Arbitrage-free construction of the swaption cube

Simon Johnson
Bereshad Nonas
Financial Engineering
Commerzbank Corporates and Markets
60 Gracechurch Street
London EC3V 0HR

5th January 2009

Abstract

In this paper we look at two areas in the interest rate options market where arbitrage could be hiding. In the first section we derive a no-arbitrage condition for swaption prices with complementary expiry dates and tenors within the swaption cube. In the second section we propose an alternative European option approximation for the widely used SABR dynamics that reduces the possibility of arbitrage for long maturities and low strikes.

A clear understanding of the arbitrage condition between vanilla options is essential for all market participants. Options market makers need to avoid the danger of being arbitraged by their customers; proprietary traders seek to identify and exploit arbitrage or near-arbitrage prices; exotics traders use models which can of course only be successfully calibrated to arbitrage-free vanilla option prices. The arbitrage conditions on the equity and FX volatility surface are well-known [5] and correspond to the constraint that the prices of all butterflies, call-spreads, and calendar-spreads must be positive.

However in the interest rates market the matrix of liquidly traded vanilla options forms a more complicated object, the *swaption cube* which is indexed by the expiry date of the option, the strike, and the tenor of the underlying swap. The implied volatilities of caplets are also stored in the swaption cube since a caplet can be considered as an option on a one-period swap.

This paper makes two contributions to the construction of an arbitrage-free swaption cube.

Firstly, a novel arbitrage condition is derived between a triangle of European swaptions with the same strike. In the strike direction, the butterfly arbitrage and call-spread arbitrage conditions can be applied to the swaption cube just as they can to the equity volatility surface. However instead of the calendar-spread arbitrage condition, a different condition is derived and discussed in section 1. We are not aware of any previous discussion of this arbitrage condition between vanilla swaptions in the mathematical finance literature although there is a clear relationship with the well known result that "a basket of options is worth more than an option on the basket".

Secondly, we derive a new approximation for the SABR model[1]

$$\begin{aligned} dF &= \sigma F^{\beta} dV \\ d\sigma &= \alpha \sigma dW \\ \langle dV dW \rangle &= \rho dt \end{aligned} \tag{1}$$

which has become a market standard for interpolation of swaption volatilities in the strike direction. Unfortunately there is no known exact analytics for European options in the SABR model so it is probably more accurate to say that the Hagan et al [6] analytic approximation has become a market standard. The Lee moment formula [7] shows that the implied volatility smile must flatten at large maturities, however this is violated by the Hagan et al approximation. The result is that butterflies with low strikes can have negative value when priced in this approximation. The origin of the problem is that the Hagan approximation is derived using a short-time asymptotic approximation like those introduced in [3, 4], whereas interest rates traders need to calculate prices and risks for options with maturities as high as 30 years. In section 2 we derive a novel approximation for the SABR model which reduces dramatically this problem of negative densities. There have been a number of recent papers in which sophisticated methods of differential geometry are applied to derive short-time approximations for stochastic volatility models. By contrast the approximation derived in this paper is more heuristic; however in our view what it lacks in elegance of derivation it more than makes up in practicality for pricing long dated options with realistic market data.

The methods described in this paper are particularly timely because the market turmoil of 2008 has resulted in two changes in the interest rate volatility market:

[1]Warning: a number of authors use different notation. In particular Hagan et al [6] use α to represent the volatility and ν to represent the volatility of volatility. At least one paper [9] uses β to represent the volatility of volatility. We find it more intuitive to choose a notation in which σ represents the instantaneous volatility.

- The implied volatility of short dated options on long dated swaps has increased much more sharply than the implied volatility of longer dated options on short dated swaps. The reason for this is unclear, although it is thought that it may be related to hedging activities by issuers of options on CMS spread. This increases the danger of violating the no-arbitrage condition described in section 1,
- Forward rates have decreased, whilst volatility skew has increased significantly. This has increased the degree of arbitrage at low strike introduced by the use of the Hagan approximation and makes the kind of approximation described in section 2 more important.

