



HAL
open science

Local volatility calibration using an adjoint proxy

Gabriel Turinici

► **To cite this version:**

| Gabriel Turinici. Local volatility calibration using an adjoint proxy. 2008. hal-00306187v1

HAL Id: hal-00306187

<https://hal.science/hal-00306187v1>

Preprint submitted on 25 Jul 2008 (v1), last revised 5 Dec 2008 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Local volatility calibration using an adjoint
proxy

Gabriel Turinici

CEREMADE, Université Paris Dauphine,

Place du Marechal de Lattre de Tassigny, 75016 Paris, France

`gabriel.turinici@dauphine.fr`

+ 33 1 44 05 48 58

July 24, 2008

Local volatility calibration using an adjoint proxy

Abstract

We document the calibration of the local volatility in a framework similar to Coleman, Li and Verma. The quality of a surface is assessed through a functional to be optimized; the specificity of the approach is to separate the optimization (performed with any suitable optimization algorithm) from the computation of the functional where we use an adjoint (as in L. Jiang et. al.) to obtain an approximation; moreover our main calibration variable is the implied volatility (the procedure can also accommodate the Greeks). The procedure performs well on benchmarks from the literature and on FOREX data.

Keywords: calibration, local volatility, implied volatility, Dupire formula, adjoint

1 Motivation: the local volatility surface

Let us consider a security S_t (e.g. a stock, a FOREX rate, etc.) whose price, under the risk-neutral [Musielà and Rutkowski, 2005, Hull, 2006] measure, follows the stochastic differential equation

$$dS_t/S_t = r(t)dt + \sigma dW_t \tag{1}$$

with $r(t)$ being the time dependent risk-free rate and σ the volatility (we will make explicit its dependence later) and W_t a Brownian motion.

Let us consider (for now) plain vanilla call options contingent on S_t and recall that when the volatility (and the discount rate r) are constant the Black-Scholes model [Black and Scholes, 1973] gives a closed formula for the price $C(S, t)$ of such claims. It is standard to note that the reverse is also true, i.e., provided r is constant and known, from the observed market prices denoted C_{K_l, T_l}^{market} (with strikes K_l and maturities T_l , $l = 1, \dots, L$) one can find (i.e. calibrate) the unique *implied volatilities* σ_{K_l, T_l}^I that, when introduced in the Black-Scholes formulae, match the observed market prices C_{K_l, T_l}^{market} . However the implied volatilities σ_{K_l, T_l}^I thus obtained are not the same for all K_l and T_l (the *smile* effect) which is inconsistent with the initial model. To address this issue it was independently proposed by Rubinstein [Rubinstein, 1994], Dupire [Dupire, 1994] and Derman and Kani [Derman and Kani, 1994] to take the volatility σ as depending on the time and the security price S : $\sigma = \sigma(S, t)$; the model is named *local volatility*. Historically the proposals in [Rubinstein, 1994, Derman and Kani, 1994] build on the Cox-Ross-Rubinstein binomial tree [Cox et al., 1979] and are described as *implied trees*.

Let us make clear that we do not discuss here the local volatility model itself nor its dynamics. We only see the local volatility as a way to express

the non-arbitrage relationships between the set of derivatives contracts contingent on the same (set of) underlying instruments (much similar to the way one uses the *risk neutral* probability measure as a tool to compute prices but does not necessarily want to assign it to any real world probabilities).

Matching the observed prices, i.e. calibrating the local volatility $\sigma(S, t)$ is not straightforward as no closed formula exists to express the dependence $\sigma \rightarrow C$. The problem becomes now an inverse problem [Bouchouev and Isakov, 1997, 1999].

When the number of quoted market prices C_{K_l, T_l}^{market} is large enough (i.e. K_l, T_l cover well the range of S and t) the local volatility can be expressed using the Dupire formula [Dupire, 1994, Hull, 2006, Achdou and Pironneau, 2005] or different asymptotics [Berestycki et al., 2002]. However, when only a few prices are known, the Dupire formula is less effective and other methods have to be used [Avellaneda et al., 1997, Bodurtha and Jermakyan, 1999]. Among those, Coleman, Li & Verma [Coleman et al., 2001] introduced a parametric procedure which we refine in this contribution. Further, L. Jiang, and co-authors established a mathematical grounding for formulating this problem as a control problem [Jiang et al., 2003]; we will retain in this paper the adjoint state technique that we adapt to take into account the constraints (see [Lagnado and Osher, 1997, 1998] for related endeavors). Our procedure combines the approaches above and is accelerated by the use of an approxi-

mation of the functional through the use of the adjoint (7). A particularity of the procedure is to calibrate directly the implied volatility (and can accommodate any Greeks); this choice enhance not only the efficiency of the numerical procedure but, in some extreme cases, its selection of adequate local surface as was confirmed in numerical experiments. This approach (rather natural since option traders often only quote the implied volatility and not the price) is especially useful in markets that heavily rely on Greeks (as is the case in the FOREX market that quotes *risk reversals* which involve Deltas and the implied volatility. Further, since in general only limited data is available, the local surface is non-unique: to eliminate improper candidates we set lower and upper bounds on the volatility. The resulting procedure is stable with respect to the number of price information used and in particular no interpolation is required to fill this information when missing.

