A Pairwise Local Correlation Model

Frank Koster*, Daniel Oeltz†

Abstract

In this paper we develop a local correlation model which uses a given correlation matrix and a generic function $g(t, m_i, m_j)$ to compute the local correlation between any asset-asset pair (i, j) of a basket of underlyings. The arguments m_i, m_j are spot moneynesses. The generic function is calibrated to fit the implied volatilities of an equity index like DAX or EUROSTOXX50. The advantage of this approach is that we do not need to simulate the complete index basket when pricing options on a usually small subset of this index. The approach does also not suffer from the so-called chewing gum effect of correlation models where the local correlation depends on just the index value.

The main part of our work is to show how to calibrate the generic function for each time step with constraints on the positive definiteness of the resulting correlation matrices and the smoothness of g to allow for stable evaluation. In a second part we present numerical experiments which show the impact on prices and deltas for synthetic auto-callable instruments.

Keywords local correlation, Monte Carlo, Gyöngy's theorem

1 Introduction

It is well-known that the multi-asset version of, for example, the local volatility model with a constant correlation matrix is not able to reproduce the smile of implied volatilities of equity index options. The simulated smile is

*DekaBank, Mainzer Landstrasse 16, D-60325 Frankfurt a.M., Germany, emails: frank.koster@dekabank.de, d.oeltz@rivacon.com

†RIVACON GmbH

usually flatter than the smile observed in the market. The reason for this miss-pricing is that the tendency of the correlation to go up in falling markets is not reflected by the constant correlation assumption. Therefore, stochastic correlation models have been developed. Here the correlation matrix follows a stochastic process which is correlated to the driving spot processes. More information on such models can be found in e.g. (Ahdida and Alfonsi 2013), (Bossu 2014), (Brockhaus 2016) or (Zetocha 2015). Another approach to solve the index smile problem are local correlation models where the correlation matrix depends on the spots of the index constituents via a deterministic function. To explain these models let us briefly introduce some basic notion which will be used throughout in this paper. The time t constituent spots are denoted by $S_t := (S_t^1, ..., S_t^n)$ and the ith constituent contributes with weight α_i to index I

$$I_t := I(S_t) := \sum_{i=1}^{n} \alpha_i S_t^i \ . \tag{1}$$

In the non-quanto case the correlated spot processes read

$$\begin{aligned} dS_t^i &:= \mu_t S_t^i dt + \sigma_i(t, S_t^i) S_t^i dW_t^i \\ \langle dW_t^i, dW_t^j \rangle &= \rho_{ij}(t, S_t) dt \end{aligned} \tag{2}$$

where $\rho(t, S_t)$ is the instantaneous local correlation matrix. Such models have been considered, for example, in (Langnau 2010) where the local correlation is constructed to match the instantaneous local variance of the simulated index with the local variance of the index at each time point for each path. (Reghai 2010) models the correlation matrix as a linear combination of two limiting correlation matrices, for example a historical correlation matrix ρ^0 and the 'worst case scenario matrix' ρ^1 with all entries equal to one

$$\rho(t, S_t) = \rho^0(1 - \lambda(t, S_t)) + \rho^1 \lambda(t, S_t) \ . \tag{3}$$

As he points out it is essential for the calibration problem that in his approach λ depends on S_t via an aggregator function $L : \mathbf{R}^n \mapsto \mathbf{R}$ only

$$\lambda(t, S_t) = f(t, L(S_t)) \ . \tag{4}$$

For our problem of fitting the correlation model to (liquid) market quotes of the index, a natural choice for the aggregator is the index $I(.)$. This approach

has been coined in (Guyon 2014) as 'local in index' – a notion we will also use in the present paper. (Reghai 2010) uses an iterative fix point algorithm to fit λ where the difference between index implied volatility and simulated index implied volatility is used to update the λ iterates. A different calibration procedure is developed in (Guyon 2014), (Jourdain and Sbai 2009) where λ is bootstrapped time step by time step matching expected values of the local variance which results from (1) and (2) and the local variance computed directly from index option prices or implied volatilities via the Dupire formula (Dupire 1994) or (Gatheral 2006).

Once one has fitted such a function f one needs to simulate all index constituents (or a meaningful large subset as in (Ahdida 2011)) to price instruments even if the instrument depends on just two constituents of the index. This may be a disadvantage if a large number of multi-asset instruments shall be priced and one can not re-use the simulated paths of the index constituents to price all these instruments.

Another feature of the local in index approach is the so-called chewing gum effect. It basically means that the correlation does not change if, for example, the movements of the constituents offset each other, but one is interested in the movements of the worst and second worst performing constituents where one expects a change in the correlation between these two assets.

To address these two points we propose a local correlation model where the element $\rho_{ij}(t,.)$ is computed from only S^i_t and S^j_t by means of a generic function g. From a practical point of view this has also the advantage that it is possible to use a calibrated local correlation model for pricing instruments whose constituents are not part of the index, for example, if a constituent leaves the index during the lifetime of the instrument. In (Zetocha 2015) this very desirable feature is called transferability. Closely related is the question how different the models will be for two calibration indices with overlapping constituents and what the price differences will be for basket instruments on common members? In this paper we can not give answers to these questions as we haven't done any analysis or experiments in this direction so far.

For data sets with high index implied volatility skew the local in index/pairwise correlation extensions of the local volatility model may fail to fit the index smile, usually on the down side. Stochastic volatility models generate extra skew and may reduce this issue. Technically it is straight forward to extend the local in index/pairwise correlation model to stochastic (local) volatility models as long as one uses a simple methodology to

extend the spot-spot correlation matrix to a full correlation matrix with spot-variance and variance-variance correlations like (Jäckel and Kahl 2009) or (Delanoe 2014). We did some initial tests with the Heston local volatility model (Heston 1993) and found that the additional skew may be not enough to sufficiently lift the basket local variance for a correlation of one above the index local variance. This might be different for other stochastic volatility models we have not yet considered. For that, we deliberately do not cover stochastic volatility models in this paper and focus on the local volatility model.

Another, even more general class of local correlation models was introduced in (Guyon 2015) which promises to fit the index smile under all circumstances. Here the constituent leverage functions and the local correlation function are jointly fitted by a fixed point iteration. Of course, the forward volatility dynamics may be different from what one may be used to. To maintain a consistent view on the portfolio one probably needs to price even simple instruments like American vanilla options using the full multi-dimensional process.

This paper is organized as follows. In the next section we introduce our model which we call pairwise local correlation model in this paper. We describe the respective calibration method in the third section and in the fourth section we present numerical experiments on the fitting quality for DAX market data and compare our model with results for the local in index approach.

In the last section we compare prices and deltas for synthetic auto-callable instruments for the constant correlation model, the pairwise model and the local in index correlation model.

2 Local Correlation Model

The basic idea for the pairwise correlation model is that the elements $\rho_{ij}(t,.)$ of the correlation matrix are functions of t, S_t^i and S_t^j only. In the spirit of (3) we could introduce 'mixing' functions λ_{ij} for each pair i, j

$$\rho_{ij}(t, S_t) = \rho_{ij}^0(1 - \lambda_{ij}(t, S_t^i, S_t^j)) + \rho_{ij}^1 \lambda_{ij}(t, S_t^i, S_t^j) \ .$$

But, this gives a large number $n(n-1)/2$ of mixing functions λ_{ij} to be fitted and the calibration problem is likely to be under-determined.