1 In-plane triangular arbitrage condition

Constructing an arbitrage-free volatility surface for an equity or FX rate involves checking for calendar spread arbitrage and removing this if necessary. The non-arbitrage condition is that the Black-Scholes implied variance $\hat{\sigma}^2 t$ should always be increasing, for options with a fixed moneyness K/F. However in the interest rate vanilla options market, the same condition does not exist. A caplet with exercise date 3y and a caplet with exercise date 4y are options on quite different underlyings - the two Libor rates are made up from different discount bonds. So there is nothing to stop the implied variance of the 4y caplet from being less than the implied variance of the 3y caplet[2]. For this reason, practitioners often consider that each (expiry/tenor) point on the swaption matrix can be treated independently and assume that the only requirement is to ensure that butterfly and call-spread arbitrages are avoided for each point.

However this is not correct. Although the calendar spread arbitrage condition for caplets / swaptions is quite different from that for equity options, there is nonetheless a non-trivial arbitrage constraint between European swaptions with the same strike and different option expiries. Because the three swaptions involved must all have the same strike, we call this an in-plane triangular arbitrage condition.

Let us start with some notation. $S(t, T_s, T_e, K)$ is the PV observed at time t of a payers swap whose first fixing date is T_s and whose last payment date is T_e. The coupon paid on the fixed side of the swap is K, $A(t, T_s, T_e)$ is the PV of the associated annuity and $F(t, T_s, T_e)$ is the forward swap rate. So the definition of the forward swap rate is

$$S(t, T_s, T_e, K) = A(t, T_s, T_e)(F(t, T_s, T_e) - K) \tag{2}$$

[2] It is worth noting that conventional 'short rate' models such as extended Vasicek, Black-Karasinski, quasi-Gaussian and so on, can only be calibrated to non-increasing caplet variances by the use of undesirably excessive time-inhomogeneities in their volatility and mean reversion parameters. This is a fault of the model rather than a genuine arbitrage.

Let us then use the notation $Sw(T_{ex}, T_s, T_e, K)$ to represent the PV observed at time 0 of a European option, with exercise date T_{ex}, to enter into the payers swap $S(T_{ex}, T_s, T_e, K)$. So

$$Sw(t, T_s, T_e, K) = N(0)\mathbf{E}\left[\left(\frac{S(t, T_s, T_e, K)}{N(t)}\right)^+\right] \tag{3}$$

where $N(t)$ is the value of the numeraire at time t. Standard vanilla swaptions have $t = T_s$, but our notation also allows us to describe options on a forward starting swap rate, for which $t < T_s$.

We will start by stating the triangular arbitrage condition and then work through the proof. The arbitrage condition relates the prices of 3 vanilla swaptions with the same strikes as follows:

$$Sw(T_1, T_1, T_2, K) + Sw(T_2, T_2, T_3, K) \geq Sw(T_1, T_1, T_3, K) \quad \forall\, T_1, T_2, T_3, K \tag{4}$$

The proof starts by noting that, if M_t is a martingale and if $t \leq T$ then

$$\mathbf{E}[(M_T - K)^+] \geq \mathbf{E}[(M_t - K)^+] \tag{5}$$

This result is an example of the conditional Jensen inequality (see Theorem 2.3.2 of [11]) and is the starting point for proving the standard condition for calendar spread arbitrage in an equity or FX volatility surface.

We then use this to demonstrate that if $t \leq T$ then

$$Sw(T, T_s, T_e, K) \geq Sw(t, T_s, T_e, K) \tag{6}$$

In other words, the value of a European option on a particular forward swap decreases monotonically as the exercise date becomes earlier and earlier. This result is proved by re-writing it as

$$N(0)\mathbf{E}\left[\left(\frac{S(T, T_s, T_e, K)}{N(T)}\right)^+\right] \geq N(0)\mathbf{E}\left[\left(\frac{S(t, T_s, T_e, K)}{N(t)}\right)^+\right] \tag{7}$$

$$N(0)\mathbf{E}\left[\left(\frac{A(T, T_s, T_e)(F(T, T_s, T_e) - K)}{N(T)}\right)^+\right] \geq N(0)\mathbf{E}\left[\left(\frac{A(t, T_s, T_e)(F(t, T_s, T_e) - K)}{N(t)}\right)^+\right] \tag{8}$$

and then by choosing $N(t) = A(t, T_s, T_e)$, in other words the natural numeraire for the forward starting swap in question.

$$A(0, T_s, T_e)\mathbf{E}_A\left[(F(T, T_s, T_e) - K)^+\right] \geq A(0, T_s, T_e)\mathbf{E}_A\left[(F(t, T_s, T_e) - K)^+\right] \tag{9}$$

Using the fact that $F(t,T_s,T_e)$ is a martingale in the $A(t,T_s,T_e)$-measure, we can see that the result (6) is therefore equivalent to (5).