2 Adjoint formulas and the cost functional

Under the local volatility model, the price $C(S, t)$ of a derivative contract on S_t with pay-off $h(S)$ at maturity $t = T$, will satisfy the (Black-Scholes) equation [Hull, 2006] for all $S \geq 0$ and $t \in [0, T]$:

$$\partial_t C + r \partial_S C + \frac{\sigma^2 S^2}{2} \partial_{SS} C - rC = 0 \quad (2)$$

$$C(S, t = T) = h(S) \quad (3)$$

Remark 1 *Similar considerations apply if the security S_t distributes dividends at a known proportional rate $q(t)$ or if S_t is a FOREX spot (in this case r is the domestic discount rate and $q(t)$ is the foreign rate).*

The price at $t = 0$ of the contract is $C(S_{t=0}, t = 0)$; recall that the payoff of an European call of strike K is $h(S) = (S - K)_+$ (with the notation $x_+ = \max\{x, 0\}$). Note the retrograde nature of the equation (2)-(3).

We will use the technique of the adjoint state and view the price as a implicit functional of σ (here δ is the Dirac operator):

$$C(t = 0; S = S_0) = \langle \delta_{t=0, S=S_0}, C(S, t) \rangle. \quad (4)$$

Then the variation $\frac{\delta C}{\delta(\sigma^2)}$ of C with respect to σ^2 (and respectively the variation with respect to σ) will be

$$\frac{\delta C}{\delta(\sigma^2)} = \frac{S^2}{2} (\partial_{SS} C) \chi, \quad (5)$$

$$\frac{\delta C}{\delta \sigma} = 2\sigma \frac{S^2}{2} (\partial_{SS} C) \chi. \quad (6)$$

Here the adjoint state χ is the solution of:

$$\partial_t \chi + \partial_S (rS\chi) + \partial_{SS} \left(\frac{\sigma^2 S^2}{2} \chi \right) + r\chi = 0 \quad (7)$$

$$\chi(S, t = 0) = \delta_{t=0, S=S_0} \quad (8)$$

Same technique works for any other quantity dependent on the price. A very important example of such quantity is the implied volatility, denoted

here σ^I . Recall that an explicit formula links the price to the implied volatility $\sigma^I = \sigma^I(C)$ and as such $\frac{\partial \sigma^I}{\partial \sigma} = \frac{\partial \sigma^I}{\partial C} \frac{\partial C}{\partial \sigma}$. We recognize in the term $\frac{\partial \sigma^I}{\partial C}$ the inverse of the Black-Scholes vega, that we will denote ν^I . We obtain

$$\frac{\partial \sigma^I}{\partial \sigma} = \frac{1}{\nu^I} \frac{\partial C}{\partial \sigma}. \quad (9)$$

Remark 2 *Both problems (2) and (7) can be solved e.g. through a Crank-Nicholson finite-difference scheme [Hull, 2006, Andersen and Brotherton-Ratcliffe, 1998]; is is best to use for (7) the numerical adjoint of (2).*

To illustrate the nature of this gradient we display an example in Figure 1 where we note two singularities appearing in $(t = 0, S = S_0)$ (from eqn (8)) and $(t = 1, S = K)$ (from $\partial_{SS}(S - K)_+$ (see also [Avellaneda et al., 1997] for similar conclusions).

Since in general several option prices (or Greeks) are available and have to be accounted in the calibration, we introduce a cost functional (depending on σ) which is the sum of relative errors of the prices computed with a given σ and the market prices. Moreover, depending on the market (e.g. the FOREX market quotes risk-reversals in terms of implied volatility and deltas directly) one would also want to fit the implied volatility. Of course, if a perfect calibration is achieved, both results will give the same implied volatility; in practice fitting the implied volatility in addition or instead of the prices give better numerical stability of the procedure. Numerical tests

(not shown here) display, for the FOREX market, a clear improvement in the calibration quality when the implied volatilities are used instead of just prices.

The cost functional so far is

$$\eta_1 \sum_{l=1}^L \left(\frac{C_l(0; S_0)}{C_{K_l, T_l}^{market}} - 1 \right)^2 + \eta_2 \sum_{l=1}^L \left(\frac{\sigma^I(K_l; T_l)}{\sigma_{K_l, T_l}^{I; market}} - 1 \right)^2. \quad (10)$$

Here $\sigma_{K_l, T_l}^{I; market}$ is the market implied volatility while $\sigma^I(K_l; T_l)$ is the implied volatility corresponding to the local volatility σ ; η_1 and η_2 are some positive weights.

