-\beta)^2}{24}\frac{\hat{\sigma}(0)^2}{F^{2-2\beta}} + \frac{1}{4}\frac{\hat{\rho}_{F,\sigma}\beta\hat{\gamma}\hat{\sigma}(0)}{F^{1-\beta}} + \frac{2-3\hat{\rho}_{F,\sigma}^2}{24}\hat{\gamma}^2\right)T\right)$$

Change to display style OK? Risk Journals prefers to avoid using stacked fractions within stacked fractions.

This equation is inverted, as in West (2005), to calculate $\hat{\sigma}(0)$ as a root of the cubic equation:

$$\frac{(1-\beta)^2 T}{24F^{2-2\beta}}\hat{\sigma}(0)^3 + \frac{\hat{\rho}_{F,\sigma}\beta\hat{\gamma}T}{4F^{1-\beta}}\hat{\sigma}(0)^2 + \left(1 + \frac{2-3\hat{\rho}_{F,\sigma}^2}{24}\hat{\gamma}^2 T\right)\hat{\sigma}(0) - \sigma_{\text{ATM}}F^{1-\beta} = 0$$

For typical parameters, the above cubic equation has only one real-valued root (and two imaginary roots), but it is in general possible to have three real-valued roots. In such cases, the smallest positive root should be chosen (West (2005)). As a result of the cubic equation formulation, $\hat{\sigma}(0)$ is not a free variable anymore, but a function of the parameters $\hat{\rho}_{F,\sigma}$, $\hat{\gamma}$ and the market ATM implied volatility, σ_{ATM}. Subsequently, the calibration only has to be performed over the parameters $\hat{\rho}_{F,\sigma}$ and $\hat{\gamma}$, which can be done very efficiently.

From Equation (3.7), we know that:

$$\hat{\sigma}(0) = \frac{\sigma(0)}{T}\int_0^T v(s)\,\mathrm{d}s$$

where:

$$v(t) \approx \sqrt{\xi^{2(\beta-1)}(t) + \left(\frac{\eta B_r(t,T)}{\sigma(0)F^{\beta-1}(0)}\right)^2 - 2\rho_{x,r}\frac{\eta B_r(t,T)\xi^{\beta-1}(t)}{\sigma(0)F^{\beta-1}(0)}}$$

So, $\hat{\sigma}(0)$ itself is a function of $\sigma(0)$ (since the other parameters and functions in the equation for $v(t)$ have been determined in earlier steps). Applying a numerical root-finding routine provides us with a value for $\sigma(0)$. Similarly, we can find the solution for γ via Equations (3.8) and (3.5). With formula (2.16) we can rewrite the correlation $\hat{\rho}_{F,\sigma}$ as a function of $\rho_{x,\sigma}$. The numerical inversion of this expression gives us parameter $\rho_{x,\sigma}$. After this, all parameters of the SABR–HW system, $\sigma(0)$, γ, $\rho_{x,\sigma}$, $\rho_{x,r}$, λ, $\theta(t)$ and η, have been determined.

This stage of the calibration procedure is highly efficient as most of the evaluations are based on analytic expressions. The numerical root-finding procedure is used four times. The overall computational time is less than one second.

We present an example of the calibration of the parameters $\hat{\rho}_{F,\sigma}$ and $\hat{\gamma}$ to the five-year and fifteen-year DAX options with equally spaced strike values from 40% to 220% with 10% intervals (nineteen strikes in total), from September 27, 2010, based

Change OK?

TABLE 4 Calibrated parameters for the SABR (parameters with a hat) and SABR–HW (without a hat) models, to five-year and fifteen-year DAX options.

(a) $T = 5$ years

β	$\hat{\sigma}(0)$	$\hat{\gamma}$	$\hat{\rho}_{F,\sigma}$	Squared sum error	$\sigma(0)$	γ	$\rho_{F,\sigma}$
0.25	0.280	0.033	−1	3.82×10^{-4}	0.274	0.034	−0.964
0.5	0.276	0.113	−0.955	2.22×10^{-6}	0.269	0.117	−0.922
0.75	0.274	0.224	−0.824	9.94^{-6}	0.267	0.234	−0.800
1	0.279	0.319	−0.840	6.53^{-5}	0.272	0.333	−0.815

(a) $T = 15$ years

β	$\hat{\sigma}(0)$	$\hat{\gamma}$	$\hat{\rho}_{F,\sigma}$	Squared sum error	$\sigma(0)$	γ	$\rho_{F,\sigma}$
0.25	0.402	10^{-7}	−1	1.16×10^{-4}	0.369	1.0134×10^{-7}	−0.942
0.5	0.356	10^{-7}	−1	1.75×10^{-5}	0.327	1.0134×10^{-7}	−0.937
0.75	0.328	10^{-7}	−1	1.73×10^{-6}	0.301	1.0134×10^{-7}	−0.933
1	0.319	0.033	−0.758	3.89×10^{-8}	0.292	0.033	−0.722

on the procedure described. The calibration has been performed with four sets of parameters with different *a priori* chosen values for β. The whole procedure (for the four sets of parameters) takes approximately 0.3 seconds computational time on a desktop computer. The resulting parameters and squared sum errors are presented in Table 4. Different values of β result in a qualitatively similar fit to the market implied volatilities. The fit of the SABR–HW model (based on Hagan's formula) to the market implied volatilities is presented in Figure 2 on the facing page.

The last three columns in Table 4 (the $\sigma(0)$, γ and $\rho_{F,\sigma}$ columns) present the constant-parameter SABR–HW model parameters obtained from the inversion of the projection formulas described above. We see that constant parameters produce a very good fit for individual maturity times, but that the resulting parameters differ for different maturities.

4.2 Stage II: calibration of the dynamic SABR–HW model

Calibration of the constant-parameter SABR–HW model results in a series of independent implied volatility smiles across several maturities, which do not show coherent dynamics over a longer time period. We therefore describe the calibration of the time-

FIGURE 2 Calibration results for the SABR model with different *a priori* chosen β parameters to the implied volatilities of (a) five-year and (b) fifteen-year maturity.