We can then go on to prove (4) as follows

$$
\begin{aligned}
Sw(T_1,T_1,T_2,K)+Sw(T_2,T_2,T_3,K) &\geq Sw(T_1,T_1,T_2,K)+Sw(T_1,T_2,T_3,K) && (10)\\
&= N(0)\mathbf{E}\left[\left(\frac{S(T_1,T_1,T_2,K)}{N(T_1)}\right)^+ + \left(\frac{S(T_1,T_2,T_3,K)}{N(T_1)}\right)^+\right] && (11)\\
&\geq N(0)\mathbf{E}\left[\left(\frac{S(T_1,T_1,T_2,K)+S(T_1,T_2,T_3,K)}{N(T_1)}\right)^+\right] && (12)\\
&= N(0)\mathbf{E}\left[\left(\frac{S(T_1,T_1,T_3,K)}{N(T_1)}\right)^+\right] && (13)\\
&= Sw(T_1,T_1,T_3,K) && (14)
\end{aligned}
$$

Equation (10) made use of (6), whilst we moved from (11) to (12) using Jensen's inequality. In practice we find that this arbitrage condition is not often violated, at least not for moderate strikes and not once bid-offer spread is taken into account. Nonetheless we have observed cases where the calibration of exotics models has failed due to the violation of the triangular in-plane arbitrage condition.

2 A new approximation for the SABR model

2.1 Literature review

Berestycki, Busca and Florent [4] derive the short-time asymptotics for general stochastic volatility models. In the limit $T \to 0$, they show that the implied volatility of the lognormal SABR model (in other words with $\beta = 1$) is given by

$$
\begin{aligned}
\hat{\sigma}_{bbf}(\zeta) &= \sigma\zeta / \ln\left(\frac{\zeta-\rho+\sqrt{1-2\rho\zeta+\zeta^2}}{1-\rho}\right)\\
\zeta &= \frac{\alpha}{\sigma}\ln\frac{F}{K} && (15)
\end{aligned}
$$

and in the case of general beta is given by

$$\begin{aligned}\hat{\sigma} &= \alpha \ln \frac{F}{K} / \ln \left(\frac{\zeta - \rho + \sqrt{1 - 2\rho\zeta + \zeta^2}}{1 - \rho} \right) &(16)\\ \zeta &= \frac{\alpha}{\sigma} \frac{F^{1-\beta} - K^{1-\beta}}{1 - \beta} &(17)\end{aligned}$$

Our challenge is to extend this approximation in a systematic way so that it can be used for large times (at least 30y). The reason that this is non-trivial is that equation (15) suggests that $\hat{\sigma} \sim \zeta$ for large ζ, whereas the Lee moment formula [7] states that it can grow no faster than $\zeta^{1/2}$ for finite expiries.

Hagan et al [6] derives a first order correction term. In the lognormal case this is given by

$$\hat{\sigma} = \sigma\zeta / \ln \left(\frac{\zeta - \rho + \sqrt{1 - 2\rho\zeta + \zeta^2}}{1 - \rho} \right) \Pi(t)$$

where the first term in a Taylor expansion of $\Pi(t)$ is calculated. This $O(t)$ correction improves the accuracy of at-the-money options for larger t, but does nothing to deal with the problem of extreme strikes.

In the case of $\beta < 1$, Obloj [10] correctly states that the Hagan formula does not match (16, 17) because Hagan in some places replaces equation (17) with the expression

$$\frac{\alpha}{\sigma} \frac{F - K}{(FK)^{\beta/2}} \tag{18}$$

whose Taylor expansion in powers of $\ln K/F$ matches that of (17) to second order. Obloj therefore proposes a modification to the Hagan formula which avoids negative densities for the parameter set $F = 0.0801, \alpha = 0.245, \beta = 0.6, \sigma = 0.0155, \rho = -0.37, T = 20$. However increasing σ to a more realistic value such as 0.055, his approximation again gives highly arbitrageable prices.