Repeated application of the chain rule and the formulas (6) and (9) allow to compute the variation $\frac{\delta J_e}{\delta \sigma}$ of the J_e with respect to σ . Note that for each index l one needs to solve a PDE for the price C_l and a corresponding PDE for the adjoint χ_l and use them as in (6).

Remark 3 *Other forms of the cost functional can also be treated, for instance the distances*

$$\sum_{l=1}^L (C_l(0; S_0) - C_{K_l, T_l}^{market})^2. \quad (11)$$

or, when bid/ask quotes are available, i.e. $C_l(0; S_0) \in [C_{K_l, T_l}^{bid}, C_{K_l, T_l}^{ask}]$ one can use as in [Coleman et al., 2001]

$$\sum_{l=1}^L [(C_l(0; S_0) - C_{K_l, T_l}^{bid})_+]^2 + [(C_{K_l, T_l}^{ask} - C_l(0; S_0))_+]^2. \quad (12)$$

Remark 4 *A naive approach is to use a standard optimization algorithm [Bonans et al., 2006]; for instance, a fixed step ($\rho > 0$) gradient algorithm would read:*

$$\sigma_{n+1} = \sigma_n - \rho \frac{\partial J_e}{\partial \sigma}(\sigma_{n+1}). \quad (13)$$

In this case the singularities of $\frac{\partial J_e}{\partial \sigma}$ will propagate into the solution which will have a full list of singularities at $(0, S_0)$ and (T_l, K_l) , $l = 1, \dots, L$. Such properties are not natural for the local volatility surface $\sigma(t, S)$ and the inversion procedure has to address them. Note that obtaining a smoother local surface is possible because of its underdetermination : in the extreme situation $L = 1$ only one price C_{K_l, T_l}^{market} is available which brings a limited information on the volatility surface that will not be unique; in this case the most natural volatility surface will a constant, equal to the Black-Scholes implied volatility.

A traditional choice to avoid singularities and address the non-uniqueness is to parametrize the surface $\sigma(S, t)$ [Achdou and Pironneau, 2005, Coleman et al., 2001]; the result will be the optimal surface in the class.

In order to ensure smoothness we add to the cost functional terms that avoid large variations of σ by penalizing its gradient with respect to S and t (η_3 and η_4 are positive weights):

$$\eta_3 \left\| \frac{\partial \sigma(S, t)}{\partial S} \right\|_{L^2_{S,t}}^2 + \eta_4 \left\| \frac{\partial \sigma(S, t)}{\partial t} \right\|_{L^2_{S,t}}^2 \quad (14)$$

(recall that $\|F(x)\|_{L_x^2}^2 = \int F^2(x)dx$). The final cost functional is

$$\begin{aligned}
J_e(\sigma) = & \eta_1 \sum_{l=1}^L \left(\frac{C_l(0; S_0)}{C_{K_l, T_l}^{market}} - 1 \right)^2 + \eta_2 \sum_{l=1}^L \left(\frac{\sigma^I(K_l; T_l)}{\sigma_{K_l, T_l}^{I; market}} - 1 \right)^2 \\
& + \eta_3 \left\| \frac{\partial \sigma(S, t)}{\partial S} \right\|_{L_{S, t}^2}^2 + \eta_4 \left\| \frac{\partial \sigma(S, t)}{\partial t} \right\|_{L_{S, t}^2}^2.
\end{aligned} \tag{15}$$

3 Surface space and the optimisation procedure

Continuing the arguments of the previous section, we give here a possible choice to describe the space of available surface shapes. We consider continuous affine functions with degrees of freedom being the values on some grid ($S_i = S_0 + i\Delta S, t_j = t_0 + j\Delta t$), $i \leq I, j \leq J$. We denote by $f_{ij}(S, t)$ the unique piecewise linear and continuous function that has value of 1 at (t_i, S_j) , and is zero everywhere else. The surfaces are linear combinations

$$\sigma(S, t) = \sum \alpha_{ij} f_{ij}(S, t). \tag{16}$$

of the shapes $f_{ij}(S, t)$.

The advantage of linear interpolation is that the shape functions have nice localisation properties: the scalar product of two such functions (or their gradient) is zero except if they are neighbors i.e. matrices (22)-(23) are sparse. Also setting constraints e.g. $\sigma(S, t) > \sigma_{min}$ for all S, t is equivalent to asking that all α_{ij} are larger than σ_{min} .

However we also tested cubic splines interpolation and it performed equally satisfactory.