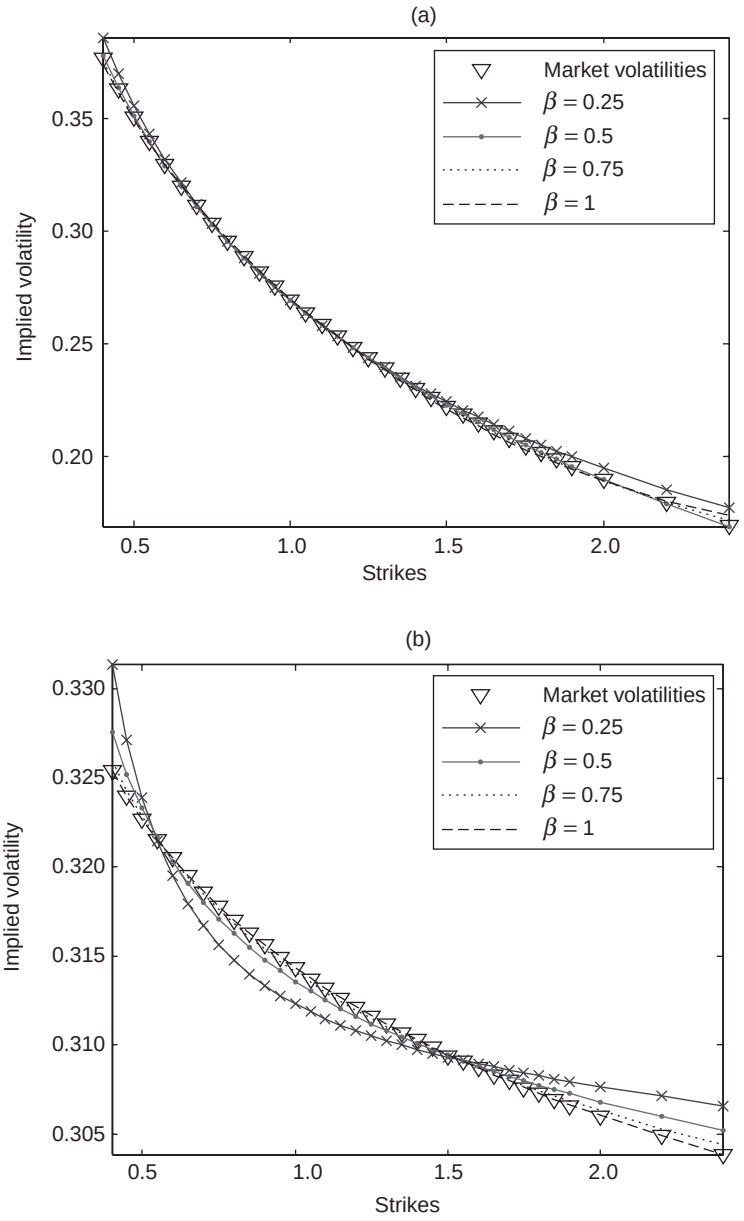

dependent functions in the dynamic SABR–HW system (2.1), and start with the β and $\rho_{x,\sigma}$ parameters for this dynamic SDE system. The value for β that fits optimally for all maturities (eg, the optimal value from Table 4 on page 16) is chosen, and we simultaneously average the calibrated correlation parameters, $\rho_{x,\sigma}$, for the corresponding β value, over the different maturity times.

Then the time-dependent function $h(t)$ in system (2.1) is parameterized in the form proposed by Rebonato (2006):

$$h(t) = (a_1 + b_1 t)\mathrm{e}^{-c_1 t} + d_1 \tag{4.1}$$

The parameters a_1, b_1, c_1 and d_1 are determined, as in Rebonato and White (2007), by solving the following system of equations for all maturities T_i included in the calibration instruments:

$$\gamma^{T_i} - \frac{1}{\sigma^{T_i}(0)T_i}\sqrt{2\int_0^T g(t)^2\hat{h}(t)^2 t\,\mathrm{d}t} = 0 \tag{4.2}$$

Here, superscript T_i denotes the maturity for which the parameter is determined, and $\hat{h}(t)$ denotes the mean value of $h(\cdot)$ up to time t, ie:

$$\hat{h}(t) = \sqrt{\frac{1}{t}\int_0^t (h(s))^2\,\mathrm{d}s}$$

Equation (4.2) can the product of $g(\cdot)$ and $h(\cdot)$ (the squared difference will be even more complicated). In this case, be best dealt with using a numerical root-finding technique.

It seems like something has gone wrong in this sentence and the next one. Can you reword both for clarity please?

For the time-dependent function $g(t)$, a common parametrization is:

$$g(t) = (a_2 + b_2 t)\mathrm{e}^{-c_2 t} + d_2 \tag{4.3}$$

which can also be found in Brigo and Mercurio (2007) or Rebonato (2002). More precisely, we obtain a_2, b_2, c_2 and d_2 by a minimization of the sum of squared errors:

$$\min_{a_2,b_2,c_2,d_2}\sum_{i=1}^{M}[\sigma^{T_i}(0) - \hat{g}(t)]^2, \quad \hat{g}(t) = \sqrt{\frac{1}{T_i}\int_0^{T_i}[(a_2 + b_2 t)\mathrm{e}^{-c_2 t} + d_2]^2\,\mathrm{d}t}$$

with M the number of option maturity times.

The time-dependent functions $g(t)$ (4.3) and $h(t)$ (4.1) are then fitted for all maturity times to the parameters $\sigma(0)$ and γ obtained from the constant-parameter calibration. The resulting parameters are presented in Table 5 on the facing page. The functions are illustrated in Figure 3 on the facing page. We fix the parameter d_1 in function $h(t)$ to the value of the volatility-of-volatility parameter of the longest maturity to prevent it from attaining negative values.

TABLE 5 Parameters a_1, b_1, c_1 and d_1 for calibrated function $h(\cdot)$ and a_2, b_2, c_2 and d_2 for function $g(\cdot)$ for DAX option implied volatilities on September 27, 2010.

Parameters	$a_1(a_2)$	$b_1(b_2)$	$c_1(c_2)$	$d_1(d_2)$
$h(\cdot)$	0.5928	−0.1943	1.1936	0.1080
$g(\cdot)$	0.0949	0.0673	0.1297	0.0858

FIGURE 3 The calibrated $g(\cdot)$ and $h(\cdot)$ functions.

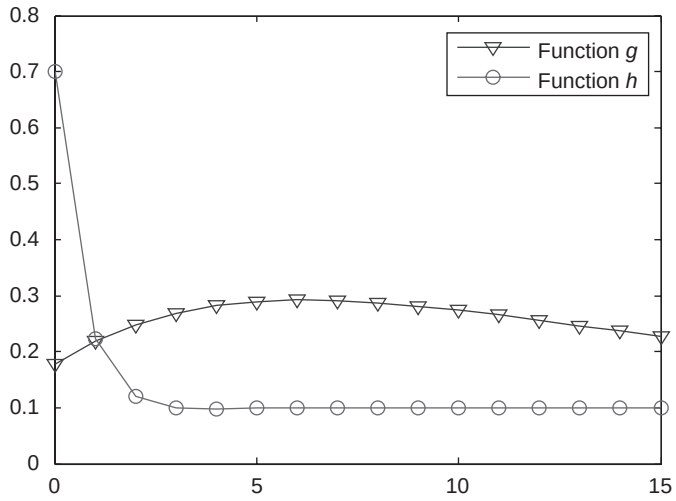

REMARK 4.1 The Hagan implied volatility function (Hagan *et al* (2002)) is based on asymptotic expansions that have a limited range of applicability. The formula is not exact, for deep-out-of-the-money strikes, for example, particularly for strikes close to zero, and for long maturities. Thus, the model dynamics that are simulated by the Monte Carlo technique may not resemble the parameters determined during calibration. In the next section we propose a method to eliminate such approximation-error-induced calibration error.

4.3 The weighted Monte Carlo technique

We employ a nonparametric approach to further improve the SABR–HW model calibration. The general idea is to perturb the weights of the individual Monte Carlo paths so that calibration instruments such as options, forwards and bonds resemble the corresponding market prices more closely.

Changes to sentence OK?

Most often, one deals with ordinary Monte Carlo methods, which are governed by the fact that the same weight (ie, $1/N$, with N the total number of paths) is assigned to each sampled path. For a claim with a payoff, ϕ, the derivative value at $t = 0$ is then determined as:

$$V = \frac{1}{N} \sum_{i=1}^{N} \phi(\omega_i)$$

where ω_i denotes the ith Monte Carlo path.

In addition, weighted Monte Carlo methods have been developed by Avellaneda (1998) and Avellaneda *et al* (1997, 2000), where different "probability" weights, $p_1, p_2, \ldots, p_N$, are assigned to the individual Monte Carlo paths. The value of the claim then reads:

$$V = \sum_{i=1}^{N} \phi(\omega_i) p_i$$

The weights are determined so that the model values of the calibration instruments match well with the market prices and these weights should be kept as close as possible to the uniform weights ($p_i = 1/N$).