Medvedev [9] discusses various systematic ways to extend the asymptotic result (15) to finite t in the case of $\beta = 1$. In particular section 6 of his paper is extremely interesting as it gives a mechanism for deriving series expansions in powers of t. He writes

$$\begin{aligned}\Upsilon(\zeta) &= \zeta / \ln \left(\frac{\zeta - \rho + \sqrt{1 - 2\rho\zeta + \zeta^2}}{1 - \rho} \right) \\ \zeta &= \frac{\alpha}{\sigma} \ln \frac{F}{K} \\ \hat{\sigma} &= \sigma \Upsilon(\zeta \Lambda(\zeta, t)) \Pi(t) &(19)\end{aligned}$$

where $\Lambda(\zeta, t)$ and $\Pi(t)$ are some functions which are to be determined as a power series in t.

Medvedev assumes that he can drop the ζ-dependence in $\Lambda(\zeta, t)$ and to expand $\Lambda(0, t)$ and $\Pi(t)$ as power series in t. The effect of this assumption is that his approximation amounts to reducing the effective value of α at long maturities, thereby reducing significantly the problem of negative densities. However the fundamental problem - that $\Upsilon(\zeta)$ grows faster than $\zeta^{1/2}$ for large $|\zeta|$ - is not removed. The danger is therefore that users will use higher α values to match the observed kurtosis in the market and that the problem will remain exactly as before.

Benhamou and Croissant [2] propose an alternative approximation based on the idea of local time. Their expression for the value of a call option in SABR (equation 3.10) involves a numerical integral which can be evaluated in various different ways including expressing it as the error function of a complex number. We have performed the integral using a change of variables (to deal with a singularity at $u = 0$) followed by simple trapezoidal integration.

We reproduced the results in Benhamou and Croissant's figure 1, which shows that for the parameter set $F = 0.05, \alpha = 0.2, \beta = 0.7, \sigma = 0.11, \rho = -0.5, T = 20$, their approximation, unlike that of Hagan et al, gives positive probability densities at low strikes in this case. Unfortunately we also found that for other realistic parameter sets, such as $F = 0.05, \alpha = 0.2, \beta = 0.3, \sigma = 0.033, \rho = -0.5, T = 20$ the Benhamou-Croissant approximation gives highly arbitrageable prices.

2.2 Truncated approximation for the lognormal SABR model

Let us start with equation (15). Although this formula does not obey the Lee moment formula, we can extend it by noting that the following truncated function

$$\hat{\sigma}_{trunc}(\zeta) = \sigma \left(1 - 4\rho\zeta + (\frac{4}{3} + 5\rho^2)\zeta^2\right)^{1/8} \tag{20}$$

has the following properties. Firstly, its Taylor expansion matches that of $\hat{\sigma}_{bbf}(\zeta)$ to $O(\zeta^2)$. Therefore the ATM volatility, skew and kurtosis match the short time approximation for $\beta = 1$. However as $|\zeta| \to \infty$, $\hat{\sigma}_{trunc}(\zeta) \propto \zeta^{1/4}$. Therefore the implied volatility will grow much less quickly in the wings than the constraints of the Lee moment formula, and is therefore less likely to suffer from the problems with negative densities.[3]

[3] If the approximation was only to be used in the case $\beta = 1$ then it would be sensible to choose instead $\hat{\sigma}_{trunc}(\zeta) = \sigma \left(1 - 2\rho\zeta + (4 + 3\rho^2)\zeta^2/6\right)^{1/4}$ which has the asymptotic behaviour $\hat{\sigma}_{trunc}(\zeta) \propto \zeta^{1/2}$ instead of $\hat{\sigma}_{trunc}(\zeta) \propto \zeta^{1/4}$. This is supported by Benaim et al [1] who show that the extreme strike behaviour of the lognormal SABR model takes this form. However if the approximation is to be used in conjunction