Remark 5 *A possible procedure would be to optimize the cost functional (15) expressed as a function of the coefficients α_{ij} of σ in (16). But this dependence may be highly nonlinear and the resulting optimization will have many unwanted local extrema.*

Chain rule gives the gradient of any derivative contract $C(S, t)$ (among C_l , $l = 1, \dots, L$) with respect to variations of the local surface σ inside the admissible surface space. This is in fact just a matter of projecting the exact gradient (6) onto each shape f_{ij} . We obtain an approximation formula around the current local volatility σ :

$$C \left(\sigma + \sum_{ij} \alpha_{ij} f_{ij}(S, t) \right) \simeq C(\sigma) + \sum_{ij} \left\langle \frac{\partial C}{\partial \sigma}, f_{ij} \right\rangle_{L^2_{S,t}} \alpha_{ij}. \quad (17)$$

Same works for the implied volatility

$$\sigma^I \left(\sigma + \sum_{ij} \alpha_{ij} f_{ij}(S, t) \right) \simeq \sigma^I(\sigma) + \sum_{ij} \left\langle \frac{\partial \sigma^I}{\partial \sigma}, f_{ij} \right\rangle_{L^2_{S,t}} \alpha_{ij}. \quad (18)$$

In discrete formulation the cost functional will employ the matrices

$$M_{ij;rs}^C = \sum_l \left\langle \frac{\partial C_l}{\partial \sigma}, f_{ij} \right\rangle_{L^2_{S,t}} \left\langle \frac{\partial C_l}{\partial \sigma}, f_{rs} \right\rangle_{L^2_{S,t}} \quad (19)$$

for the first part of (10) and

$$M_{ij;rs}^\sigma = \sum_l \left\langle \frac{\partial \sigma_l^I}{\partial \sigma}, f_{ij} \right\rangle_{L^2_{S,t}} \left\langle \frac{\partial \sigma_l^I}{\partial \sigma}, f_{rs} \right\rangle_{L^2_{S,t}} \quad (20)$$

for the second part.

Note that (18) and (17) already provide (some) second order information for J_e ; also note that for $\sigma = \sum_{ij} \beta_{ij} f_{ij}(S, t)$ the smoothness terms (14) can be written as

$$\eta_3 < \beta + \alpha, Q_S(\beta + \alpha) > + \eta_4 < \beta + \alpha, Q_S(\beta + \alpha) > \quad (21)$$

with

$$(Q_S)_{ij;kl} = \int \int \frac{\partial f_{ij}(S, t)}{\partial S} \frac{\partial f_{kl}(S, t)}{\partial S} dS dt \quad (22)$$

and

$$(Q_t)_{ij;kl} = \int \int \frac{\partial f_{ij}(S, t)}{\partial t} \frac{\partial f_{kl}(S, t)}{\partial t} dS dt. \quad (23)$$

A last ingredient involves bounds on the local volatility surface; indeed, it seems natural that the local volatility cannot be negative. Even when this is the case, local volatilities with very low values (e.g. 3% !) are obviously not realistic. Enforcing constraints on the local volatilities is a very important step towards selecting meaningful candidates. A choice that is consistent with other observations in the literature [Rubinstein, 1994, Derman and Kani, 1994] is to ask

$$\sigma_{min} \leq \sigma(t, S) \leq \sigma_{max} \text{ with}$$

$$\sigma_{min} = \frac{1}{2} \min\{\sigma_{K_l, T_l}^{I;market}; l = 1, \dots, L\}, \quad \sigma_{max} = 2 \max\{\sigma_{K_l, T_l}^{I;market}; l = 1, \dots, L\} \quad (24)$$

3.1 Optimization procedure

The algorithm operates as follows: first we choose as initial guess σ_0 to be the (projection on the space $Vect\{f_{ij}\}$) of the implied volatility surface (eventually corrected to be between bounds σ_{min} and σ_{max}). One can also use more specific formulas relating implied and local volatility see e.g. [Gatheral, 2006] formula 1.10 page 13.

The iterative procedure operates at each step in the following order:

1/ computes the gradient of the price and implied volatility with respect to variations of σ in the admissible space i.e. formula (17) and (18);

2/ constructs and solves the (quadratic) optimization problem

$$\min_{\sigma_{min} \leq \alpha \leq \sigma_{max}} \frac{1}{2} \langle \alpha^T, (\eta_1 M^C + \eta_2 M^\sigma + \eta_3 Q_S + \eta_4 Q_t) \alpha \rangle + w^t \alpha \quad (25)$$

3/ update the local volatility σ ; if the replication error J_e is too high return in 1/ otherwise exit.

In practice very few cycles 1/-3/ are necessary. We tested on several indices and in the FOREX markets and the numbers varied between 5 and 10 cycles.

Remark 6 *The quadratic problem (25) can be solved by any suitable algorithm; for instance Matlab uses by default a subspace trust-region method based on the interior-reflective Newton method described in [Coleman and*

Li, 1996]. We also tested a simple projected gradient which performed very satisfactory. The advantage of the approach is precisely to separate the optimization itself from the formulation of the problem.