We denote the market prices of M calibration instruments by $C_1, \ldots, C_M$ and represent the present values of the derivative products of the jth calibration instruments along path ω_i by ϕ_{ij}, $j = 1, 2, \ldots, M$. The first index represents the Monte Carlo path number and the expression is short notation for $\phi_{ij} \equiv \phi_j(\omega_i)$.

The path weights, or probabilities $p = (p_1, p_2, \ldots, p_N)$, have to be determined, so that:

$$\sum_{i=1}^{N} p_i \phi_{ij} = C_j \tag{4.4}$$

or so that the difference between the left-hand side and the right-hand side is minimized. A criterion (which is adopted here) to find these weights is the minimization of the relative entropy of a nonuniformly sampled probability with respect to a uniform distribution.

The concept of relative entropy is not new in computational finance. Buchen and Kelly (1996) proposed the use of the minimization of relative entropy to determine the Arrow–Debreu probability in a single-period model. This method was generalized to dynamic models by Avellaneda (1998) and Avellaneda *et al* (1997, 2000).

Based on two sets of N discrete probabilities, $p = (p_1, p_2, \ldots, p_N)$ and $q = (q_1, q_2, \ldots, q_N)$, the relative entropy of p, with respect to q, is defined as:

$$D(p \,\|\, q) := \sum_{i=1}^{N} p_i \log\left(\frac{p_i}{q_i}\right)$$

In the case of a Monte Carlo simulation, in which $q_i = 1/N$, $\forall i$, we have:

$$D(p \,\|\, q) = \sum_{i=1}^{N} p_i \log(p_i) - \sum_{i=1}^{N} p_i \log\left(\frac{1}{N}\right)$$
$$= \sum_{i=1}^{N} p_i \log(p_i) + \log(N) \tag{4.5}$$

where we used that $\sum_{i=1}^{N} p_i = 1$. The objective is to minimize Equation (4.5) under the linear constraints implied by Equation (4.4). A true advantage of the relative entropy objective function lies in the fact that Equation (4.5) is convex in all p_i-values.[3] It is well-known in optimization theory that the above minimization problem has a unique global minimum solution, if it exists (Boyd and Vandenberghe (2004)), and that the Lagrange multiplier technique determines this solution in an efficient way.

Here we present results obtained by the weighted Monte Carlo approach (see Avellaneda *et al* (2000) for details). Changes to sentence OK?

In probability and in information theory, the relative entropy, or Kullback–Leibler divergence, is a so-called nonsymmetric measure of the difference between two probability distributions p and q. The relative entropy measure is an indication of the difference between any two models. In our case, it quantifies the consistency, or inconsistency, between the calibrated true SABR–HW model and the SABR–HW$_1$ model obtained from the first stages of the calibration. The relative entropy distance is defined as $D(p \,\|\, u)$ in Equation (4.5), in which u denotes a uniform probability of N Monte Carlo samples. Since the term $\sum_{i=1}^{N} p_i \log p_i$ in Equation (4.5) is negative (as $p_i \leqslant 1$), $D(p \,\|\, u) \in [0, \log N]$. The minimum value, $D(p \,\|\, u) = 0$, corresponds to $p_i = 1/N$, ie, the calibrated vector p equals the prior u. The maximum value, $D(p \,\|\, u) = \log N$, is realized when the probability is concentrated at a single path, ie, $p_i = 1$. Consider a probability distribution that is uniformly distributed on a subset of paths of size N^α, with $0 < \alpha < 1$. Substitution of the corresponding probabilities gives (Avellaneda *et al* (2000)): Change OK?

$$D(p \,\|\, u) = \log N + \log\left(\frac{1}{N^\alpha}\right)$$
$$= (1 - \alpha) \log N \tag{4.6}$$

[3] It is straightforward to show that:

$$\frac{\partial D(p \,\|\, q)}{\partial p_i} = \log(p_i) + 1, \qquad \frac{\partial^2 D(p \,\|\, q)}{\partial p_i^2} = \frac{1}{p_i} > 0$$

using $\sum_{i=1}^{N^\alpha} N^{-\alpha} = 1$.

The relative entropy distance can therefore be connected to the effective number of paths, N^α, supported by the prior probability measure. The effective number of paths can be obtained, as $\alpha = 1 - D(p \,\|\, u)/\log N$, with $D(p \,\|\, u)/\log N \in [0, 1]$. If $D(p \,\|\, u)/\log N \ll 1$, the number of significant paths is close to N, whereas $D(p \,\|\, u)/\log N \approx 1$ is connected to a measure with "thin support" (Avellaneda *et al* (2000)). Thin support implies that a large number of paths are discarded in the calibration, which is inefficient from a computational point of view.

REMARK 4.2 It is possible that a solution to the minimum entropy problem does not exist, when the initial problem parameters result in prices of the calibration instruments that are very different from the market prices. In such a case, the minimum entropy algorithm will not work, but one may use a quadratic difference function:

$$D^Q(p \,\|\, u) = \sum_{i=1}^{N} \left(p_i - \frac{1}{N} \right)^2$$

instead of the relative entropy distance, which guarantees a solution (see Avellaneda and Jäckel (2010)).

4.4 Stage III: calibration by weighted Monte Carlo method

Here, we use the DAX one-year, five-year and ten-year implied volatilities from September 27, 2010 with equally spaced strike values from 40% to 220% with 10% intervals. After the computation of the weighted Monte Carlo weights for these financial derivatives, the weighted Monte Carlo method perfectly replicates the prices of these calibration instruments (see Figure 4 on the facing page). Word added – OK?

We plot the resulting weighted Monte Carlo weights in a log scale in part (a) of Figure 5 on page 24. The weights seem to be randomly distributed around their mean value of 10^{-5}. Certain paths are given a small weight, which means that these paths are effectively discarded. The histogram of the weights in part (b) of Figure 5 on page 24 indicates that only a small fraction of the weights is in the left tail and most of them are distributed around the mean. The resulting effective number of paths, obtained by Equation (4.6), is 9.3517×10^4, so that 6.48% of the paths are discarded. This is efficient from a computational point of view, given the excellent weighted Monte Carlo calibration results.

REMARK 4.3 The weighted Monte Carlo calibration procedure is highly efficient, but it is nonparametric, and this may hamper its practical application. If either the model or the model parameters are not carefully chosen, too many paths will be discarded and the weighted Monte Carlo efficiency would be lost. However, the weighted Monte Carlo technique can also be used as an *a posteriori* check of the quality of

FIGURE 4 (a) Five-year and (b) ten-year option implied volatilities produced by weighted Monte Carlo compared with the input market implied volatilities.

