2. The volatility cube

Andrew Lesniewski

February 4, 2008

Contents

1 Dynamics of the forward curve

The forward curve continuously evolves. Ultimately, the goals of interest rate modeling are to

(a) capture the dynamics of the curve in order to price and risk manage portfolios of fixed income securities,

(b) identify trading opportunities in the fixed income markets.

We have already taken the first step in this direction, namely learned how to construct the current snapshot of the curve. This current snapshot serves as the starting point for the stochastic process describing the curve dynamics. The next step is to construct the *volatility cube*, which is used to model the uncertainties in the future evolution of the rates. The volatility cube is built out of implied volatilities of a number of liquidly trading options.

2 Options on LIBOR based instruments

Eurodollar options are standardized contracts traded at the Chicago Mercantile Exchange. These are short dated (8 quarterly and two serial contracts) American style calls and puts on Eurodollar futures. Their maturities coincide with the maturity dates of the underlying Eurodollar contracts[1]. The exchange sets the strikes for the options spaced every 25 basis points (or 12.5 bp for the front contracts). The options are cash settled.

Caps and floors are baskets of European calls (called *caplets*) and puts (called *floorlets*) on LIBOR forward rates. They trade over the counter.

Let us consider for example, a 10 year spot starting cap struck at 5.50%. It consists of 39 caplets each of which expires on the 3 month anniversary of today's date. It pays $\max$ (current LIBOR fixing $- 5.50\%, 0) \times$ act/360 day count fraction. The payment is made at the end of the 3 month period covered by the LIBOR contract and follows the modified business day convention. Notice that the very first period is excluded from the cap: this is because the current LIBOR fixing is already known and no optionality is left in that period.

In addition to spot starting caps and floors, *forward starting* instruments trade. For example, a 1 year $\times$ 5 year (in the market lingo: "1 by 5") cap struck at 5.50% consists of 16 caplets struck at 5.50% the first of which matures one year from today. The final maturity of the contract is 5 years, meaning that the last

[1] In addition to the quarterly and serial contracts, a number of *midcurve* options trade which, for our purposes, are exotic instruments and do not enter the volatility cube construction.

caplets matures 4 years and 9 months from today (with appropriate business dates adjustments). Unlike in the case of spot starting caps, the first period is included into the structure, as the first LIBOR fixing is of course unknown. Note that the total maturity of the $m \times n$ cap is n years.

The definitions of floors are similar with the understanding that a floorlet pays $\max$ (strike − current LIBOR fixing$\%, 0) \times$ act/360 day count fraction at the end of the corresponding period.

Swaptions are European calls and puts (in the market lingo: *payers* and *receivers*, respectively) on forward swap rates. They trade over the counter.

For example, a 5.50% 1Y $\rightarrow$ 5Y ("1 into 5") receiver swaption gives the holder the right to receive 5.50% on a 5 year swap starting in 1 year. More precisely, the option holder has the right to exercise the option on the 1 year anniversary of today (with the usual business day convention adjustments) in which case they enter into a receiver swap starting two business days thereafter. Similarly, a 5.50% 5Y $\rightarrow$ 10Y ("5 into 10") payer swaption gives the holder the right to pay 5.50% on a 10 year swap starting in 5 year. Note that the total maturity of the $m \rightarrow n$ swaption is $m + n$ years.

Since a swap can be viewed as a particular basket of underlying LIBOR forwards, a swaption is an option on a basket of forwards. This observation leads to the popular relative value trade of, say, a $2 \rightarrow 3$ swaption straddle versus a 2×5 cap / floor straddle. Such a trade my reflect the trader's view on the correlations between the LIBOR forwards or a misalignment of swaption and cap / floor volatilities.

2.1 Black's model

The standard way of quoting prices on caps / floors and swaptions is in terms of *Black's model* which is a version of the Black-Scholes model adapted to deal with forward underlying assets. In order to fix the notation we briefly discuss this model now, deferring a more indebt discussion of interest rate modeling to later parts of these lectures.

We assume that a forward rate $F(t)$, such as a LIBOR forward or a forward swap rate, follows a driftless lognormal process reminiscent of the basic Black-Scholes model,

$$dF(t) = \sigma F(t)\, dW(t)\,. \tag{1}$$

Here $W(t)$ is a Wiener process, and σ is the *lognormal volatility*. It is understood here, that we have chosen a numeraire $\mathfrak{N}$ with the property that, in the units of that numeraire, $F(t)$ is a tradable asset. The process $F(t)$ is thus a martingale, and we let Q denote the probability distribution.