To reduce the number of unknown functions we do one more step and model all λ_{ij} by a single generic function g which depends on t and moneynesses S_t^i/F_t^i only

$$\lambda_{ij}(t, S_t^i, S_t^j) = g(t, \frac{S_t^i}{F_t^i}, \frac{S_t^j}{F_t^j}) \ . \tag{5}$$

The scales F_t^i may be, for example, the current spots S_0^i, the time t forwards or exponentially time-averaged spots using the idea of anchoring (Wilmott et al. 2014), (Guyon 2014b). Function g can be seen as an average relationship between any pairwise correlation and the respective two spot levels.

g is subject to a number of constraints, for example, symmetry

$$g(t, x, y) = g(t, y, x). \tag{6}$$

If we assume that $\rho_{ij}^0 < \rho_{ij}^1$, then we get from $|\rho_{ij}| \leq 1$ the following lower and upper bounds

$$\underline{g} := \max_{ij} \frac{-(1+\rho_{ij}^0)}{\rho_{ij}^1 - \rho_{ij}^0} \ \leq \ g \ \leq \ \min_{ij} \frac{1-\rho_{ij}^0}{\rho_{ij}^1 - \rho_{ij}^0} =: \overline{g} \ . \tag{7}$$

But, the most important constraint is that the resulting matrices $\rho(t, S_t)$ should be positive definite. It is possible to write this as a continuum of linear constraints on g

$$0 \leq z^t \cdot \rho(t, S) \cdot z \quad \text{for all} \ \ z \in R^n \ . \tag{8}$$

One can easily show that (8) implies (7), but for the calibration we will use a largely simplified version of (8) which does not always lead to positive definite ρ. In this case (7) may not be automatically fulfilled. Moreover, our model requires an additional projection of the matrices ρ computed from g onto the set $\mathcal{C}$ of n-dimensional correlation matrices. We denote this operation by $\mathcal{P}_{\mathcal{C}}$. This 'projection' may lead to a discrepancy between the actually used correlation matrix $\mathcal{P}_{\mathcal{C}}(\rho)$ and the matrix $\rho(t, S)$ calibrated by the method in Section 3.

The projection is an additional step of eigenvalue clipping and diagonal scaling as in (Jäckel 2002) if the regular Cholesky-decomposition of ρ fails. This procedure may not give the nearest correlation matrix like (Higham 2002), but it is often faster by a factor 10...100 and the projection difference $\rho - \mathcal{P}_{\mathcal{C}}(\rho)$ seems to be not much larger than the optimal projection difference in practice.

3 Calibration

In this section we describe the calibration of the local correlation function g.

The calibration procedure uses the Monte-Carlo-based particle method from (Guyon 2014). For each time step we model the generic function $g(t^m, ., .)$ by a bi-linear interpolating spline

$$g(t^m, u, v) = \sum_{l,l'=1}^{L} g_{ll'}^m \phi_l(u)\phi_{l'}(v) \ . \tag{9}$$

The basis functions $\phi_{l'}$ are the usual hat-like linear B-splines w.r.t. moneyness grid $\{\zeta_l\}_{1\le l\le L}$.

Our aim is to fit the coefficients $g_{ll'}^m$ such that the instantaneous (i.e. total local) variance of the index simulated from the constituents matches the local variance $(\sigma^I)^2$ computed from the index implied volatility

$$(s\sigma^I(t,s))^2 = E[\sum_{i,j} \alpha_i S_t^i \sigma_i(t, S_t^i)\rho_{ij}(t, S_t)\sigma_j(t, S_t^j)S_t^j\alpha_j \mid I(S_t) = s] \ . \tag{10}$$

If this condition holds for all times and strikes then the marginal distributions of the simulated and the market index will agree according to Gyöngy's Theorem (Gyöngy 1986). This approach is often called Markovian projection (Piterbarg 2006). For convenience we dropped the time index m and set $t = t^m$. For the ease of presentation we define

$$c(t, S_t, \rho) := \sum_{i,j} \alpha_i S_t^i \sigma_i(t, S_t^i)\rho_{ij}(t, S_t)\sigma_j(t, S_t^j)S_t^j\alpha_j \tag{11}$$

as a short notation for the instantaneous variance of the simulated index for a given local correlation matrix ρ. We estimate the right hand side of (10) from Monte-Carlo simulated spots and volatilities by means of smoothing kernels ψ_k which must be a decomposition of unity, i.e.

$$\sum_k \psi_k(s) \equiv 1 \quad \text{for all } s \ . \tag{12}$$

Here, the kernels approximate the indicator functions with respect to a grid $\{s_k\}_{1\le k\le K}$ of index spots. For example, the conditional expected value of a random variable $X(S_t)$ is approximated by

$$E[X(S_t) \mid I(S_t) = s_k] \approx E_k[X(S_t)] := \sum_{\omega=1}^{N} X(S_t^\omega) \frac{\psi_k(I(S_t^\omega))}{\sum_{\omega'=1}^{N} \psi_k(I(S_t^{\omega'}))}$$

where N is the number of Monte-Carlo paths and X^ω the simulated value which is in our case the instantaneous variance $c(t, S_t^\omega, \rho)$.

In general we may not satisfy the fitting equation (10) for all s_k on the given grid. In addition, we are interested in smooth solutions, as in our experience the smoothness of the diffusive terms plays a significant role for the convergence of a Monte-carlo scheme. Therefore, we use a weighted least square method with an additional penalty term. The cost function of the respective optimization problem consists of the least square version of (10) and a penalty term to force sufficient smoothness of g

$$Q(g^m) := \sum_{k=1}^{K} w_k^2 \left((s\sigma^I(t, s_k))^2 - E_k[c(t, S_t, \rho(g^m))] \right)^2 + \mu \sum_{l,l'=1}^{L} (\Delta_{ll'} g^m)^2 \tag{13}$$

where

$$w_k := \frac{1}{s_k^2} \left(\frac{1}{N} \sum_{\omega=1}^{N} \psi_k(I(S_t^\omega)) \right)^\eta . \tag{14}$$

$\Delta_{ll'}$ is the usual five-point finite difference approximation of the Laplace operator on the moneyness grid $\{\zeta_l\} \times \{\zeta_{l'}\}$ and μ is a suitable penalty parameter. The fitting weights w_k are computed from the probability hitting s_k to the power of η divided by s_k^2. This makes the calibration focused on the relevant index strikes. Parameter $\eta = 2/3 < 1$ was chosen after some initial development tests to make the weights not too localized near the index forward. It was not changed since then. As an additional trick we replace the term $(s\sigma^I(t, s_k))^2$ in (13) by

$$\min((s\sigma^I(t, s_k))^2, E_k[c(t, S_t, \rho^1)])$$

because we had many test data sets were the index total local variance was (much) higher than the maximum possible local variance of the basket attained for a correlation of 1. This discrepancy dominates the cost function and makes the optimizer insensitive to discrepancies for e.g. upside strikes.

Note that the simulated index local variance c depends *linearly* on the coefficients g^m. For that reason, $Q(g^m)$ is a quadratic function of g^m. The quadratic optimization problem is augmented by the following linear constraints:

$$\begin{aligned} g_{ll'}^m &= g_{l'l}^m \quad \forall l, l' \quad \text{(see (6))} \\ \underline{g} \leq\ & g_{ll'}^m \leq \overline{g} \quad \text{(see (7))} \end{aligned}$$

and a large number of linear constraints for the heuristics to force approximate positive definiteness

$$0 \leq z^{m,\omega} \cdot \rho(t^m, S^{\omega}_{t^m}, g^m) \cdot z^{m,\omega} \quad \forall \omega \in \Omega \subseteq \{1, ..., N\} \ . \tag{15}$$

Here $\rho(t^m, S^{\omega}_{t^m}, g^m)$ is the local correlation matrix for the spot vector $S^{\omega}_{t^m}$ and the coefficients g^m.

The test vectors $z^{m,\omega}$ are computed by the following scheme: for $m = 0$ we compute N random $N(0, id)$-distributed vectors. Alternatively we could have used the eigenvectors of the original correlation matrix. For $m > 0$ we use the eigenvectors for the minimal eigenvalues of the correlation matrices $\rho(t^{m-1}, S^{\omega}_{t^{m-1}}, g^{m-1})$ of the last time step and the same path. The set Ω is defined as every 10th path. Larger sets did not really reduce the number of non positive definite correlation matrices in the following time step, but significantly increased the costs to compute the minimal vectors and to solve the optimization problem.

The optimization problem has a quadratic cost function and only linear constraints, i.e. it is convex and has a unique solution. To solve this optimization problem we use an active set method (Goldfarb and Idnani 1983, DiGaspero 2007). This method keeps and updates a QR-decomposition of the active subset of the constraints matrix. It is therefore important to eliminate the symmetry constraints (6) beforehand to speed up the optimization procedure. Moreover, one can significantly improve the performance by a clever matrix-vector multiply for the constraints matrix which is sparse for the present calibration problems and making sure that the constraints matrix is equilibrated with respect to e.g. the Euclidean norm. These optimizations will be included in future versions of DiGaspero's QuadProg++ software.