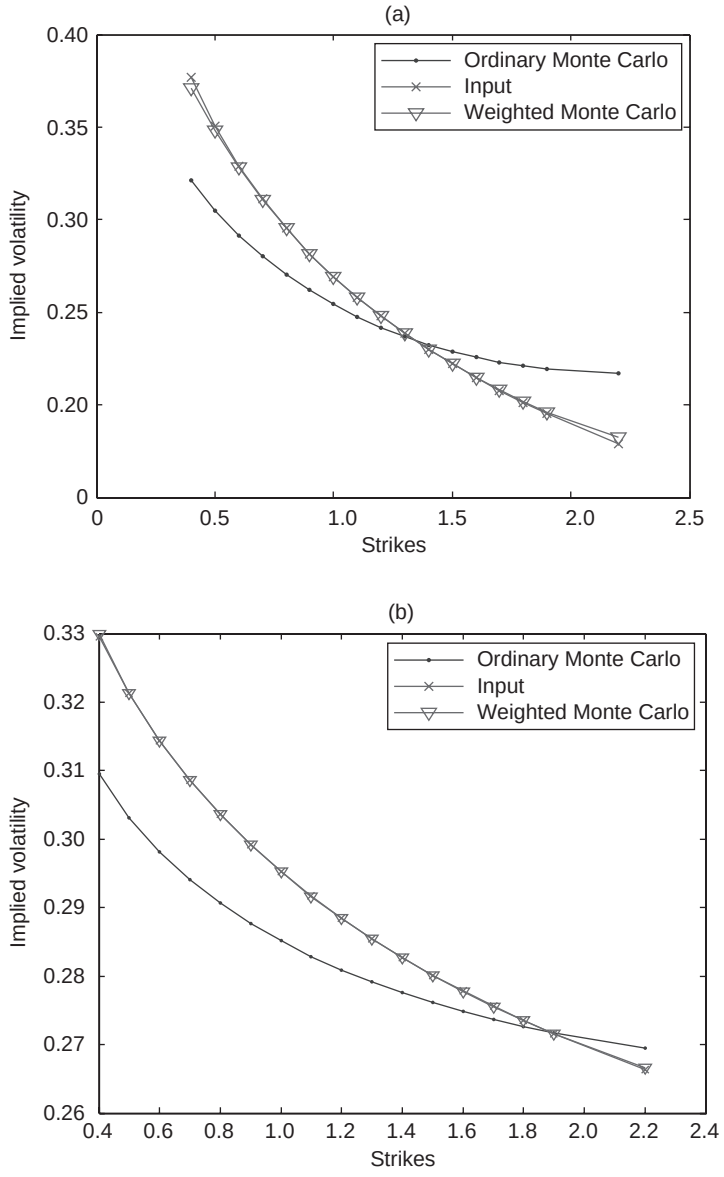

FIGURE 5 (a) Cloud plot of the 10^5 weights (in log scale) of the weighted Monte Carlo paths obtained after the calibration, and (b) histogram of the values of the calibrated weights in log scale.

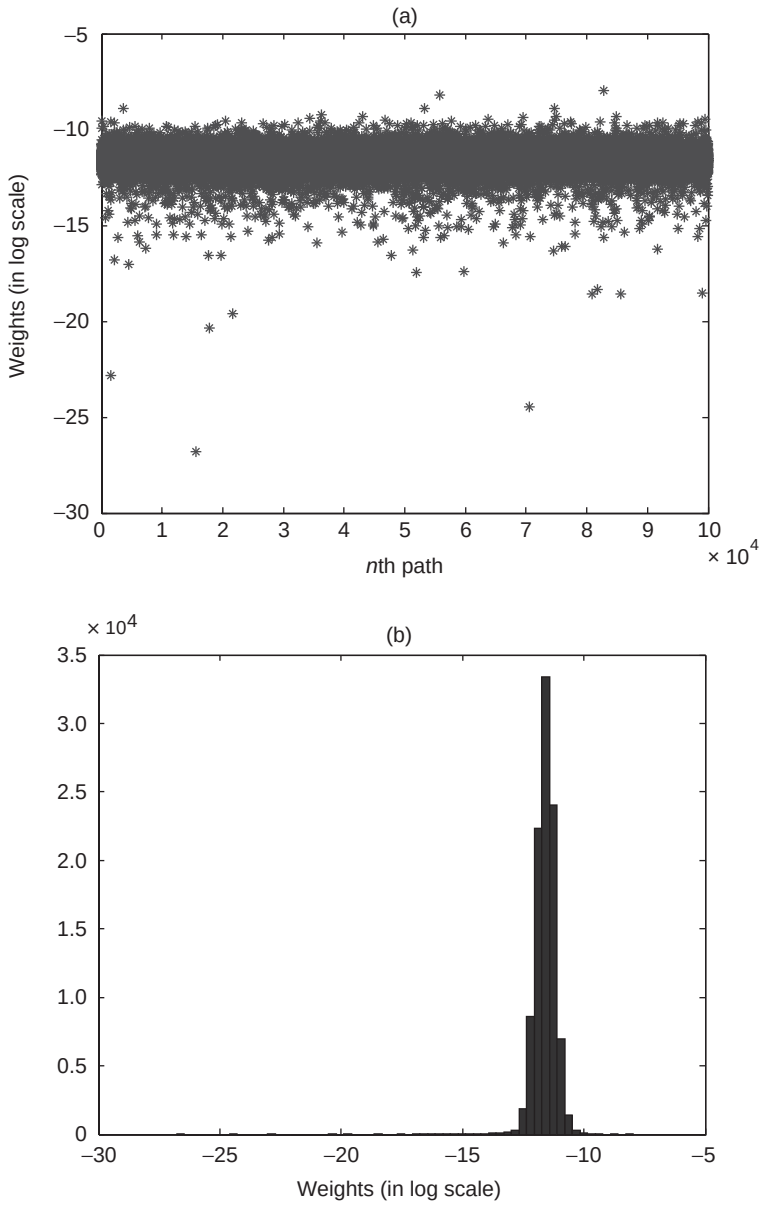

parameters obtained from calibration to Hagan's implied volatilities. If, after the first stage of calibration, the Monte Carlo weights are equally distributed and very close to $p_i = 1/N$, this may give confidence in the quality of the first stage of calibration.

5 PRICING OPTIONS UNDER THE SABR–HULL–WHITE MODEL

On the basis of the calibrated SABR–HW model, we are now ready to apply a (weighted) Monte Carlo simulation for the pricing and hedging of exotic derivatives. We present an advanced time-stepping scheme for the Monte Carlo simulation, leading to a low-bias Monte Carlo simulation. This scheme is also accurate when only a few time steps are employed. It has also been used within the weighted Monte Carlo part of the calibration procedure described earlier.

5.1 Low-bias time-discrete scheme

Applications of the SABR–HW model include the pricing of long-maturity equity options, equity-linked structured notes (like cliquet options) and equity-linked hybrid derivatives. Structured products usually have a long time horizon and a complicated payoff. It is difficult to find analytic approximations for these product prices, and often one has to rely on Monte Carlo methods to obtain prices and hedge ratios (eg, price sensitivities and Greeks). If we apply an Euler discretization scheme to the SABR–HW system, the discrete bias has to be analyzed with care. For example, in Equation (2.11) the drift term is stochastic and driven by two stochastic factors. In this case, an Euler approximation for the drift term:

$$\int_0^\Delta r(s)S(s)\,\mathrm{d}s \approx r(0)S(0)\Delta$$

is biased in general and a large number of time steps is required to reach an acceptable level of accuracy. An Euler scheme may therefore be inefficient for pricing long-term equity-linked structured products.

Here we adapt the low-bias Monte Carlo scheme proposed for the SABR model in Chen *et al* (2012) to discretize the SABR–HW model. The approach is to map the asset price process onto a square-root process using a series of spatial and time transformations. In the Monte Carlo simulation we will draw samples from the analytic distribution function of the square-root process (ie, the noncentral chi-squared distribution), as described in Section 5.2 (and described in full detail in Chen *et al* (2012)).