Negative probability densities are only observed in the standard lognormal approximation $\hat{\sigma}_{bbf}(\zeta)$ when the dimensionless parameter $\alpha^2 T|\ln F/K|$ becomes of order unity. We can therefore blend between $\hat{\sigma}_{bbf}(\zeta)$ and $\hat{\sigma}_{trunc}(\zeta)$ as follows

$$\frac{1}{\hat{\sigma}_{blend}(\zeta)} = \frac{w(t)}{\hat{\sigma}_{bbf}(\zeta)} + \frac{1-w(t)}{\hat{\sigma}_{trunc}(\zeta)} \tag{21}$$

$$w(t) = \max(1/\alpha\sqrt{t}, 1) \tag{22}$$

which ensures that we continue to use the lognormal Hagan approximation for the cases of short times or moderate strikes for which it is valid. The blending function $w(t)$ ensures that the curvature of the smile flattens realistically at high maturities.

Of course this method of truncating the short-time asympototic approximation at extreme strikes is completely heuristic, rather than being derived via a systematic expansion in a small parameter. In particular it is only designed to reduce the problems with negative density rather than being designed to improve the accuracy of the approximation at large t. We like to think of this method as being analogous to the method of Padé approximants in that it matches the first few terms of the Taylor expansion whilst enforcing some global analyticity requirements.

2.3 Extending to $\beta < 1$

The case $\alpha = 0$ is known as the constant elasticity of variance (CEV) model and in this case the density is exactly given by the non-central chi-squared distribution [8]. So in this case semianalytic option prices can be calculated using an infinite series. We wish to find a new approximation for the SABR model which reduces to the correct limit of the exact CEV analytics in the limit $\alpha \to 0$. We achieve this by first noting the well-known result that the CEV process

$$dF = \sigma F^{\beta} dW \tag{23}$$

can be well approximated for moderate strikes using the displaced lognormal process

$$dF = \sigma_{DD}(F + \Delta)dW \tag{24}$$

$$\sigma_{DD} = \sigma\beta F_0^{\beta-1} \tag{25}$$

$$\Delta = F_0 \frac{1-\beta}{\beta} \tag{26}$$

with the transformation into a CEV volatility described in section 2.3 for extending to $\beta < 1$, then we find that the slower growth in the wings is required to prevent negative densities.

So if we have an analytic formula for the lognormal SABR model with $\beta = 1$, we can extend this to the general case $0 < \beta < 1$ by the following 3-step process

1. Convert from the SABR model (in which F follows a CEV process) to a modified model in which F follows a displaced lognormal process.
2. Use a lognormal SABR formula to calculate the multiplicative renormalisation factor which should be applied to σ_{DD} as a result of stochastic volatility. We use the approximation described in section 2.2 but in principle other approximations could be used instead.
3. Convert back from displaced lognormal into CEV and use the exact analytics for the CEV model.

It is important to emphasise that whilst there might be a significant error when approximating a CEV process with a displaced lognormal, there will be a massive cancellation of errors between the approximations made in steps 1 and 3. We therefore find that in cases of small or moderate α that this method is very much more accurate that methods in which the local volatility term is handled using an asymptotic expansion.

Collecting everything together and using the first order time approximation first suggested in [6], the new general SABR formula is:

$$\sigma_{DD} = \sigma\beta F_0^{\beta-1} \tag{27}$$

$$\Delta = (1-\beta)/\beta F_0 \tag{28}$$

$$\zeta = \frac{\alpha}{\sigma_{DD}} \ln\frac{F_0+\Delta}{K+\Delta} \tag{29}$$

$$\hat{\sigma}_{bbf}(\zeta) = \sigma_{DD}\zeta/\ln\left(\frac{\zeta-\rho+\sqrt{1-2\rho\zeta+\zeta^2}}{1-\rho}\right) \tag{30}$$

$$\hat{\sigma}_{trunc}(\zeta) = \sigma_{DD}\left(1-4\rho\zeta+(\frac{4}{3}+5\rho^2)\zeta^2\right)^{1/8} \tag{31}$$

$$\hat{\sigma}_{blend}(\zeta) = \left(\frac{w(t)}{\hat{\sigma}_{bbf}(\zeta)}+\frac{1-w(t)}{\hat{\sigma}_{trunc}(\zeta)}\right)^{-1}(1+(\frac{\alpha^2}{12}+\frac{\alpha\sigma_{DD}\rho}{4}-\frac{\alpha^2\rho^2}{8})t) \tag{32}$$

$$w(t) = \max(1/\alpha\sqrt{t},1) \tag{33}$$

$$\hat{\sigma}_{CEV} = \hat{\sigma}_{blend}(\zeta)F_0^{1-\beta}/\beta \tag{34}$$

and then we use the infinite series for the non-central chi-squared distribution to price an option of the required strike and maturity in the CEV model with volatility $\hat{\sigma}_{CEV}$. This approximation has the following properties:

- For $\alpha = 0$, it matches the CEV analytics perfectly and there is no arbitrage;

- For $\beta = 1$, we observe no arbitrage for reasonable parameter values;
- The case $\beta = 0, \alpha = 0$ does *not* match the Bachelier formula since the underlying is not permitted to go negative. Instead it matches a normal model with an absorbing boundary at zero;
- For at-the-money options and maturities less than 10y, the difference from the Hagan approximation is relatively small;
- For low strikes, although arbitrage is sometimes observed, the degree of arbitrage is very much less than in the Hagan approximation.

These properties are illustrated in the following figures, where we collect some results comparing the results from a Monte Carlo simulation[4] against the approximation described in this paper and against the approximations from [2] and [10].

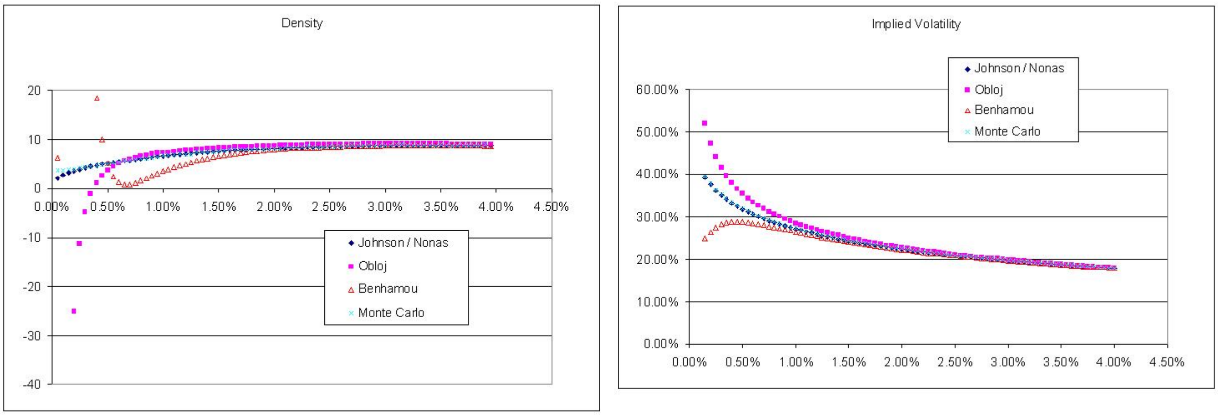

Figure 1: Implied volatility and density for the parameter set $\sigma = 0.02, \alpha = 0., \beta = 0.1, \rho = 0., T = 30, F = 0.05$. In this case the SABR model reduces to the CEV model and the new analytics matches the Monte Carlo perfectly.

3 Conclusions

The swaption volatility cube forms the basis for derivative pricing in the fixed income world. Using this data to build consistent and arbitrage free pricing models for vanilla products is therefore of the utmost importance for any more complex

[4]The Monte Carlo simulation made use of log-Euler differencing in both the volatility process and in the forward rate process. Paths were generated using Sobol numbers and Brownian Bridge path generation. 100 timesteps and $2^{20} - 1 = 1,048,575$ Monte Carlo paths were found to give sufficient numerical convergence.