A somewhat unwanted property of the above calibration is that the fitted correlation functions $g(t, ., .)$ may be quite different for consecutive time steps as there is no constraint or penalty which directly forces smoothness with respect to time. At or near expiries of listed index or constituent options the local volatilities jump and therefore the local correlation function will also jump. But, between expiries the correlation function should be smooth. The non-smooth behavior is closely related to the number of active constraints in the solution of the optimization problem. If there are not many of them, or if just some of the box constraints (7) are active, then the local correlation function is usually smooth with respect to time. Otherwise this may not hold. In the next section we show some examples. One option is just to

accept the non-smooth behavior. However, we propose a simple smoothing technique as post-processing step.

Let $T^1, ..., T^M$ be the expiries of listed index and constituent options. Then define the index sets $J^k := \{m | t^m \in]T^{k-1}, T^k]\}$. The coefficients of the smoothed bi-linear interpolating spline $\bar{g}^m$ are computed by least squares regression to a linear function

$$\bar{g}^m_{ll'} = f^k_{ll'}(t^m) \text{ if } m \in J_k$$

for each (l, l') where $f^k_{ll'}$ is the regression function through the nodes $(t^{m'}, g^{m'}_{ll'}), m' \in J^k$. If J_k contains few time indices only, then we use just the mean value $\sum_{m'} g^{m'}_{ll'}/|J^k|$, i.e. constant regression. We also tested piecewise constant regression for all periods as this even further reduces the complexity of the local correlation function. Note, regression with respect to the index strike direction is also used in (Guyon 2016). More details on B-splines and regression in the context of the particle method can be found in (Corlay 2016).

4 Numerical Experiments

In this section we present numerical results for the calibration of the pairwise correlation model and compare them with calibration results for the local in index model. Therefore we start by introducing how we implemented the calibration of the local in index correlation. The values $\lambda(t, s_k)$ satisfying

$$(s_k \sigma^I(t, s_k))^2 = \\ E[\sum_{i,j} \alpha_i S^i_t \sigma_i(t, S^i_t)(\rho^0_{ij}(1 - \lambda(t, s_k)) + \rho^1_{ij}\lambda(t, s_k)\sigma_j(t, S^j_t) S^j_t \alpha_j \mid I(S_t) = s_k]$$

are computed by solving a weighted least square problem. We also used a penalty $\mu \|\partial_s \lambda(t, s)\|^2$ to control the smoothness of $\lambda(t, .)$ with respect to the strike. The index strike grid and the least square weights are those of (14), i.e. the same as for the calibration of the pairwise correlation model. The penalty weight μ is 10^{-6}. We also include simple box constraints on the $\lambda(t, s_k)$ values in the optimization

$$-\lambda^* \leq \lambda(t, s_k) \leq 1.$$

The parameter λ^* is computed in advance as the maximum value λ such that $\rho^0(1-\lambda)+\lambda\rho^1$ is still positive semi-definite.

By these means our implementation of the local in index calibration is very close to the method used for the pairwise model which helps to compare these local correlation models.

(Market data) We applied the calibration methods to market data for the German DAX index and the respective 30 stocks. The first batch is market data from 2015 of which we show here the results from 2015/02/15. The second batch is data from 2017 where the index implied volatilities have a much higher skew which makes the calibration of our model harder. This is a problem we did underestimate initially. The results for 2017/01/04 and 2017/02/20 show the range of fitting qualities with the high skew data. All implied volatilities are auto-fitted SSVI parametrizations (Gatheral et.al. 2014).

As the DAX is a performance index the weights α_k are adjusted according to the dividend payments of the constituents, see (DeutscheBoerse 2014). The correlation matrix was computed from historical stock returns and shifted so that the simulated at the money index implied volatility roughly fits the market.

(Calibration parameters) The calibration period for the local correlation models spans five years. The Monte-Carlo scheme uses an Euler time discretization with 260 time steps and 100.000 paths. The index strike grid used for the smoothing kernels (12) has 400 points with logarithmic spacing and spans a range of 30% to 200% of the current index spot. For reasons of simplicity we use linear B-splines as kernels. We used a regularization parameter $\mu = 10^{-8}$ which is a compromise between fitting accuracy and smoothness of g.

The moneyness grid for the pairwise local correlation function in (9) also has a logarithmic spacing and spans a range of 0.15 to 3.5 with 35 points.

(Is a perfect fit possible?) For a given distribution of the spots S_t, the right hand side of the calibration equation (10) takes its maximum at $\rho = \rho^1$. If the index local variance on the left hand side is larger than this maximum then the calibration equation can not be fulfilled for the particular time and strike. As all the 31 implied volatilities are fitted completely independent of each other there is no guarantee that this does not happen. And for the high skew data it happens almost all the time. In Figure 1 we show the local variances of the index (i.e. the l.h.s. of (10) divided by s^2), the local variance of the basket for $\rho = \rho^1$ and for the fitted local correlations. For the 2015

data the maximum local variance of the basket always exceeds the index local variance except for short expiries and low strikes not shown in the Figure, while for the 2017 data there is always a large region of strikes where the maximum basket local variance is too low. It is clear that a perfect calibration of (any) local-volatility/local correlation model to the index smile is not possible in such situations. Moreover, the calibration of the pairwise local correlation model is more complicated then, because the fitted correlation matrices are close to ρ^1 and, thus, likely to be not positive definite.

(Calibration quality) Next, we consider the quality of fitting the index implied volatility smile. We show the results of the pairwise local correlation model and also the local in index model. For the pairwise model the DAX constituents are simulated using the projected, positive definite correlation matrices. This means that the simulated implied volatility smile includes the projection differences.

Figure 2 shows the implied volatility smiles for expiries $T = 1$ (left) and 3 years (right). The strike range of 30% to 200% is quite large to give a comprehensive picture of the fitting quality. The top row shows the results for 2015/02/15, the middle row for 2017/01/04 and bottom those for 2017/02/20. For the 2015/02/15 data both local correlation models fit the index implied volatility smile quite well. For the 2017 data this is not possible as explained before. Nevertheless, within the given limitations the local in index model produces good fits. The pairwise model produces similar fits, but usually they are somewhat worse, especially on the upside. Our impression is that the pairwise model sometimes can not so easily switch between full correlation on the downside and de-correlation on the upside.

(Smoothness of local correlation function) Besides accuracy we are interested in the smoothness of the calibrated local correlation functions g, in particular with respect to time, as we may need to evaluate the local correlation function at time points which are not contained in the calibration time grid when pricing instruments. In Figure 3 the diagonal values of $g(t, m, m)$ are shown vs. time for the three considered market data, left without smoothing, right with smoothing by piecewise linear regression. Note, we actually plot $1 - g$, i.e. the factor for the original correlations, because this makes it easier to grasp the features visually.

Obviously the non smoothed g is quite noisy for the 2017 market data and there are large regions where $1 - g$ is zero and, thus, a full correlation of 1 will be used by the local correlation model. A side effect if this is that more p.s.d. constraints are active which seems to be the reason for the quite