The SABR model considered in Chen *et al* (2012) was developed for a system without a drift term. A stochastic interest rate can be incorporated by a technique described in Goldenberg (1991, p. 28), which was introduced for a constant elasticity

of variance (CEV) process with drift term $rS(t)\,\mathrm{d}t$:

$$\mathrm{d}S^{(r)}(t) = rS^{(r)}(t)\,\mathrm{d}t + \sigma S^{(r)}(t)^{\beta}\,\mathrm{d}W(t) \tag{5.1}$$

We use the superscript "(r)" to distinguish the process with drift from the process without drift, which does not have a superscript. The distribution of the CEV process in (5.1) can be obtained from its sister without drift via a time change:

$$S^{(r)}(t) = \mathrm{e}^{rt} S(\tau(t)), \qquad \tau(t) = \frac{1}{2r(\beta-1)}(\mathrm{e}^{2r(\beta-1)t} - 1) \tag{5.2}$$

The validity of this transformation can easily be explained as the limit $r \to 0$ recovers the original clock, ie:

$$\lim_{r\to 0} \tau(t) = \lim_{r\to 0} \frac{1}{2r(\beta-1)}(\mathrm{e}^{2r(\beta-1)t} - 1) = t$$

This result is not restricted to the constant interest rate case. The time transformation in (5.2) also applies to stochastic interest rates (Goldenberg (1991)).

In an SDE system with a stochastic interest rate, the time transformation is different for each path, due to the randomness of the rates:

$$\tau(t,\omega) = \frac{\Delta}{2(\beta-1)\int_0^{\Delta} r(s)\,\mathrm{d}s}\left(\exp\left(2(\beta-1)\int_0^{\Delta} r(s)\,\mathrm{d}s\right) - 1\right) \tag{5.3}$$

where $\omega \in \Omega$ denotes a random scenario. Expression (5.3) suggests that the pathwise time transformation, $\tau(t,\omega)$, can be determined without all details of the interest rate path, ω, as long as we have knowledge of $\int_0^{\Delta} r(s)\,\mathrm{d}s$ at each path.

Following the arguments by Andersen (2008), we focus on the evolution of the system over a small time interval $[0,\Delta]$ and repeat the one-period Δ scheme to produce a complete time-discrete path. Note that we consider the SDE system in the τ timescale, so that the time interval for the system equals $[0,\tau(\Delta,\omega)]$.

The CEV system in (5.2) is then simulated on a timescale $\tau(t,\omega)$, induced by a stochastic interest rate. Subsequently, the result is multiplied by an exponentially integrated interest rate:

$$S(\tau(\Delta,\omega)) = S(0) + \int_0^{\tau(\Delta,\omega)} \sigma(\tau(s,\omega))S(\tau(s,\omega))^{\beta}\,\mathrm{d}W_x(\tau(s,\omega))$$

$$S^{(r)}(\Delta) = \exp\left(\int_0^{\Delta} r(s)\,\mathrm{d}s\right)S(\tau(\Delta,\omega))$$

Although a transformed model based on time changes may not be intuitively clear, it is easy to implement numerically. We just replace the uniform time interval Δ by τ, based on the computation on each interest rate path.

5.2 Discretization of the SDE system

For the SABR–HW system, we also consider the system without drift with timescale $\tau(t,\omega)$, and a low-bias Monte Carlo simulation scheme (see Section 5.2 and Chen *et al* (2012)). We then multiply the result by the exponentially integrated interest rate.

In this section we describe the low-bias time discretization scheme to simulate the plain SABR system in the interest rate-dependent timescale $\tau(t,\omega)$:

$$\begin{aligned}\mathrm{d}S(\tau) &= \sigma(\tau)S(\tau)^{\beta}\,\mathrm{d}W_x(\tau)\\ \mathrm{d}\sigma(\tau) &= \gamma\sigma(\tau)\,\mathrm{d}W_\sigma(\tau)\end{aligned}$$

where we denote $\tau(t,\omega)$ simply by τ. Let us consider a system of three Brownian motions correlated with the following correlation matrix:

$$\begin{pmatrix} 1 & \rho_{x,\sigma} & \rho_{x,r}\\ \rho_{x,\sigma} & 1 & 0\\ \rho_{x,r} & 0 & 1\end{pmatrix}$$

Applying the Cholesky decomposition gives:

$$\left.\begin{aligned}\mathrm{d}W_x(\tau) &= \rho_{x,\sigma}\,\mathrm{d}W_1(\tau) + \rho_{x,r}\,\mathrm{d}W_2(\tau) + \sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\,\mathrm{d}U(\tau)\\ \mathrm{d}W_\sigma(\tau) &\equiv \mathrm{d}W_1(\tau)\\ \mathrm{d}W_r(\tau) &\equiv \mathrm{d}W_2(\tau)\end{aligned}\right\} \tag{5.4}$$

Second and third equation numbers deleted and brace added so that first number applies to all three lines as it appeared from the coding that this was the intention – OK or revert to original?

where the Brownian motions $W_1(\tau)$, $W_2(\tau)$ and $U(\tau)$ are mutually independent.

Based on an argument from Schroder (1989), we consider the invertible transformation of variables $X(\tau) = S(\tau)^{1-\beta}/(1-\beta)$, $\beta \neq 1$, such that:

$$\mathrm{d}X(\tau) = \sigma(\tau)\,\mathrm{d}W_x(\tau) - \frac{\beta\sigma(\tau)^2}{(2-2\beta)X(\tau)}\,\mathrm{d}\tau \tag{5.5}$$

We substitute Equation (5.4) into Equation (5.5) and integrate from 0 to $\tau(\Delta,\omega)$, which gives:

$$\begin{aligned}&X(\tau(\Delta,\omega))\\ &\quad= X(0,\omega) + \rho_{x,\sigma}\int_0^{\tau(\Delta,\omega)}\sigma(s)\,\mathrm{d}W_1(s) + \rho_{x,r}\int_0^{\tau(\Delta,\omega)}\sigma(s)\,\mathrm{d}W_2(s)\\ &\qquad- \int_0^{\tau(\Delta,\omega)}\frac{\beta\sigma(s)^2}{(2-2\beta)X(s)}\,\mathrm{d}s + \sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\int_0^{\tau(\Delta,\omega)}\sigma(s)\,\mathrm{d}U(s)\end{aligned}$$

In Chen *et al* (2012) it was shown that:

$$\int_0^{\tau(\Delta,\omega)}\sigma(s)\,\mathrm{d}W_1(s) = \frac{\sigma(\tau(\Delta,\omega)) - \sigma(0)}{\gamma}$$

It is also not difficult to show that:

$$\int_0^{\tau(\Delta,\omega)} \sigma(s)\,\mathrm{d}W_2(s)$$

is a Gaussian integral with variance $\int_0^{\tau(\Delta,\omega)} \sigma(s)^2\,\mathrm{d}s$, because of the independence of W_1 and W_2. Hence, we can replace the Gaussian integral:

$$\int_0^{\tau(\Delta,\omega)} \sigma(s)\,\mathrm{d}W_2(s) = \zeta(\tau(\Delta,\omega))W_2(\tau(\Delta,\omega)) \tag{5.6}$$

where we have defined:

Stacked fraction here OK?

$$\zeta(\tau(\Delta,\omega)) := \sqrt{\int_0^{\tau(\Delta,\omega)} \frac{\sigma(s)^2\,\mathrm{d}s}{\tau(\Delta,\omega)}}$$

Based on these results, we can sample the SABR system without drift term and use the conditional scheme proposed in Chen *et al* (2012), conditional on the terminal volatility, the integrated volatility and the realization of W_2:

$$\begin{aligned} &X(\tau(\Delta,\omega)) \\ &\quad = X(0) + \frac{\rho_{x,\sigma}}{\gamma}\{\sigma(\tau(\Delta,\omega)) - \sigma(0)\} + \rho_{x,r}\zeta(\tau(\Delta,\omega))W_2(\tau(\Delta,\omega)) \\ &\qquad - \int_0^{\tau(\Delta,\omega)} \frac{\beta\sigma(s)^2}{(2-2\beta)X(s)}\,\mathrm{d}s + \sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\int_0^{\tau(\Delta,\omega)} \sigma(s)\,\mathrm{d}U(s) \end{aligned}$$

where we have used (5.6).