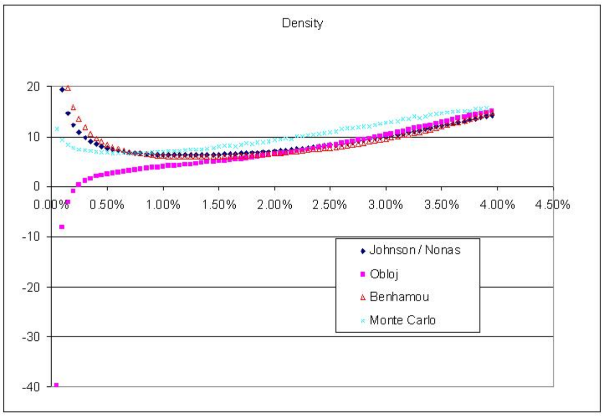

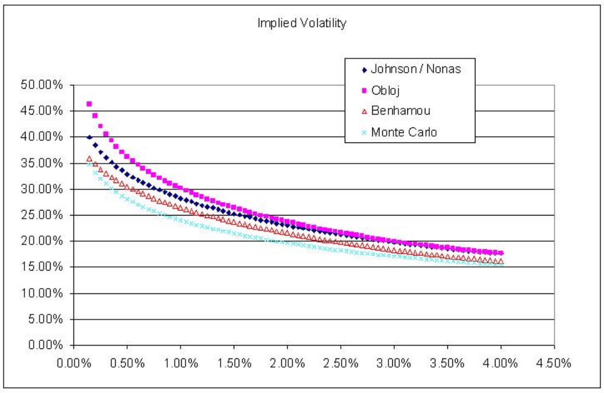

Figure 2: Implied volatility and density for the parameter set $\sigma = 0.15, \alpha = 0.3, \beta = 1.0, \rho = -0.2, T = 20, F = 0.05$. This is the lognormal case.

pricing model, as well as being an end in itself for vanilla option traders. In the first part of this paper we showed a connection between payer swaptions with the same strike and complementary maturity/tenor buckets that forms a no arbitrage condition within the swaption cube. In the second part we looked into the well known problem of implicit arbitrage due to the breakdown of the approximate analytics for European options in the market standard SABR smile parameterization for long maturities/low strikes. We suggest a three stage approximation using the representation of a CEV process by a shifted lognormal process and a robust approximation for the lognormal SABR model that reflects the Lee moment formula in the low strike wing of the smile surface. The approach has proved to produce very reliable analytical prices for a wide range of parameters with only very minor instabilities under extreme settings.

4 Acknowledgements

We acknowledge valuable discussions with Matthew Haigh, Peter Jaeckel, Martin Forde, and with all the members of the Commerzbank Financial Engineering team. The views expressed in this article are personal and do not represent the views of Commerzbank.

References

[1] S. Benaim, R. Friz and R. Lee, On the Black-Scholes Implied Volatility at Extreme Strikes, Forthcoming in *Frontiers in Quantitative Finance: Volatility*

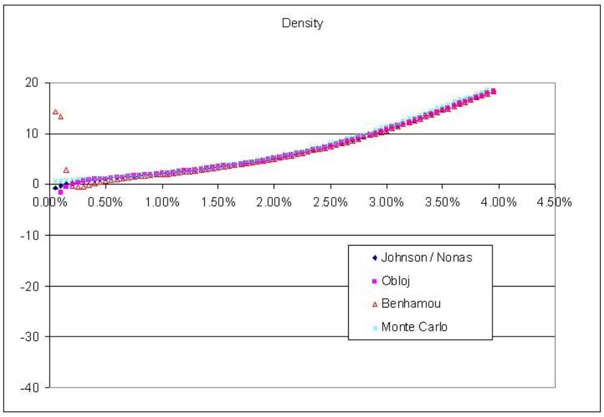

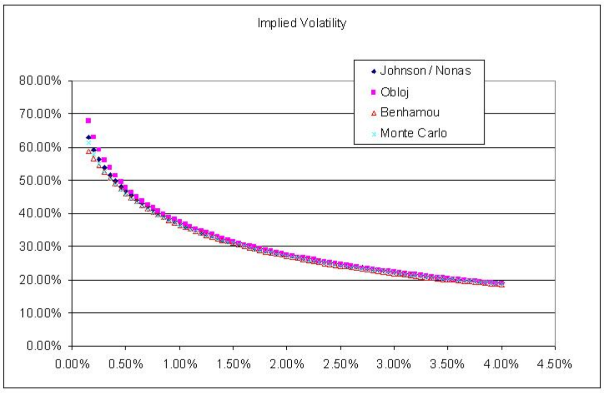

Figure 3: Implied volatility and density for the parameter set $\sigma = 0.02, \alpha = 0.25, \beta = 0.3, \rho = -0.2, T = 5, F = 0.05$. This is a realistic parameter set for the interest rate swaptions market. Note that for moderate maturities such as this one, there is very little difference between any of the approximations.

and Credit Risk Modeling, 2008.