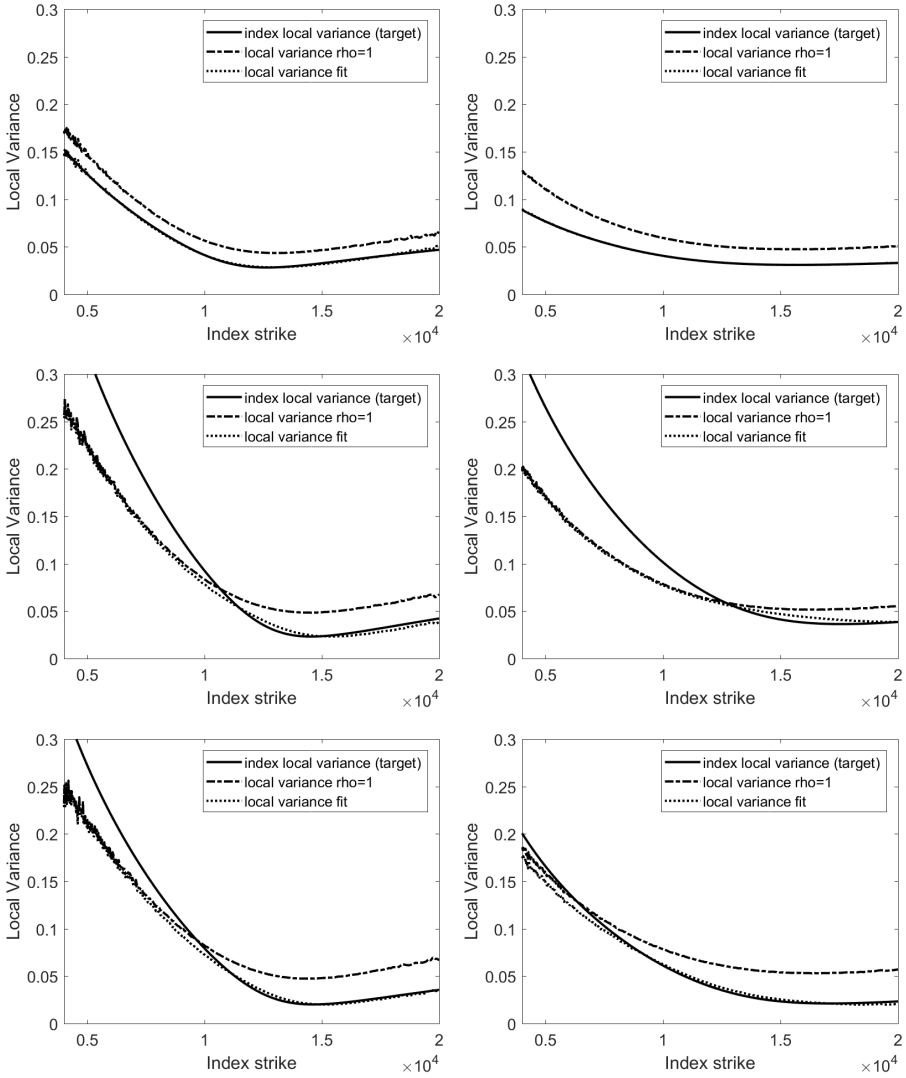

Figure 1: Local variance at $t = 1y$ (left) and $t = 3y$ (right) for 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom) data.

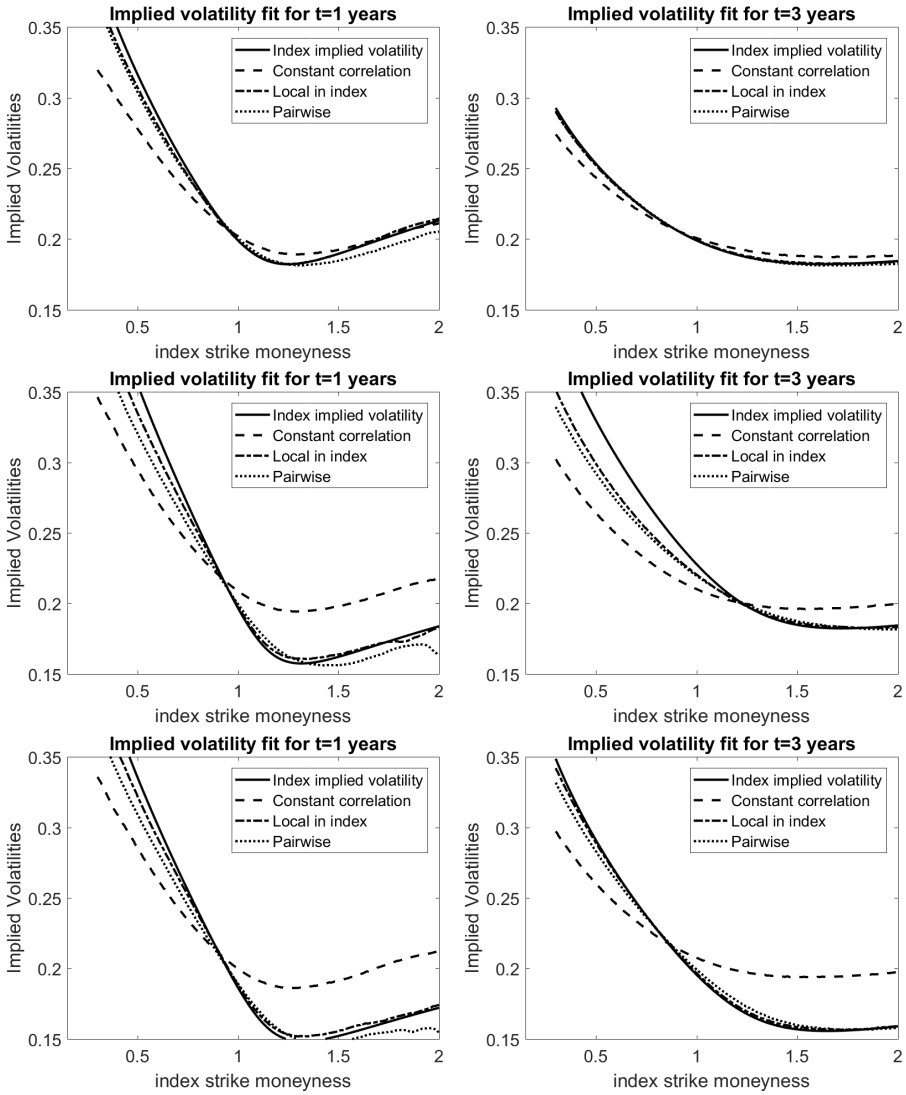

Figure 2: Implied volatilities for $t = 1y$ (left) and $t = 3y$ (right) for 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom) data.

different time slices of g. For the 2015 data the box constraints are almost never active and the much faster cpu times for the optimizer (in comparison to 2017 data, see comments below) indicate that also the number of active p.s.d. constraints is much lower so that the time slices of g are basically the solutions of unconstrained least squares problems with input data which depends more or less smooth on time, except at option expiries of course.

The smoothed local correlation functions on the right side of Figure 3 look much more friendly and aside from the full correlation regions they are not qualitatively different for the 2015 and 2017 data. The question is of course how smoothing affects the fitting quality? In Figure 4 we show again the implied volatilities for $t = 1y$ (left) and $3y$ (right) for no smoothing, the piecewise constant and the piecewise linear regression method. The fitting quality is nearly as high as without smoothing. This was the case for all the data sets we used for our tests. Finally, Figure 5 shows the local correlation functions $g(t, ., .)$ for $t = 1y$ to give an impression how the correlation depends on both moneynesses.

(Positive definiteness) As a measure for non positive definite correlation matrices we consider the difference between the correlation matrices computed directly from the local correlation function and the projection to $\mathcal{C}$

$$|\rho(t, S_t) - \mathcal{P}_{\mathcal{C}}(\rho(t, S_t))|_{l_1}/n^2$$

for each time step and each path. A reasonable norm for this difference is the l_1- norm divided by n^2 which gives the average discrepancy of a single pairwise correlation. In Figure 6 we show the maximum and mean values of the projection difference over all simulated paths for each time step. Moreover, we plot these measures for the original and smoothed (linear regression) local correlation functions. For the 2015 data less than 1% of the correlation matrices are not positive definite and there are visible projection differences for $t < 1$ only, but for the 2017 data indefinite correlation matrices occur all the time. For the non smoothed local correlation up to 65% of the correlation matrices are not positive definite which is quite much. However, the time average of the mean projection difference is less than 0.003 which means that $\mathcal{P}_{\mathcal{C}}$ adjusts a single pairwise correlation ρ_{ij} by roughly 0.003 on average – which is small. Given that the calibration quality is still good this seems to be acceptable. For the smoothed local correlation the differences are significantly smaller. The time average of the mean projection difference is