Conditional on the volatility, the integrated variance and Brownian motion W_2, process X is a shifted Bessel process, $\tilde{X}$, with dynamics:

$$\begin{aligned} \mathrm{d}\tilde{X}(\tau) &:= \sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\sigma(\tau)\,\mathrm{d}U(\tau) - \frac{\beta\sigma(\tau)^2}{(2-2\beta)\tilde{X}(\tau)}\,\mathrm{d}\tau \\ \tilde{X}(0) &= X(0) + \frac{\rho_{x,\sigma}}{\gamma}\{\sigma(\tau(\Delta)) - \sigma(0)\} + \rho_{x,r}\zeta(\tau(\Delta))W_2(\tau(\Delta)) \end{aligned}$$

We define another change of variables, $Y(\tau) := \tilde{X}(\tau)^2$, and apply Ito's lemma:

$$\begin{aligned} \mathrm{d}Y(\tau) &= 2\tilde{X}(\tau)\,\mathrm{d}\tilde{X}(\tau) + \mathrm{d}\tilde{X}(\tau)^2 \\ &= 2\tilde{X}(\tau)\left(\sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\sigma(\tau)\,\mathrm{d}U(\tau) - \frac{\beta\sigma(\tau)^2}{(2-2\beta)\tilde{X}(\tau)}\,\mathrm{d}\tau\right) \\ &\qquad + (1-\rho_{x,\sigma}^2-\rho_{x,r}^2)\sigma(\tau)^2\,\mathrm{d}\tau \\ &= 2\sqrt{Y(\tau)}\sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\sigma(\tau)\,\mathrm{d}U(\tau) \\ &\qquad + \left(\frac{1-2\beta-(\rho_{x,\sigma}^2+\rho_{x,r}^2)(1-\beta)}{(1-\beta)(1-\rho_{x,\sigma}^2-\rho_{x,r}^2)}\right)(1-\rho_{x,\sigma}^2-\rho_{x,r}^2)\sigma(\tau)^2\,\mathrm{d}\tau \end{aligned} \tag{5.7}$$

Let us define the time change:

$$\nu(\tau(\Delta,\omega)) = (1-\rho_{x,\sigma}^2-\rho_{x,r}^2)\int_0^{\tau(\Delta,\omega)}\sigma(s)^2\,\mathrm{d}s$$

Due to the independence of the Brownian motion, U, and the volatility process, a Brownian motion under clock $\nu(\tau(\cdot,\cdot))$ has the same distribution as:

$$\sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\int_0^{\tau(\Delta,\omega)}\sigma(s)\,\mathrm{d}U(s)$$

ie:

$$\begin{aligned}U(\nu(\tau(\Delta,\omega))) &\equiv \int_0^{\nu(\tau(\Delta,\omega))}\mathrm{d}U(s)\\ &= \sqrt{1-\rho_{x,\sigma}^2-\rho_{x,r}^2}\int_0^{\tau(\Delta,\omega)}\sigma(s)\,\mathrm{d}U(s)\end{aligned}$$

We substitute this time change into Equation (5.7), and obtain:

$$\mathrm{d}Y(\nu(\tau)) = 2\sqrt{Y(\nu(\tau))}\,\mathrm{d}U(\nu(\tau)) + \left(\frac{1-2\beta-(\rho_{x,\sigma}^2+\rho_{x,r}^2)(1-\beta)}{(1-\beta)(1-\rho_{x,\sigma}^2-\rho_{x,r}^2)}\right)\mathrm{d}\nu(\tau) \tag{5.8}$$

which gives us a time-changed squared Bessel process of dimension:

$$\delta = \frac{1-2\beta-(\rho_{x,\sigma}^2+\rho_{x,r}^2)(1-\beta)}{(1-\beta)(1-\rho_{x,\sigma}^2-\rho_{x,r}^2)}$$

with initial value $Y(0)=\tilde{X}(0)^2$. The time change depends on the interest rate path via $\tau(\cdot,\cdot)$ and depends on the volatility path as a result of the time-change function $\nu(\cdot)$. Each change of time is conditioned on specific path information of the volatility and the interest rate paths.

The stochastic volatility induced time change, $\nu(\cdot)$, can be computed by an asymptotic expansion (see Chen *et al* (2012, Section 3.4)). In order to determine the change of timescale $\tau(\cdot,\cdot)$, we have to compute the integral $\int_0^\Delta r(s)\,\mathrm{d}s$ in Equation (5.3). There are several ways to approximate this integral. A straightforward approach is the following discrete approximation:

$$\int_0^\Delta r(s)\,\mathrm{d}s \approx \Delta[w_1 r(0) + w_2 r(\Delta)]$$

The constants w_1 and w_2 can be chosen in different ways. For example, in an Euler scheme we have $w_1=1$ and $w_2=0$. A central discretization employs $w_1=w_2=\frac{1}{2}$. This scheme is computationally efficient and sufficiently accurate if the underlying

stochastic process is slowly varying. This is the case for these interest rate processes. The calibrated interest rate process usually includes a volatility parameter $\eta < 1\%$.

The cumulative distribution function of process Y can be obtained by the properties of squared Bessel processes. By a mapping:

$$h\colon\ s \to \frac{s^{2-2\beta}}{(1-\beta)^2}, \quad s \geqslant 0$$

with its inverse:

$$h^{-1}\colon\ y \to \frac{y^{2-2\beta}}{(1-\beta)^2}, \quad y \geqslant 0$$

we can define:

$$S(\tau(\Delta,\omega)) = h(Y(\nu(\tau(\Delta,\omega)))), \qquad Y(0) = h^{-1}(S(0)) = \frac{S(0)^{2(1-\beta)}}{(1-\beta)^2}$$

We now have the following proposition.

PROPOSITION 5.1 (Cumulative distribution for the conditional SABR process) *The cumulative distribution for $S(\tau(\Delta,\omega))$, conditional on:*

$$\sigma(\tau(\Delta,\omega)) \quad \textit{and} \quad \int_0^{\tau(\Delta,\omega)} \sigma_s^2 \,\mathrm{d}s$$

with an absorbing boundary at $S = 0$ reads:

$$\Pr[S(\tau(\Delta,\omega)) \leqslant x \mid S(0)] = 1 - \chi^2(a; b, c) \tag{5.9}$$

where:

$$\left.\begin{aligned} a &= \frac{1}{\nu(\tau(\Delta,\omega))}\left[\frac{S(0)^{1-\beta}}{(1-\beta)} + \frac{\rho}{\alpha}\{\sigma(\tau(\Delta,\omega)) - \sigma(0)\} \right.\right. \\ &\qquad \left.\left. + \rho_{x,r}\zeta(\tau(\Delta,\omega))W_2(\tau(\Delta,\omega))\right]^2 \\ b &= 2 - \frac{1 - 2\beta - (\rho_{x,\sigma}^2 + \rho_{x,r}^2)(1-\beta)}{(1-\beta)(1-\rho_{x,\sigma}^2 - \rho_{x,r}^2)} \\ c &= \frac{S(\tau(\Delta,\omega))^{2(1-\beta)}}{(1-\beta)^2\nu(\tau(\Delta,\omega))} \\ \nu(\tau(\Delta,\omega)) &= (1-\rho^2)\int_0^{\tau(\Delta,\omega)} \sigma(s)^2 \,\mathrm{d}s \end{aligned}\right\} \tag{5.10}$$

and $\chi^2(x;\delta,\lambda)$ is the noncentral chi-squared cumulative distribution function for random variable x with noncentrality parameter λ and degree of freedom δ.