Remark 7 *Should a bid/ask functional (e.g. as in (12)) be used then the problem will not be quadratic any more but (17) is still used; the constraints arise from the requirement that α_k be in $[\sigma_{min}, \sigma_{max}]$; additional constraints, in a "trust-region" style, can be put to remain in a region where the approximation (17) holds.*

4 Results and conclusions

A specificity of the approach is that instead of a unique optimization in the parametric space we perform one optimization around each current point; this reduces the number of computations of the PDE (2). But, equally importantly, the separation between the optimization and the approximation of the functional provides flexibility in the information that can be fitted, e.g. we can readily accommodate any derivative contract (as soon as an gradient formula like (5) exists for it; when it does not one can use Malliavin calculus) such as options on futures, strategies, structured products etc. This allows for instance to be very flexible in the information available and to ignore some prices should them not be available or if one wants to arbitrage against

them (in contrast with the pioneering approaches [Rubinstein, 1994, Dupire, 1994, Derman and Kani, 1994] that need a uniform set of data to perform the inversion); in particular no interpolation is required to fill this information when missing.

The use of the gradient not in an optimization procedure but to obtain an approximation of the functional around the current point is a acknowledgement of the fact that the main difficulty is not finding a solution but choosing one among all compatible surfaces (i.e. ill-posedness).

We noted that in practice the implied volatility term in the cost functional i.e. $\eta_2 M^\sigma$ in (25) is more helpful to orient the optimization procedure than the price term $\eta_1 M^C$. In fact in all cases we tested putting $\eta_2 \neq 0$ and $\eta_1 = 0$ gave better results than the reverse.

We used throughout a grid with $I = 24$ values of S and $J = 13$ values of t i.e. 312 shapes f_{ij} , cf. eqn. (16).

Let us now iterate through several benchmarks from the literature; we begin with the European call data on the S& P index from [Andersen and Brotherton-Ratcliffe, 1998, Coleman et al., 2001]. Similar to [Andersen and Brotherton-Ratcliffe, 1998, Coleman et al., 2001], we use only the options with no more than two years maturity in our computation. The initial index, interest rate and dividend rate are the same. We first checked (not shown) that for $L = 1$ the problem recovers the implied volatility; it did so with only

one cycle. When we took all the $L = 70$ data the resulting local volatility surface is given Figure 3.

We next moved to a FOREX example (from [Avellaneda et al., 1997]) where synchronous option prices (based on bid-ask volatilities and risk-reversals) are provided for the USD/DEM 20,25 and 50 delta risk-reversals quoted on August 23rd 1995. The results in Figure 4 show a very good fit quality with only five cycles 1/-3/.

We remain in the FOREX market and take as the next example 10,25 and 50-Delta risk-reversal and strangles for EUR/USD dated March 18th 2008. We recall that e.g. a 25 Delta risk reversal contract consists in a long position in a call option with $\text{delta}=0.25$ and a short position in a put option with $\text{delta} = -0.25$; the contract is quoted in terms of the difference of the implied volatilities of these two options. Note that at no moment the price of the options appear in the quotes. In order to set the input implied surface we used 10 and 25 Delta strangles which are quoted as the arithmetic mean of the implied volatilities of the two options above. Of course, from this data one can next recover the implied volatilities of each option, then all other characteristics. We present in Figure 5 the implied and the calibrated local volatility from the data in Tables 1,2 and 3. The procedure was also tested (not shown here) on other currencies pairs (GBP, CHF, JPY, KRW, THB, ZAR all with respect to USD) and performed well.

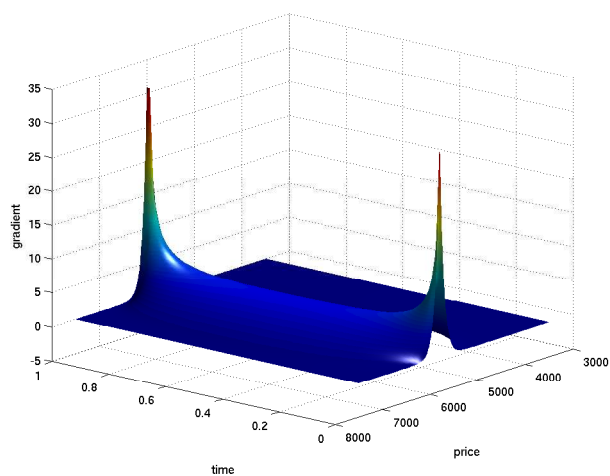


Figure 1: Gradient $\frac{\delta C}{\delta(\sigma^2)}$ (see eqn. (5)) of the price C of a derivative (e.g. a plain vanilla call) with respect to the volatility surface squared σ^2 . Note the two singularities at the initial time (around the spot price) and at the expiration around the strike. These singularities prevent the direct use of any gradient method otherwise the resulting surface will be singular.