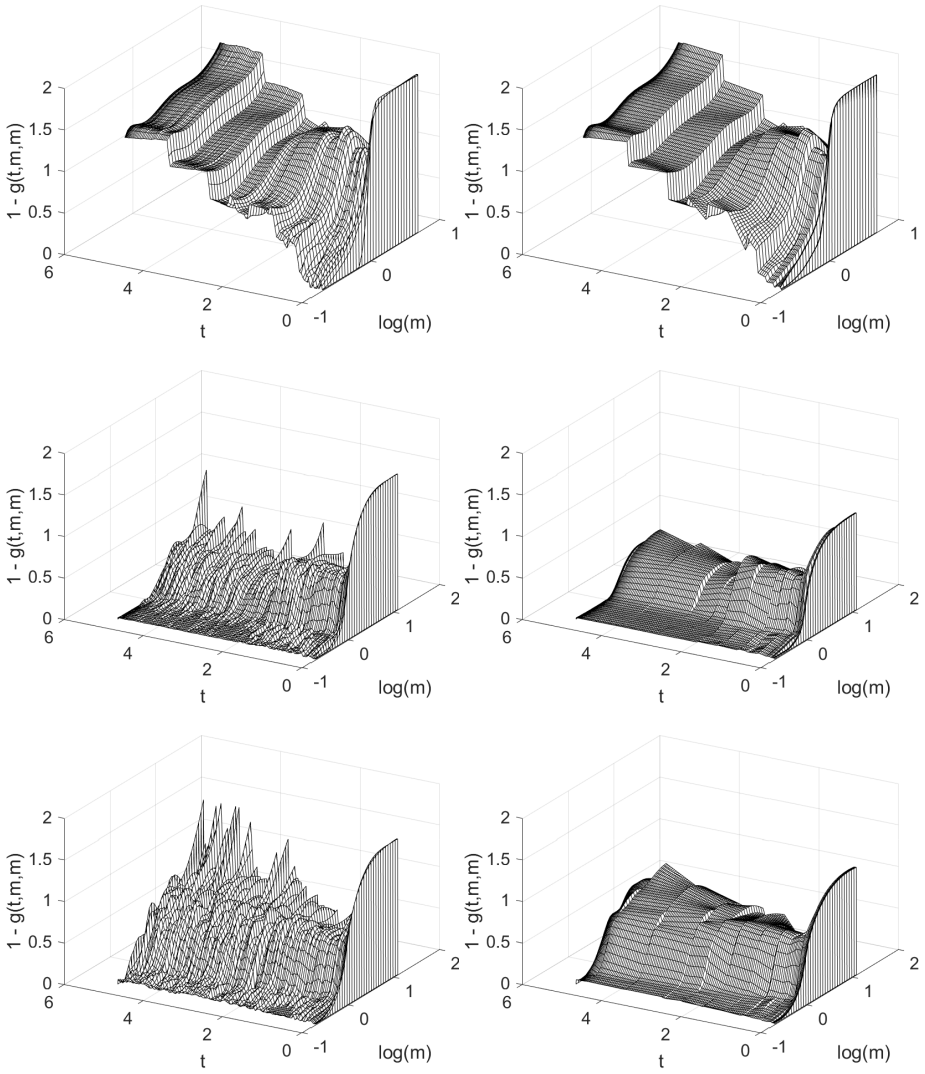

Figure 3: Diagonals of pairwise correlation functions for 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom) data. No smoothing (left), linear regression (right).

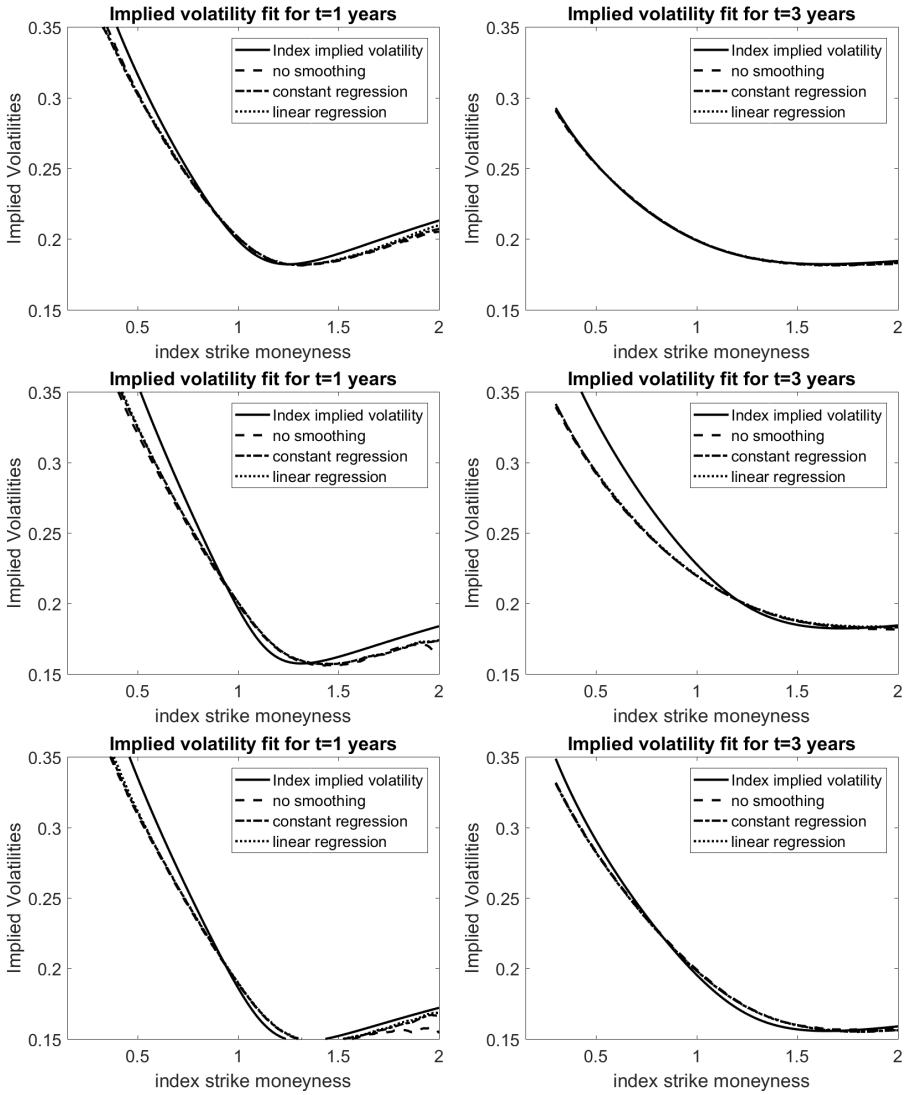

Figure 4: Implied volatilities for $t = 1y$ (left) and $t = 3y$ (right) for 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom) data.