PROOF The proof proceeds along the same lines as the work in Islah (2009) and Chen *et al* (2012). □

Summary of the algorithm

The algorithm to sample the SABR–HW system for a time interval $[0, \Delta]$ involves the following steps.

(1) Draw samples of $r(\Delta)$ from a normal distribution with mean and variance defined by Equation (2.8) and (2.9), respectively.

(2) Apply a drift interpolation:

$$\int_0^\Delta r(s)\,\mathrm{d}s \approx \Delta(\tfrac{1}{2}r(0) + \tfrac{1}{2}r(\Delta))$$

for the quantity in formula (5.3) for the pathwise (stochastic interest rate induced) rescaled time step $\tau(\Delta, \omega)$.

(3) Conditional on $\sigma(\tau(\Delta, \omega))$, draw samples of $\int_0^{\tau(\Delta,\omega)} \sigma(s)^2\,\mathrm{d}s$ by the method proposed in Section 3.4 of Chen *et al* (2012).

(4) Conditional on $\sigma(\tau(\Delta, \omega))$ and $\int_0^{\tau(\Delta,\omega)} \sigma(s)^2\,\mathrm{d}s$, draw samples of the dynamics without drift, S, in the timescale τ from the noncentral chi-squared distribution described in Proposition 5.1.

(5) Find the asset price process with drift, $S^{(r)}(\tau(\Delta, \omega))$, by multiplying the estimated $\int_0^\Delta r(s)\,\mathrm{d}s$ (step (2)) with $S(\tau(\Delta, \omega))$.

5.2.1 Numerical experiment with European options

We consider the pricing of European options in the SABR–HW model using the low-bias Monte Carlo method. We focus on a call option, maturing at time T with strike price K, and denote the exact option price at initial time by $C(K, 0)$. It is approximated by:

$$\hat{C}(K, 0) = P(t, T)\mathbb{E}^T[(S(T) - K)^+]$$

The '$P(t, T)$' in this display does not appear in the manuscript PDF provided but appeared in the coding – OK or changes needed?

where $\hat{C}(K, 0)$ is typically not equal to $C(K, 0)$, and we define the bias, e, of the simulation as a function of the time step, Δ, and analyze its behavior.

Word added – OK?

We will use the parameter sets 1, 2 and 3 from Table 1 on page 11 to carry out the numerical experiments for the call options with $T = 10$ and ATM strikes. As the benchmark, we apply a Monte Carlo simulation based on the Euler discretization with a large number of time steps (400 time steps per year), and we report the average of three runs as the reference prices. For sets 1, 2 and 3, the option prices obtained by the low-bias Monte Carlo method are listed in Table 6 on the next page for time steps Δ ranging from one to thirty-two steps a year.

Change to Arabic numbering style to match table style – OK?

Changes to sentence OK?

We plot the differences between the Monte Carlo estimate and the reference values for different step sizes in Figure 6 on page 33, from which we see that the low-bias

TABLE 6 Estimated ten-year ATM call option prices for cases 1, 2 and 3.

	Case 1		Case 2		Case 3	
Δ	Low-bias	Euler	Low-bias	Euler	Low-bias	Euler
1	0.4653	0.4699	0.4731	0.4771	0.5485	0.5374
$\frac{1}{2}$	0.4680	0.4697	0.4680	0.4703	0.5473	0.5382
$\frac{1}{4}$	0.4664	0.4662	0.4665	0.4691	0.5420	0.5479
$\frac{1}{8}$	0.4666	0.4664	0.4659	0.4681	0.5440	0.5451
$\frac{1}{16}$	0.4669	0.4675	0.4644	0.4676	0.5430	0.5449
$\frac{1}{32}$	0.4669	0.4673	0.4633	0.4667	0.5431	0.5432
Reference	0.4671	—	0.4634	—	0.5431	—

scheme is advantageous to the Euler scheme with respect to its low bias. By increasing the number of time steps, the low-bias scheme produces smaller errors than the Euler scheme for all parameter settings.

5.2.2 Time-dependent parameters

We can also discretize system (2.1) with the low-bias scheme by assuming that the functions can be approximated by piecewise constant functions on $[t, t+\Delta]$ with value $\bar{\gamma}$. According to the arguments of Andersen (2008) and Glasserman (2004, p. 130), we then use $\bar{\gamma} = \frac{1}{2}(\gamma(t) + \gamma(t+\Delta))$, which leads to a modification of system (2.1) that can easily be simulated by the low-bias scheme. More precisely, the volatility process in the dynamic SABR–HW system is sampled by the following formula:

$$\Sigma(t+\Delta) = \bar{g}k(t)\exp(-\tfrac{1}{2}\bar{h}^2\Delta + \bar{h}Z\sqrt{\Delta}), \quad k(0) = 1 \tag{5.11}$$

where:

$$\bar{g} = \tfrac{1}{2}(g(0,t) + g(0,t+\Delta)), \qquad \bar{h} = \tfrac{1}{2}(h(0,t) + h(0,t+\Delta))$$

As a result, the formula for the integrated variance has to be adapted as well: it has to be multiplied by the factor $\bar{g}^2$.

6 CONCLUSION

We have presented the dynamic SABR–HW model as an alternative model for pricing long-maturity equity options and equity–interest rate hybrid products. We have defined the model, introduced its building blocks and described several issues for the

FIGURE 6 The Monte Carlo error as a function of the number of time steps.