The solution to this stochastic differential equation reads:

$$F(t) = F_0 \exp\left(\sigma W(t) - \frac{1}{2}\,\sigma^2 t\right). \tag{2}$$

Therefore, today's value of a European call struck at K and expiring in T years is given by:

$$\begin{aligned}\mathsf{PV}_{\text{call struck at } K} &= \mathfrak{N}(0)\,\mathsf{E}^{\mathsf{Q}}\left[\max\left(F(T) - K, 0\right)\right] \\ &= \mathfrak{N}(0)\,\frac{1}{\sqrt{2\pi T}}\int_{-\infty}^{\infty}\max\left(F_0 e^{\sigma W - \frac{1}{2}\,\sigma^2 T} - K, 0\right) e^{-\frac{W^2}{2T}}\,dW,\end{aligned} \tag{3}$$

where E^{Q} denotes expected value with respect to Q. The last integral can easily be carried out, and we find that

$$\begin{aligned}\mathsf{PV}_{\text{call struck at } K} &= \mathfrak{N}(0)\left[F_0 N(d_+) - K N(d_-)\right] \\ &\equiv \mathfrak{N}(0)\,B_{\text{call}}(T, K, F_0, \sigma).\end{aligned} \tag{4}$$

Here, $N(x)$ is the cumulative normal distribution, and

$$d_\pm = \frac{\log\frac{F_0}{K} \pm \frac{1}{2}\,\sigma^2 T}{\sigma\sqrt{T}}. \tag{5}$$

The price of a European put is given by:

$$\begin{aligned}\mathsf{PV}_{\text{put struck at } K} &= \mathfrak{N}(0)\left[-F_0 N(-d_+) + K N(-d_-)\right] \\ &\equiv \mathfrak{N}(0)\,B_{\text{put}}(T, K, F_0, \sigma).\end{aligned} \tag{6}$$

2.2 Valuation of caps and floors

A cap is a basket of options on LIBOR forward rates. Recall that a forward rate $F(t, T)$ for the settlement t and maturity T can be expressed in terms of discount factors:

$$\begin{aligned}F(t, T) &= \frac{1}{\delta}\left(\frac{1}{P(t, T)} - 1\right) \\ &= \frac{1}{\delta}\,\frac{P(0, t) - P(0, T)}{P(0, T)}.\end{aligned} \tag{7}$$

The interpretation of this identity is that $F(t, T)$ is a tradable asset if we use the zero coupon bond maturing in T years as numeraire. Indeed, the trade is as follows:

(a) Buy $1/\delta$ face value of the zero coupon bond for maturity t.

(b) Sell $1/\delta$ face value of the zero coupon bond for maturity T.

A LIBOR forward rate can thus be modeled as a martingale! Choosing, for now, the process to be (1), we conclude that the price of a call on $F(t,T)$ (or caplet) is given by

$$\mathsf{PV}_{\text{caplet}} = \delta B_{\text{call}}(t, K, F_0, \sigma) P(0,T), \tag{8}$$

where F_0 denotes here today's value of $F(t,T)$.

Since a cap is a basket of caplets, its value is the sum of the values of the constituent caplets:

$$\mathsf{PV}_{\text{cap}} = \sum_{j=1}^{n} \delta_j B_{\text{call}}(T_{j-1}, K, F_j, \sigma_j) P(0, T_j), \tag{9}$$

where δ_j is the day count fraction applying to the accrual period starting at T_{j-1} and ending at T_j, and F_j is the LIBOR forward rate for that period. Notice that, in the formula above, the date T_{j-1} has to be adjusted to accurately reflect the expiration date of the option (2 business days before the start of the accrual period). Similarly, the value of a floor is

$$\mathsf{PV}_{\text{cap}} = \sum_{j=1}^{n} \delta_j B_{\text{floor}}(T_{j-1}, K, F_j, \sigma_j) P(0, T_j). \tag{10}$$

What is the at the money (ATM) cap? Characteristic of an ATM option is that the call and put struck ATM have the same value. We shall first derive a put / call parity relation for caps and floors. Let E^{Q_j} denote expected value for the probability distribution corresponding to the zero coupon bond maturing at T_j. Then,

$$\begin{aligned}
&\mathsf{PV}_{\text{floor}} - \mathsf{PV}_{\text{cap}} \\
&\quad = \sum_{j=1}^{n} \delta_j \left(\mathsf{E}^{\mathsf{Q}_j}[\max(K - F_j, 0)] - \mathsf{E}^{\mathsf{Q}_j}[\max(F_j - K, 0)] \right) P(0, T_j) \\
&\quad = \sum_{j=1}^{n} \delta_j \mathsf{E}^{\mathsf{Q}_j}[K - F_j] P(0, T_j).
\end{aligned}$$