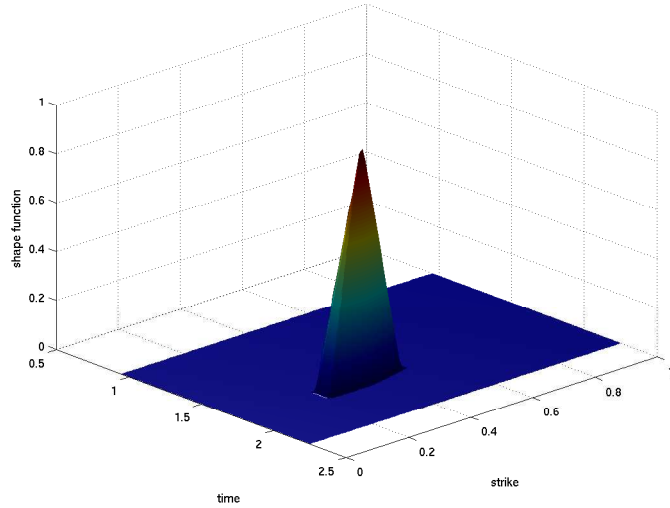


Figure 2: The local volatility $\sigma(S, t)$ is sought after as a linear combination of basic shapes $f_{ij}(S, t)$: $\sigma(S, t) = \sum_{ij} \alpha_{ij} f_{ij}$. A possible option is to take $f_{ij}(S, t)$ as the (unique) linear interpolation which is zero except in some point (S_i, t_j) (part of a grid in S and t). We display here such a shape.

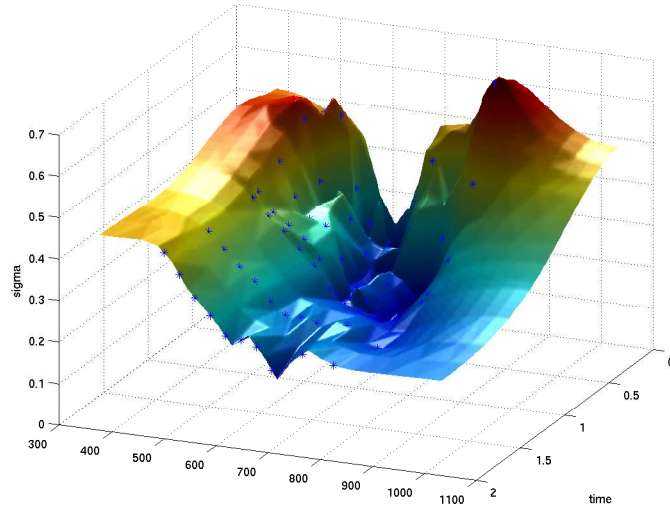


Figure 3: Local volatility surface of the S&P 500 index as recovered from the published European call options data [Andersen and Brotherton-Ratcliffe, 1998, Coleman et al., 2001]; spot price is \$590; discount rate $r = 6\%$, dividend rate 2.62% . The blue marks on the surface indicate the option prices that were used to invert i.e. the K_l and T_l ($L = 70$). After 10 iterations the prices are recovered up to $4.e - 4$ and the implied volatility up to 0.18% . Setting regularization parameters η_3 and η_4 to smaller values give better fit but less smooth surfaces.