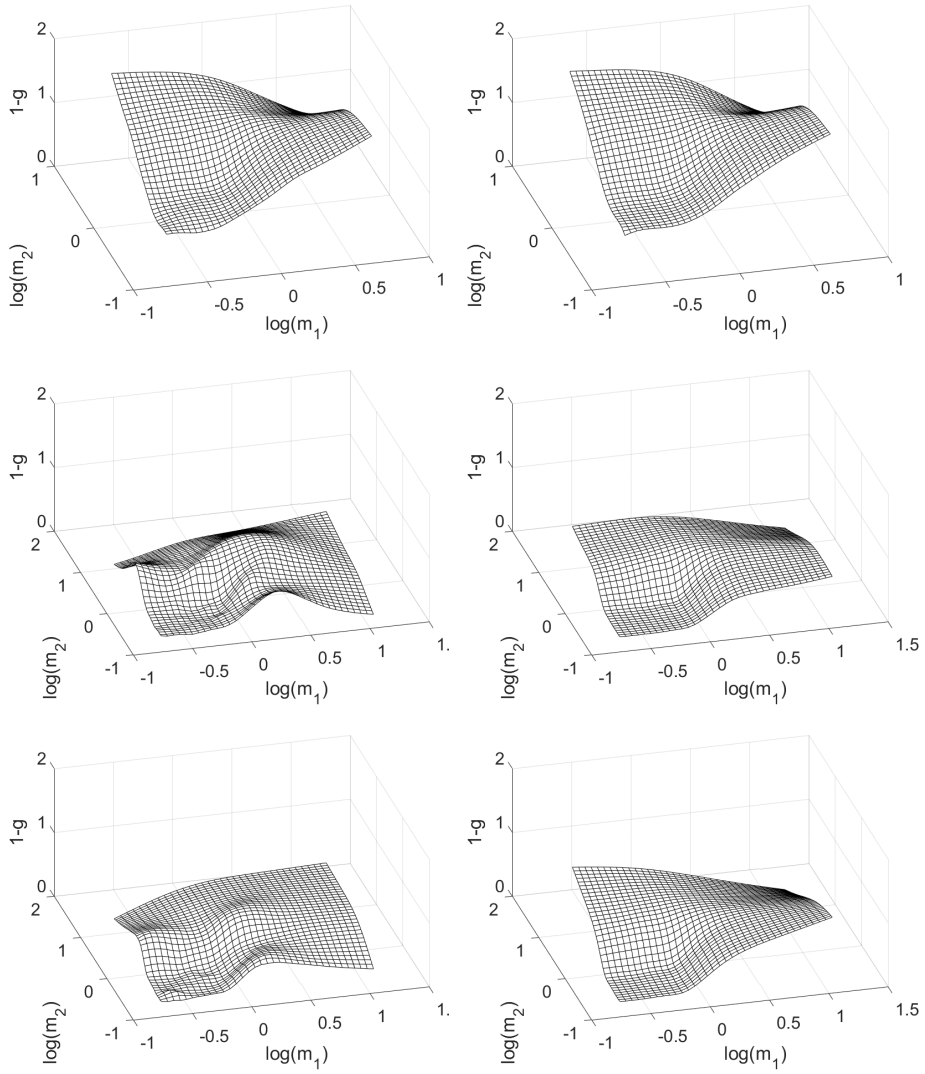

Figure 5: Pairwise correlation functions for $t = 1y$. No smoothing (left), linear temporal regression (right), as of 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom).

only about 0.0007. Here, up 45% of the correlation matrices are not positive definite.

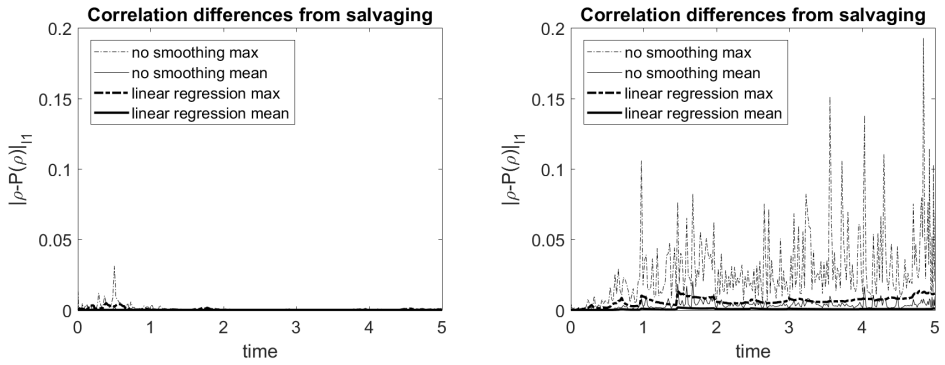

Figure 6: Maximum and mean differences from projection to $\mathcal{C}$ Left 2015/02/15. Right 2017/01/04.

(Computing Times) All tests were run on a laptop with Core i5-6200U cpu (2 cores). The calibration of λ took about 3 minutes and the calibration of g about 15 minutes for the 2015 data and between 20 and 30 minutes for the 2017 data. The longer calibration times are due to the salvaging operations and the quadratic optimization with more active constraints.

5 Impact on Pricing

In this section we use the correlation models to compute the prices of synthetic auto-callable instruments. These instruments have three yearly periods with a single monitoring date at the end of each period. The trigger functions are the worst-of or best-of the first three DAX constituents relative to the individual spots at the valuation date. If the value of the trigger function is above the redemption level at one of the three redemption dates the instrument is called and pays one Euro. Otherwise the instrument pays nothing. The redemption level varies between 0.6 and 1.3. The market data is again as of 2015/02/15, 2017/01/04 and 2017/02/20. For the local in index correlation model we simulate all 30 DAX constituents to be able to evaluate the local correlation function, but we use just the first three constituents to evaluate the trigger function.

The following Figures 7 and 8 compare prices for the worst-of and best-of trigger functions obtained with constant correlations, the local in index and the pairwise correlation model. For the later we use non-smooth and the linear smoothed local correlation functions to find out how stable the results are under such an ad hoc parameter reduction. The graphs on the right side show the differences to constant correlation prices. For the worst-of trigger the pairwise correlation model gives consistently higher prices than the local in index model which indicates that the instantaneous correlations are higher for down moving paths. For the best-of trigger both correlation models give nearly the same prices which, in our opinion, indicates that up-moving paths decorrelate in a similar way.

We found it interesting, that the spread 'maximum price difference to constant correlation minus minimum price difference' is quite similar for both correlation models. Price differences are larger by a factor of about three for the high index skew data from 2017 data than for low index skew 2015 data.

Smoothing the pairwise local correlation function has a surprisingly small effect on prices which is in line with our observations for fitting the index smile. This shows that a local correlation model with good fitting properties is possible with a much smaller number of parameters than what comes directly out of the particle method.

Next, we are interested in differences of delta which is computed by a finite difference with a $\pm 1\%$ spot shift. For simplicity, we only consider the delta with respect to the first constituent (Adidas ADS) and only for the worst-of trigger function. An important aspect of computing delta are stickiness assumptions. We use a sticky strike assumption for the implied volatility. For the correlation model we compare two variants:

(recalibration) The prices for the ADS up- and down-shift are computed using independently calibrated local correlation models from market data with just shifted spots. This includes most effects changing the correlation under a spot variation. But, of course it is more expensive and presumably less stable from a numerical point of view.

(sticky correlation) In this case we reuse the random number stream generated for the unshifted prices also for the spot-up or down shifted simulations. Of course, we miss some of the correlation effects on the delta. In a stochastic correlation model were the Brownian motions for the spot and the correlation processes are only coupled by a correlation we have a similar situation. The advantage of using sticky correlations is that we may also employ so-called path-recycling (Brockhaus 2016) where paths of unshifted constituents are