Now, the expected value $\mathsf{E}^{\mathsf{Q}_j}[F_j]$ is the current value of the forward which, by an excusable abuse of notation, we shall also denote by F_j. Hence we have arrived at the following put / call parity relation:

$$\begin{aligned}
\mathsf{PV}_{\text{cap}} - \mathsf{PV}_{\text{floor}} &= K \sum_{j=1}^{n} \delta_j P(0, T_j) - \sum_{j=1}^{n} \delta_j F_j P(0, T_j) \\
&= \mathsf{PV}_{\text{swap paying } K \text{, q, act/360}}.
\end{aligned} \tag{11}$$

This is an important relation. It implies that:

(a) It is natural to think about a floor as a call option, and a cap as a put option. The underlying asset is the forward starting swap on which both legs pay quarterly and interest accrues on the act/360 basis. The coupon dates on the swap coincide with the payment dates on the cap / floor.

(a) The ATM rate is the break-even rate on this swap. This rate is close to but not identical to the break-even rate on the standard semi-annual swap.

2.3 Valuation of swaptions

Let $S\left(t, T_{\text{start}}, T_{\text{mat}}\right)$ denote the forward swap rate observed at time $t < T_{\text{start}}$ (in particular, $S\left(T_{\text{start}}, T_{\text{mat}}\right) = S\left(0, T_{\text{start}}, T_{\text{mat}}\right)$). We know from Lecture Notes 1 that the forward swap rate is given by

$$S\left(t, T_{\text{start}}, T_{\text{mat}}\right) = \frac{P\left(t, T_{\text{start}}\right) - P\left(t, T_{\text{mat}}\right)}{L\left(t, T_{\text{start}}, T_{\text{mat}}\right)}, \tag{12}$$

where $L\left(t, T_{\text{start}}, T_{\text{mat}}\right)$ is the forward level function:

$$L\left(t, T_{\text{start}}, T_{\text{mat}}\right) = \sum_{j=1}^{n_{\text{fixed}}} \alpha_j P\left(t, T_j\right). \tag{13}$$

The forward level function is the time t PV of an *annuity* paying $\$1$ on the dates $T_1, T_2, \ldots, T_n$. As in the case of a simple LIBOR forward, the interpretation of (12) is that $S\left(t, T_{\text{start}}, T_{\text{mat}}\right)$ is a tradable asset if we use the annuity as numeraire. Indeed, the trade is as follows:

(a) Buy $\$1$ face value of the zero coupon bond for maturity T_{start}.

(b) Sell $\$1$ face value of the zero coupon bond for maturity T_{mat}.

A forward swap rate can thus be modeled as a martingale! Choosing, again, the lognormal process (1), we conclude that the value of a receiver swaption is thus given by

$$\mathsf{PV}_{\text{rec}} = LB_{\text{put}}\left(T, K, S_0, \sigma\right), \tag{14}$$

and the value of a payer swaption is

$$\mathsf{PV}_{\text{pay}} = LB_{\text{call}}\left(T, K, S_0, \sigma\right), \tag{15}$$

where S_0 is today's value of the forward swap rate $S\left(T_{\text{start}}, T_{\text{mat}}\right)$.

The put / call parity relation for swaptions is easy to establish:

$$\mathsf{PV}_{\text{rec}} - \mathsf{PV}_{\text{pay}} = \mathsf{PV}_{\text{swap paying } K\text{, s, 30/360}}. \tag{16}$$

Therefore,

(a) It is natural to think about a receiver as a call option, and a payer as a put option.

(a) The ATM rate is the break-even rate on the underlying forward starting swap.

3 Beyond Black's model

The basic premise of Black's model, that σ is independent of K and F_0, is not supported by the interest volatility markets. In particular, for a given maturity, option implied volatilities exhibit a pronounced dependence on their strikes. This phenomenon is called the *skew* or the *volatility smile*. It became apparent especially over the past ten years or so, that in order to accurately value and risk manage options portfolios refinements to Black's model are necessary.

An improvement over Black's model is a class of models called local volatility models. The idea is that even though the exact nature of volatility (it could be stochastic) is unknown, one can, in principle, use the market prices of options in order to recover the risk neutral probability distribution of the underlying asset. This, in turn, will allow us to find an effective ("local") specification of the underlying process so that the implied volatilities match the market implied volatilities.