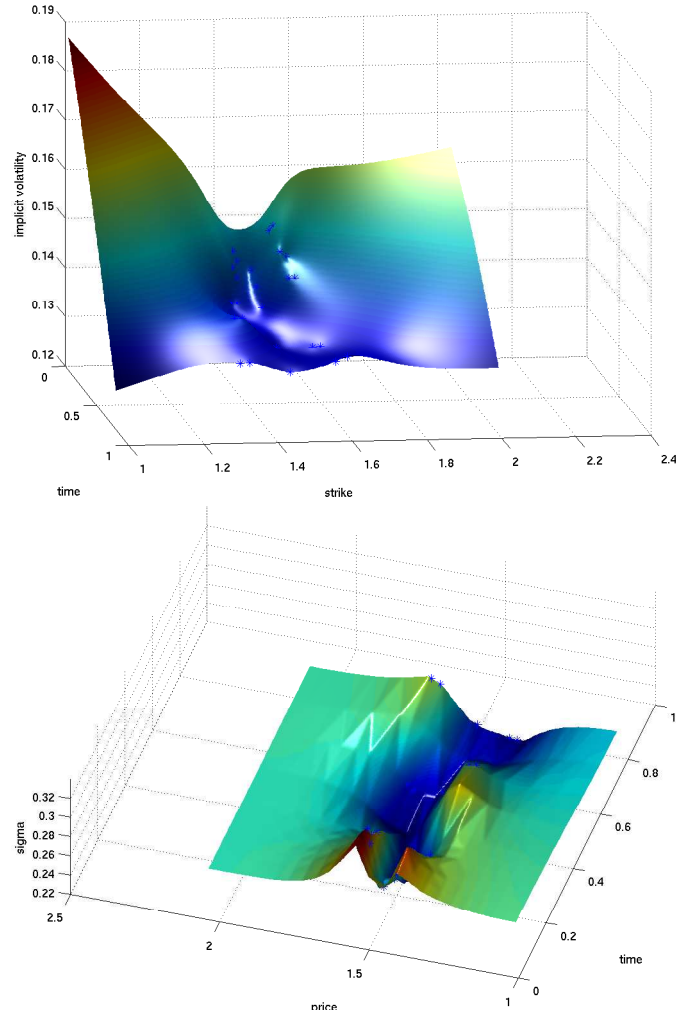


Figure 4: Top: implied volatility surface of the USD/DEM rate from [Avellaneda et al., 1997]; blue marks on the surface represent the available prices (to be matched). Bottom: local volatility surface as recovered from quoted 20,25 and 50-delta risk-reversals [Avellaneda et al., 1997]; (mid) spot price is 1.48875; *USD* discount rate $r = 5.91\%$, and *DEM* rate 4.27%. The blue marks on the surface indicate the option prices that were used to invert i.e. the K_l and T_l ($L = 30$). After 5 iterations the prices are recovered up to $3.e - 4$ (below the PDE resolution) and the implied volatility up to 0.11% (below the bid/ask spread).

Delta	0,1	0,25	0,5	0,75	0,9
Days to Expiry					
7	1,6177	1,5965	1,5753	1,5544	1,5341
31	1,6564	1,6150	1,5740	1,5335	1,4935
59	1,6804	1,6253	1,5720	1,5191	1,4653
92	1,7013	1,6333	1,5686	1,5042	1,4368
184	1,7449	1,6474	1,5592	1,4711	1,3728
365	1,8030	1,6611	1,5391	1,4164	1,2665

Table 1: Strikes of the EUR/USD data derived from March 18th 2008 10,25 and 50 Delta risk-reversals and straddles corresponding to results in Figure 5.

Delta	0,1	0,25	0,5	0,75	0,9
Days to Expiry					
7	14,8250%	14,1750%	13,9250%	14,1750%	14,8250%
31	13,5250%	12,9500%	12,7750%	13,1000%	13,8250%
59	12,7750%	12,1500%	12,0250%	12,4000%	13,2750%
92	12,4250%	11,7500%	11,6250%	12,1000%	13,1250%
184	12,0875%	11,2125%	11,0500%	11,6375%	12,9625%
365	11,9125%	10,8625%	10,7000%	11,3375%	12,8875%

Table 2: Implied volatilities of the EUR/USD data derived from March 18th 2008 10,25 and 50 Delta risk-reversals and straddles corresponding to results in Figure 5.

Delta	0,1	0,25	0,5	0,75	0,9
Days to Expiry					
7	0,0015	0,0046	0,0121	0,0253	0,0425
31	0,0029	0,0088	0,0231	0,0486	0,0824
59	0,0038	0,0113	0,0298	0,0632	0,1088
92	0,0046	0,0136	0,0358	0,0767	0,1341
184	0,0063	0,0182	0,0479	0,1039	0,1883
365	0,0086	0,0247	0,0650	0,1430	0,2724

Table 3: Premiums of the EUR/USD data derived from March 18th 2008 10,25 and 50 Delta risk-reversals and straddles corresponding to results in Figure 5.