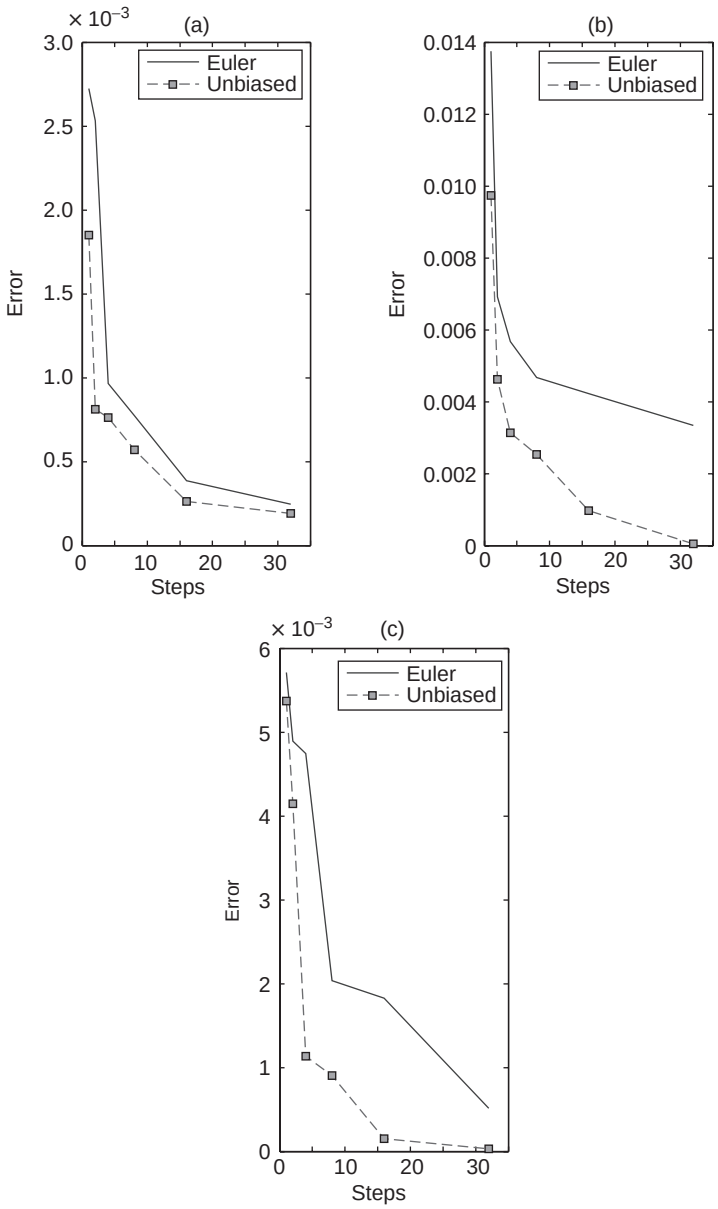

(a) Case 1. (b) Case 2. (c) Case 3. Note the different scaling of the three figures.

practical application of the SABR–HW model, such as model calibration and option pricing. At several points we have presented numeric techniques that are not commonly used by the financial industry, like the low-bias discretization scheme and a weighted Monte Carlo technique to enhance the calibration.

In particular, we have proposed an invertible projection formula for the constant-parameter SABR–HW model connecting it to the plain SABR model. The basis for this projection was a change of measure, to the T-forward measure, and a linearization. The projection formula greatly simplifies the calibration of the SABR–HW model.

The inversion of the projection formula serves as a first step in the calibration procedure, ie, it gives a rapid and fairly accurate approximation of the constant SABR–HW parameters at each maturity. Based on these parameters, we have defined time-dependent functions in the dynamic SABR–HW model that are consistent with the market implied volatilities for all maturities.

In the final calibration step, nonuniform Monte Carlo weights have been determined in such a way that the implied volatilities from the market and those generated by the Monte Carlo paths of the SABR–HW model match optimally. The overall calibration procedure is highly efficient and accurate.

Exotic contracts were then priced using the weighted Monte Carlo paths generated by a low-bias time discretization scheme of the dynamic SABR–HW model. An advantage of the low-bias scheme is that accurate Monte Carlo results can be obtained for large time steps. This is particularly useful when long-maturity options are considered.

REFERENCES

Andersen, L. (2008). Simple and efficient simulation of the Heston stochastic volatility model. *The Journal of Computational Finance* **11**(3), 1–42.

Avellaneda, M. (1998). Minimum-relative-entropy calibration of asset-pricing models. *International Journal of Theoretical and Applied Finance* **1**, 447–472.

Avellaneda, M., and Jäckel, P. (2010). Weighted Monte Carlo. In *Encyclopedia of Quantitative Finance*. Wiley.

Avellaneda, M., Friedman, C., Holmes, R., and Samperi, D. (1997). Calibrating volatility surfaces via relative-entropy minimization. *Applied Mathematical Finance* **4**, 37–64.

Avellaneda, M., Buff, R., Friedman, C., Grandechamp, N., Kruk, L., and Newman, J. (2000). Weighted Monte Carlo: a new technique for calibrating asset-pricing models. *International Journal of Theoretical and Applied Finance* **4**, 91–119.

Boyd, S., and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.

Brigo, D., and Mercurio, F. (2007). *Interest Rate Models: Theory and Practice: With Smile, Inflation and Credit*, 2nd edition. Springer.

Broadie, M., and Kaya, Ö. (2006). Exact simulation of stochastic volatility and other affine jump diffusion processes. *Operations Research* **54**, 217–231.

Buchen, P. W., and Kelly, M. (1996). The maximum entropy distribution of an asset inferred from option prices. *Journal of Financial and Quantitative Analysis* **31**, 143–159.

Chen, B., Oosterlee, C. W., and Van Weeren, S. (2010). Analytic approximation to constant maturity swap convexity correction in a multi-factor SABR model. *International Journal of Theoretical and Applied Finance* **13**(7), 1019–1046.

Chen, B., Oosterlee, C. W., and Van der Weide, J. A. M. (2012). An efficient unbiased simulation scheme for SABR stochastic volatility model. Working Paper, CWI.

Year changed to match coding of bibitem here – OK?

Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. Wiley.

Geman, H., El Karoui, N., and Rochet, J. C. (1996). Changes of numeraire, changes of probability measures and pricing of options. *Journal of Applied Probability* **32**, 443–458.

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

Goldenberg, D. (1991). A unified method for pricing options on diffusion processes. *Journal of Financial Economics* **29**, 3–34.

Grzelak, L. A., and Oosterlee, C. W. (2011). On the Heston model with stochastic interest rates. *SIAM Journal of Financial Mathematics* **2**, 255–286.

Hagan, P. S. (2002). Adjusters: turning good prices into great prices. *Wilmott Magazine* **December**, 56–59.

Hagan, P. S., Kumar, D., Lesniewski, A. S., and Woodward, D. E. (2002). Managing smile risk. *Wilmott Magazine* **November**, 84–108.

Hull, J., and White, A. (1996). Using Hull–White interest rate trees. *Journal of Derivatives* **3**(3), 26–36.

Hunter, C., and Picot, G. (2005). *Hybrid Derivatives: The Euromoney Derivatives and Risk Management Handbook*. BNP Paribas, London.

Islah, O. (2009). Solving SABR in exact form and unifying it with LIBOR market model. Working Paper, SSRN.

Kariyam, T., and Kurata, H. (2004). *Generalized Least Squares*. Wiley.

Mendoza, R., Carr, P., and Linetsky, V. (2010). Time changed Markov processes in unified credit-equity modelling. *Mathematical Finance* **20**(4), 527–569.

Piterbarg, V. (2006). Smiling hybrids. *Risk* **19**, 66–71.

Rebonato, R. (2002). *Modern Pricing of Interest-Rate Derivatives*. Princeton University Press.

Rebonato, R. (2006). A time-homogeneous, SABR-consistent extension of the LMM. *Risk* **20**, 92–97.

Rebonato, R. (2009). *The SABR/LIBOR Market Model: Pricing, Calibration and Hedging for Complex Interest-Rate Derivatives*. Wiley.

Rebonato, R., and White, R. (2007). Linking caplet and swaption prices in the LMM–SABR model. URL: www.riccardorebonato.co.uk/papers/LMMSABRSwaption.pdf.

Schroder, M. (1989). Computing the constant elasticity of variance option pricing formula. *Journal of Finance* **44**, 211–219.

West, G. (2005). Calibration of the SABR model in illiquid markets. *Applied Mathematical Finance* **12**(4), 371–385.

Wilmott, P. (2002). Cliquet options and volatility models. *Wilmott Magazine* **March**, 78–83.