Local volatility models are usually specified in the form

$$dF(t) = C(F(t), t)\, dW(t), \tag{17}$$

where $C(F, t)$ is a certain effective volatility coefficient. Popular local volatility models which admit analytic solutions are:

(a) The normal model.

(b) The shifted lognormal model.

(c) The CEV model.

We now briefly discuss the basic features of these models.

3.1 Normal model

The dynamics for the forward rate $F(t)$ in the normal model reads

$$dF(t) = \sigma dW(t), \tag{18}$$

under the suitable choice of numeraire. The parameter σ is appropriately called the *normal volatility*. This is easy to solve:

$$F(t) = F_0 + \sigma W(t). \tag{19}$$

This solution exhibits one of the main drawbacks of the normal model: with non-zero probability, $F(t)$ may become negative in finite time. Under typical circumstances, this is, however, a relatively unlikely event.

Prices of European calls and puts are now given by:

$$\begin{aligned} \mathsf{PV}_{\text{call}} &= \mathfrak{N}(0)\, B^{\text{n}}_{\text{call}}(T, K, F_0, \sigma), \\ \mathsf{PV}_{\text{put}} &= \mathfrak{N}(0)\, B^{\text{n}}_{\text{put}}(T, K, F_0, \sigma). \end{aligned} \tag{20}$$

The functions $B^{\text{n}}_{\text{call}}(T, K, F_0, \sigma)$ and $B^{\text{n}}_{\text{put}}(T, K, F_0, \sigma)$ are given by:

$$\begin{aligned} B^{\text{n}}_{\text{call}}(T, K, F_0, \sigma) &= \sigma\sqrt{T}\Big(d_+ N(d_+) + N'(d_+)\Big), \\ B^{\text{n}}_{\text{put}}(T, K, F_0, \sigma) &= \sigma\sqrt{T}\Big(d_- N(d_-) + N'(d_-)\Big), \end{aligned} \tag{21}$$

where

$$d_\pm = \pm\,\frac{F_0 - K}{\sigma\sqrt{T}}\ . \tag{22}$$

The normal model is (in addition to the lognormal model) an important benchmark in terms of which implied volatilities are quoted. In fact, many traders are in the habit of thinking in terms of normal implied volatilities, as the normal model often seems to capture the rates dynamics better than the lognormal (Black's) model.

3.2 Shifted lognormal model

The dynamics of the shifted lognormal model reads:

$$dF(t) = (\sigma_1 F(t) + \sigma_0)\, dW(t).$$

Volatility structure is given by the values of the parameters σ_1 and σ_0.

Prices of calls and puts are given by the following valuation formulas:

$$\begin{aligned} \mathsf{PV}_{\text{call}} &= \mathfrak{N}(0)\, B^{\text{sln}}_{\text{call}}(T, K, F_0, \sigma_0, \sigma_1), \\ \mathsf{PV}_{\text{put}} &= \mathfrak{N}(0)\, B^{\text{sln}}_{\text{put}}(T, K, F_0, \sigma_0, \sigma_1). \end{aligned} \tag{23}$$

The functions $B^{\text{sln}}_{\text{call}}(T, K, F_0, \sigma_0, \sigma_1)$ and $B^{\text{sln}}_{\text{put}}(T, K, F_0, \sigma_0, \sigma_1)$ are generalizations of the corresponding functions for the lognormal and normal models:

$$B^{\text{sln}}_{\text{call}}(T, K, F_0, \sigma_0, \sigma_1) = \left(F_0 + \frac{\sigma_0}{\sigma_1}\right) N(d_+) - \left(K + \frac{\sigma_0}{\sigma_1}\right) N(d_-), \tag{24}$$

where

$$d_\pm = \frac{\log \dfrac{\sigma_1 F_0 + \sigma_0}{\sigma_1 K + \sigma_0} \pm \dfrac{1}{2}\, \sigma_1^2 T}{\sigma_1 \sqrt{T}}, \tag{25}$$

and

$$B^{\text{sln}}_{\text{put}}(T, K, F_0, \sigma_0, \sigma_1) = -\left(F_0 + \frac{\sigma_0}{\sigma_1}\right) N(-d_+) + \left(K + \frac{\sigma_0}{\sigma_1}\right) N(-d_-). \tag{26}$$

The shifted lognormal model is used by some market practitioners as a convenient compromise between the normal and lognormal models. It captures some aspects of the volatility smile.