[2] Benhamou, Eric and Croissant, Olivier, "Local Time for the SABR Model: Connection with the 'Complex' Black Scholes and Application to CMS and Spread Options" (October 2007). Available at SSRN: http://ssrn.com/abstract=1064461

[3] H. Berestycki, J. Busca, and I. Florent, Asymptotics and Calibration in Local Volatility models, Quantitative Finance 2002, volume 2, 61-69.

[4] H. Berestycki, J. Busca, and I. Florent, Computing the implied volatility in Stochastic Volatility Models, Commun. Pure and Appl. Math, 2004

[5] J. Gatheral, The Volatility Surface: A Practitioner's Guide, Wiley, 2006.

[6] P. Hagan, D. Kumar, A.S. Lesniewski and D.E. Woodward, Managing Smile Risk, Wilmott Magazine, September 2002, 84-108.

[7] R.W. Lee, The Moment Formula for Implied Volatility at Extreme Strikes, Mathematical Finance vol 14, p469-480, July 2004.

[8] C.F. Lo and P.H. Yuen, Constant elasticity of variance option pricing with time-dependent parameters, Int. Journal of theoretical and applied Finance, Vol. 3, No. 4 (2000) 661-674.

[9] A. Medvedev, Asymptotic Methods for computing Implied Volatilities under Stochastic Volatility, (October 2004). NCCR-Finrisk Working Paper No. 204. Available at SSRN: http://ssrn.com/abstract=667281

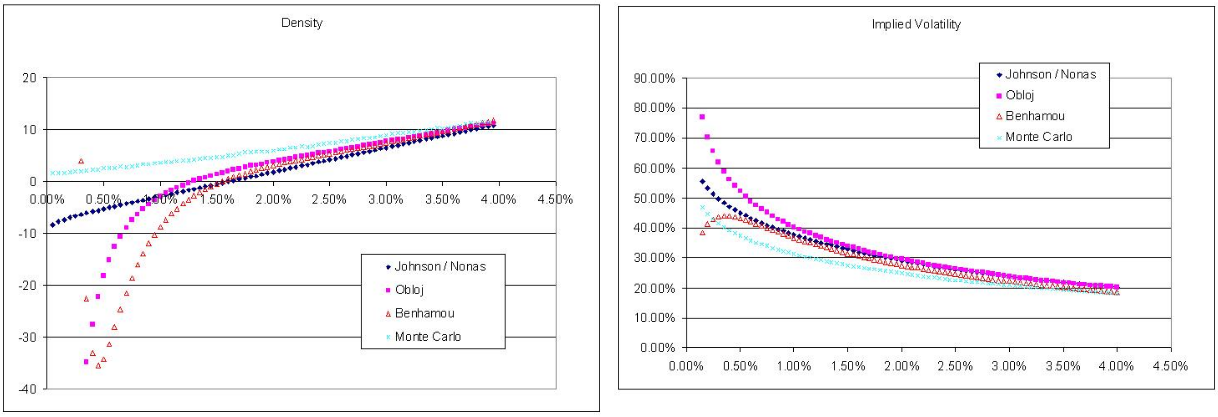

Figure 4: Implied volatility and density for the parameter set $\sigma = 0.02, \alpha = 0.25, \beta = 0.3, \rho = -0.2, T = 20, F = 0.05$. This is a realistic parameter set for the interest rate swaptions market. Note that both analytics have a region of negative densities. Negative densities occur at slightly higher strikes in the new approximation described in this paper, but the density does not go nearly as negative as in the approximations of Obloj or Benhamou.

[10] J. Obloj: Fine-tune your smile correction to Hagan et al, Wilmott magazine, May/June 2008 p102.

[11] S. Shreve, Stochastic Calculus for Finance, Volume II, Continuous-time Models, 2004, Springer-Verlag.