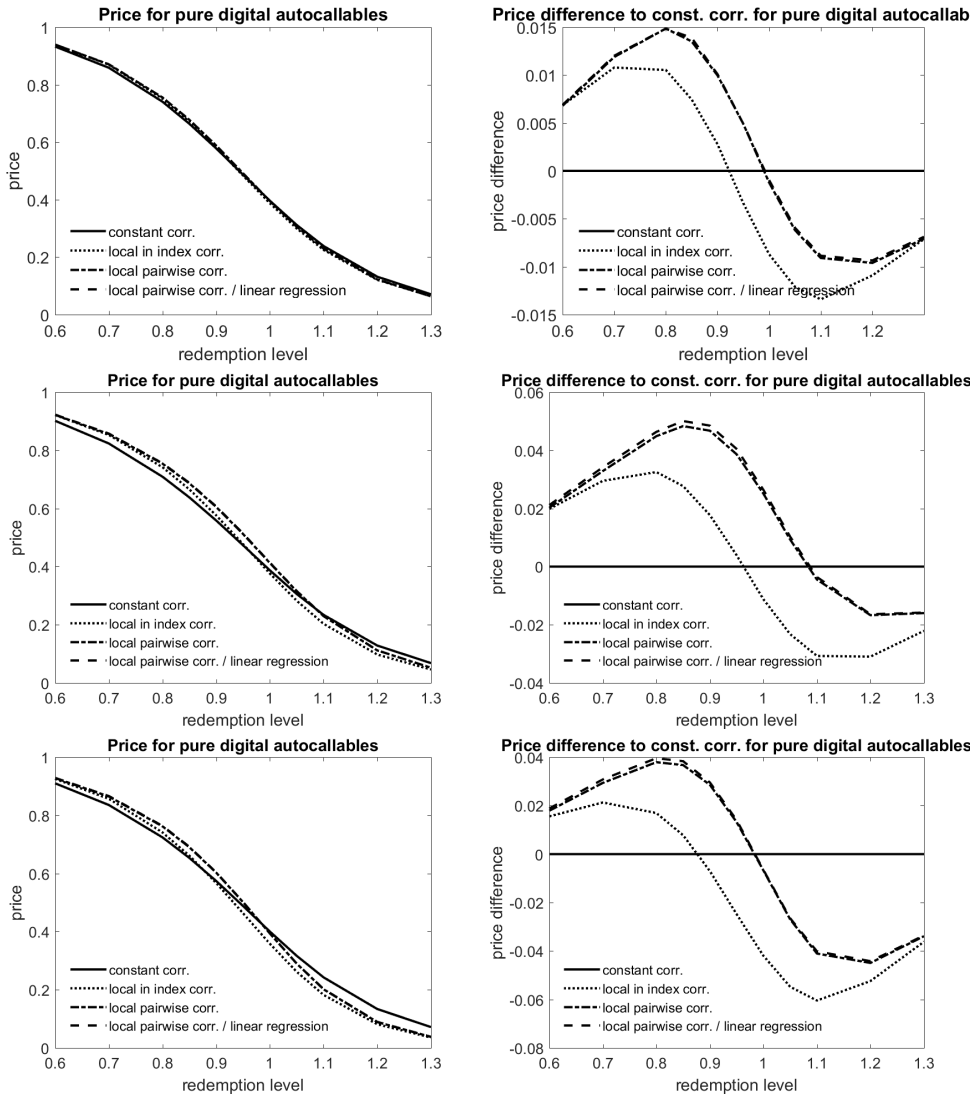

Figure 7: Prices for auto-callable(worst-of) instruments vs redemption level (left). Price differences to constant correlation model (right). 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom) data.

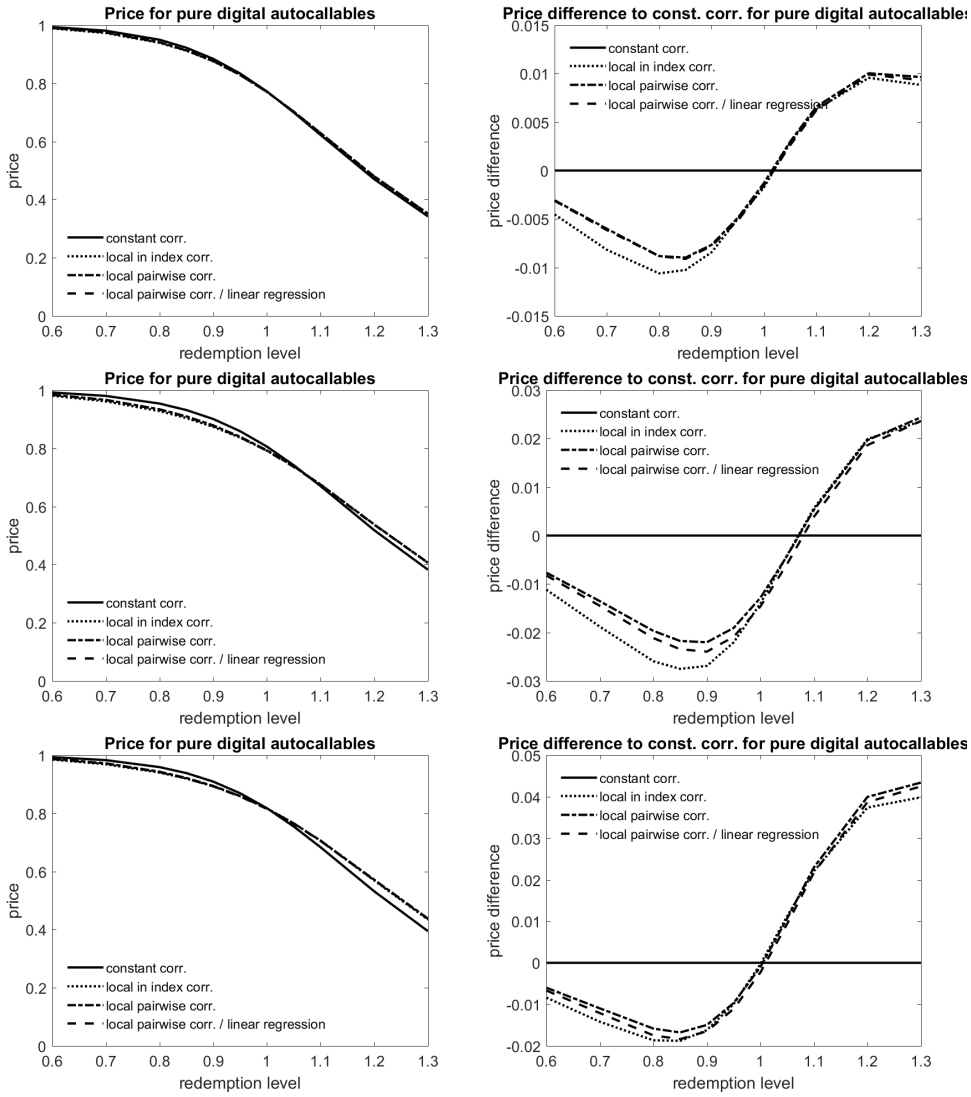

Figure 8: Prices for auto-callable(best-of) instruments vs redemption level (left). Price differences to constant correlation model (right). 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom) data.

not simulated again. This speeds up the computation of the delta vector by a factor of roughly d and of the gamma matrix by roughly $d^2/2$ where d is the dimension of the underlying basket, i.e. three in our example.

The different deltas (precisely, delta times spot of ADS) of our test instruments are shown in Figure 9 (left). The differences to the constant correlation model are shown on the right. The differences between the delta variants for each model are rather small, at least smaller than the differences between the two local correlation models. Another observation is that the recalibration delta is more noisy than the sticky correlation delta which is expected, but it may still be an option.

(Computing times) One of the reasons why we are interested in an alternative to the local in index model are computing times for pricing instruments. The local in index model can be implemented such that no extra Cholesky decompositions are needed during calibration or pricing, either by means of rank-1 updates or by means of so-called dimension expansion (Reghai 2010). This is not possible for the pairwise correlation model. Here we always need a Cholesky decomposition in each time step. For that the pairwise model can not be faster for all instruments. Table 1 shows the cpu times for pricing worst-of auto-callables with a basket dimension which varies between 2 and 10. In all cases the Monte Carlo scheme used 100K simulations and 182 time steps. The prices were computed for the 2015/02/15 and the 2017/01/04 market data. For the constant correlation model and the local in index model the market data has no impact on cpu times. But for the pairwise correlation model market data has an effect because of the number of required salvaging operations.

6 Conclusions

In this paper we have developed a new local correlation model which uses a generic function g to describe the correlation between all asset-asset pairs of a basket of underlyings. This approach has the advantage that one does not need to simulate the complete index basket when pricing options on small sub baskets and leads to significant speed ups when pricing small and medium sized baskets. Moreover, the model does not suffer from the chewing gum effect which describes the insensitivity of the correlation between two stocks when their movement is offset by other index constituents. For worst-of instruments this leads to consistently higher prices in comparison to the