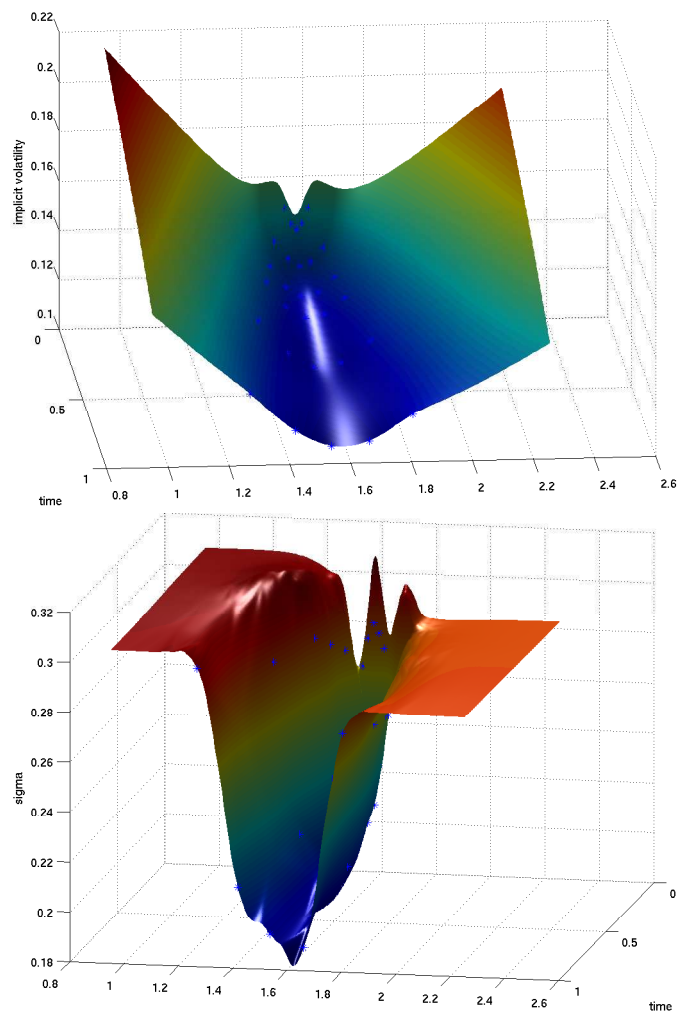


Figure 5: Top: implied volatility surface of the EUR/USD rate from Tables 1,2 and 3); marks on the surface represent the available prices (to be matched). Bottom: local volatility surface as recovered from quoted 10,25 and 50-delta risk-reversals and straddles; (mid) spot price is 1.5755; *USD* discount rate was set to $r_{USD} = 2.485\%$, and $r_{EUR} = 4.550\%$. The blue marks on the surface indicate the option prices that were used to invert i.e. the K_l and T_l ($L = 30$). After 5 iterations the prices are recovered up to $5.e - 5$ and the implied volatility up to 0.03% (below the bid/ask spread).

References

- Yves Achdou and Olivier Pironneau. *Computational methods for option pricing*, volume 30 of *Frontiers in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2005. ISBN 0-89871-573-3.
- L. Andersen and R. Brotherton-Ratcliffe. The equity option volatility smile: an implicit finite difference approach. *The Journal of Computational Finance*, 1:532, 1998.
- M. Avellaneda, C. Friedman, R. Holmes, and D. Samperi. Calibrating volatility surfaces via relative-entropy minimization. *Applied Mathematical Finance*, 4(1):37–64, 1997. URL <http://www.informaworld.com/10.1080/135048697334827>.
- H. Berestycki, J. Busca, and I. Florent. Asymptotics and calibration of local volatility models. *Quantitative Finance*, 2, 2002. URL <http://www.informaworld.com/10.1088/1469-7688/2/1/305>.
- F. Black and M. Scholes. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81, May-June 1973.
- J. Jr. Bodurtha and M. Jermakyan. Non-parametric estimation of an implied volatility surface. *J. Comput. Finance*, 2:2961, 1999.

- J. Frédéric Bonnans, J. Charles Gilbert, Claude Lemaréchal, and Claudia A. Sagastizábal. *Numerical optimization*. Universitext. Springer-Verlag, Berlin, second edition, 2006. ISBN 3-540-35445-X. Theoretical and practical aspects.
- I. Bouchouev and V. Isakov. The inverse problem of option pricing. *Inverse Problems*, 13:L117, 1997.
- I. Bouchouev and V. Isakov. Uniqueness, stability and numerical methods for the inverse problem that arises in financial markets. *Inverse Problems*, 15:R95116, 1999.
- Thomas F. Coleman and Yuying Li. A reflective Newton method for minimizing a quadratic function subject to bounds on some of the variables. *SIAM J. Optim.*, 6(4):1040–1058, 1996. ISSN 1052-6234.
- Thomas F. Coleman, Yuying Li, and Arun Verma. Reconstructing the unknown local volatility function [J. Comput. Finance **2** (1999), no. 3, 77–100]. In *Quantitative analysis in financial markets*, pages 192–215. World Sci. Publ., River Edge, NJ, 2001.
- J. Cox, S. Ross, and M. Rubinstein. Option pricing: a simplified approach. *J. Financial Economics*, 7:22963, 1979.
- E. Derman and I. Kani. Riding on a Smile. *Risk*, 7(2):329, February 1994.

- B. Dupire. Pricing with a Smile. *RISK*, 7(1):18–20, 1994.
- Jim Gatheral. *The volatility surface : a practitioner's guide*. John Wiley & Sons, 2006.
- J. Hull. *Options, Futures, and Other Derivatives*. Prentice Hall, sixth edition, 2006.
- Lishang Jiang, Qihong Chen, Lijun Wang, and Jin E. Zhang. A new well-posed algorithm to recover implied local volatility. *Quantitative Finance*, 3:451–457, 2003. URL <http://www.informaworld.com/10.1088/1469-7688/3/6/304>.
- R. Lagnado and S. Osher. Reconciling differences. *Risk*, 10(4):7983, 1997.
- R. Lagnado and S. Osher. A technique for calibrating derivative security pricing models: numerical solution of the inverse problem. *J. Comput. Finance*, 1:1325, 1998.
- Marek Musiela and Marek Rutkowski. *Martingale methods in financial modelling*, volume 36 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, second edition, 2005. ISBN 3-540-20966-2.
- M. Rubinstein. Implied binomial trees. *J. Finance*, 49:771–818, 1994.