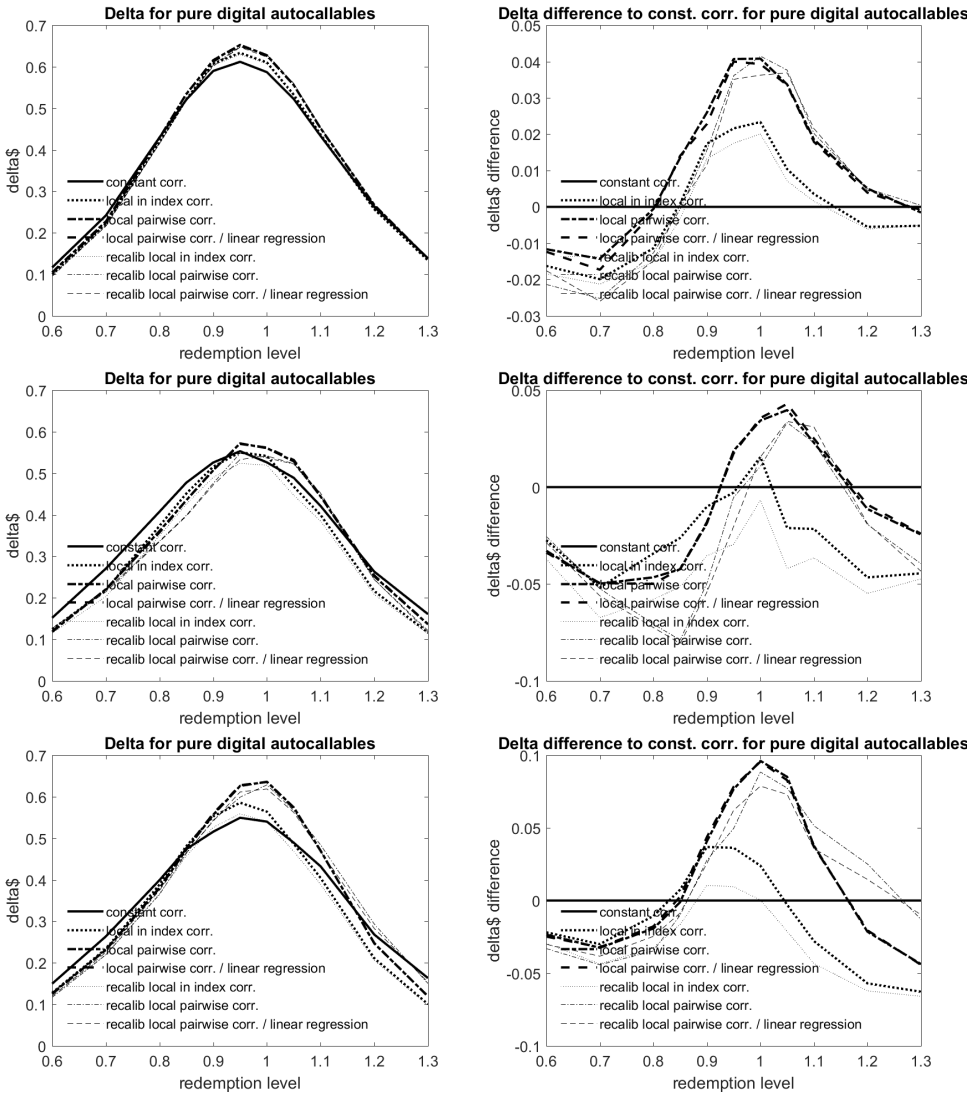

Figure 9: Delta for auto-callable instruments vs redemption level (left). Delta differences to constant correlation model (right). 2015/02/15 (top), 2017/01/04 (middle) and 2017/02/20 (bottom) data.

dim	const corr.	pairwise loc. corr.		loc. in index
2	1.3	3.6	3.6	17.1
3	1.8	4.9	4.9	”
4	2.3	6.2	6.2	”
5	3.0	7.7	7.7	”
6	3.4	9.4	9.4	”
7	3.8	11.2	11.5	”
8	4.3	13.1	13.9	”
9	5.1	15.2	17.7	”
10	5.3	17.6	21.2	”

Table 1: Cpu times [s] for pricing a worst-of auto-callable instrument with 100K simulations and 182 time steps. The times for the pairwise model correspond to the 2015/02/15 and the 2017/01/04 market data.

local in index model.

Like for some stochastic correlation models, one may also use the calibrated local correlation function for pricing instruments on baskets which are not subsets of the equity index to which the model was calibrated. The fitted local correlation function is basically a generic functional relationship between any pairwise correlation and the respective two constituent spot levels.

The calibration of the pairwise model uses the particle method of (Guyon 2014) but is slower because it is much harder to enforce positive definiteness. The fitting quality is similar to the local in index model. By means of experiments we have also shown that the fitting quality is still good using linear or constant regressions of the fitted local correlation functions, but the number of parameters is massively reduced by that. Open questions concern surely the fitting quality on the down side for high skew index implied volatilities and whether common single underlying stochastic volatility models are sufficient to improve this or if one needs the general approach of (Guyon 2015). We also think that anchoring (Wilmott et al. 2014), (Guyon 2014b) is very interesting for both, the single underlying processes and the moneyness scales for the pairwise correlation model.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

References

[1] A. Ahdida (2011) *Processus matriciels : simulation et modelisation de la dependance en finance*, Ph.D. dissertation, Universite Paris-Est, English. NNT : 2011PEST1154. pastel-00674813

[2] A. Ahdida and A. Alfonsi (2013) *A mean-reverting SDE on correlation matrices*, Stochastic Processes and their Applications, Vol. 123 (4), pp. 1472-1520. doi: 10.1016/j.spa.2012.12.008

[3] S. Bossu (2014) *Local Correlation, in Advanced Equity Derivatives: Volatility and Correlation*, John Wiley & Sons, Inc, Hoboken, doi: 10.1002/9781118835364.ch8

[4] O. Brockhaus (2016) *Equity Derivatives and Hybrids. Markets, Models and Methods*, Palgrave Macmillan.

[5] H. Buehler (2009) *Volatility and dividends - Volatility modelling with cash dividends and simple credit risk*, Available at SSRN: http://ssrn.com/abstract=1141877

[6] S. Corlay (2016) *B-spline techniques for volatility modeling*, Journal of Computational Finance 19(3), pp. 97-135.

[7] P. Delanoe (2014) *Local Correlation with Local Vol and Stochastic Vol*, Available at SSRN: http://ssrn.com/abstract=2454668

[8] Deutsche Börse AG (2014) *Leitfaden zu den DAX-global(R)-Indizes der Deutsche Börse AG Version 2.20*.

[9] L. Di Gaspero (2017) QuadProg++ software.

[10] B. Dupire (1994) *Pricing with a Smile*, Risk Magazine, Available at http://www.risk.net/data/risk/pdf/technical/2007/risk20_0707_technical_volatility.pdf

[11] J. Gatheral (2006) *The Volatility Surface: A Practitioners Guide*, John Wiley & Sons.

[12] J. Gatheral and A. Jacquier (2014) *Arbitrage-Free SVI Volatility Surfaces*, Quantitative Finance, Vol. 14, No. 1, pp. 59-71.

[13] D. Goldfarb and A. Idnani (1983) *A numerically stable dual method for solving strictly convex quadratic programs*, Mathematical Programming, Vol. 27, pp. 1-33.

[14] I. Gyöngy (1986) *Mimicking the One-Dimensional Marginal Distributions of Processes Having an Ito-Differential*, Probability Theory and Related Fields, No. 71, pp. 501-516.

[15] J. Guyon (2014) *Local correlation families*, Risk Magazine, (February 2014).

[16] J. Guyon (2014b) *Path-Dependent Volatility*, Available at SSRN: http://dx.doi.org/10.2139/ssrn.2425048

[17] J. Guyon (2015) *Cross-Dependent Volatility*, Available at SSRN: http://dx.doi.org/10.2139/ssrn.2615162.

[18] J. Guyon (2016) *Calibration of local correlation models to basket smiles*, Journal of Computational Finance 21(1), pp. 151, doi: 10.21314/JCF.2016.326

[19] S. Heston (1993) *A Closed-Form Solution for Options with Stochastic Volatility with Application to Bond and Currency Options*, The Review of Financial Studies, Vol. 6, pp. 327-343.

[20] N. Higham (2002) *Computing the nearest correlation matrix – a problem from finance* IAM Journal of Numerical Analysis, Vol 22., pp. 329-344.

[21] B. Jourdain and M. Sbai (2009) *Coupling Index and stock*, Available at SSRN: http://ssrn.com/abstract=1320164.

[22] P. Jäckel (2002) *Monte-Carlo Methods in Finance*, Wiley.

[23] P. Jäckel and C. Kahl (2009) *Positive semi-definite correlation matrix completion for stochastic volatility models*, Available at www.jaeckel.org

[24] A. Langnau (2010) *A dynamic model for correlation*, Risk Magazine, (April 2010), pp. 74-78.

[25] V.Piterbarg (2006) *Markovian Projection Method for Volatility Calibration*, Barclays Capital.

[26] A. Reghai (2010) *Breaking correlation breaks*, Risk Magazine, (October 2010), pp. 92-97.

[27] P. Wilmott, A. Lewis, A. and D.J. Duffy (2014) *Modeling Volatility and Valuing Derivatives Under Anchoring*, Wilmott Magazine, (September 2014), pp. 4857. doi: 10.1002/wilm.10366

[28] V. Zetocha (2015) *Correlation skew via stochastic correlation and jumps.* Risk Magazine, (December 2015).