3.3 The CEV model

The dynamics in the CEV model is given by

$$dF(t) = \sigma F(t)^\beta \, dW(t),$$

where $0 < \beta < 1$. In order for the dynamics to make sense, we have to prevent $F(t)$ from becoming negative (otherwise $F(t)^\beta$ would turn imaginary!). To this end, we specify a boundary condition at $F = 0$. It can be

(a) *Dirichlet* (absorbing): $F|_0 = 0$. Solution exists for all values of β, or

(b) *Neumann* (reflecting): $F'|_0 = 0$. Solution exists for $\frac{1}{2} \leq \beta < 1$.

Pricing formulas for the CEV model are of the usual (albeit a bit more complicated) form. For example, in the Dirichlet case the prices of calls and puts are:

$$\begin{aligned} \mathsf{PV}_{\text{call}} &= \mathfrak{N}(0)\, B^{\text{CEV}}_{\text{call}}(T, K, F_0, \sigma), \\ \mathsf{PV}_{\text{put}} &= \mathfrak{N}(0)\, B^{\text{CEV}}_{\text{put}}(T, K, F_0, \sigma). \end{aligned} \tag{27}$$

The functions $B^{\text{CEV}}_{\text{call}}(T, K, F_0, \sigma)$ and $B^{\text{CEV}}_{\text{put}}(T, K, F_0, \sigma)$ are expressed in terms of the cumulative function of the non-central χ^2 distribution:

$$\chi^2(x; r, \lambda) = \int_0^x p(y; r, \lambda)\, dy, \tag{28}$$

whose density is given by a Bessel function:

$$p\left(x;r,\lambda\right)=\frac{1}{2}\left(\frac{x}{\lambda}\right)^{(r-2)/4}\exp\left(-\frac{x+\lambda}{2}\right)I_{(r-2)/2}\left(\sqrt{\lambda x}\right). \tag{29}$$

We also need the quantity:

$$\nu=\frac{1}{2\left(1-\beta\right)}\ ,\quad \text{i.e. } \nu\geq\frac{1}{2}\ . \tag{30}$$

Then

$$\begin{aligned}B_{\text{call}}^{\text{CEV}}\left(T,K,F_0,\sigma\right)=F_0&\left(1-\chi^2\left(\frac{4\nu^2K^{1/\nu}}{\sigma^2T};\ 2\nu+2,\frac{4\nu^2F_0^{1/\nu}}{\sigma^2T}\right)\right)\\&-K\chi^2\left(\frac{4\nu^2F_0^{1/\nu}}{\sigma^2T};\ 2\nu,\frac{4\nu^2K^{1/\nu}}{\sigma^2T}\right)\ ,\end{aligned} \tag{31}$$

and

$$\begin{aligned}B_{\text{put}}^{\text{CEV}}\left(T,K,F_0,\sigma\right)=F_0\chi^2&\left(\frac{4\nu^2K^{1/\nu}}{\sigma^2T};\ 2\nu+2,\frac{4\nu^2F_0^{1/\nu}}{\sigma^2T}\right)\\&-K\left(1-\chi^2\left(\frac{4\nu^2F_0^{1/\nu}}{\sigma^2T};\ 2\nu,\frac{4\nu^2K^{1/\nu}}{\sigma^2T}\right)\right)\ .\end{aligned} \tag{32}$$

4 Stochastic volatility and the SABR model

The volatility skew models that we have discussed so far improve on Black's models but still fail to reflect the market dynamics. One issue is, for example, the "wing effect" exhibited by the implied volatilities of some maturities (especially shorter dated) and tenors which is not captured by these models: the implied volatilities tend to rise for high strikes forming the familiar "smile" shape. Among the attempts to move beyond the locality framework are:

(a) *Stochastic volatility models*. In this approach, we add a new stochastic factor to the dynamics by assuming that a suitable volatility parameter itself follows a stochastic process.

(b) *Jump diffusion models*. These models use a broader class of stochastic processes (for example, *Levy processes*) to drive the dynamics of the underlying asset. These more general processes allow for discontinuities ("jumps") in the asset dynamics.

For lack of time we shall discuss an example of approach (a), namely the SABR model.

4.1 Implied volatility

The SABR model is an extension of the CEV model in which the volatility parameter σ is assumed to follow a stochastic process. Its dynamics is given by:

$$\begin{aligned} dF(t) &= \sigma(t)\, C(F(t))\, dW(t)\,, \\ d\sigma(t) &= \alpha\sigma(t)\, dZ(t)\,. \end{aligned} \tag{33}$$

Here $F(t)$ is the forward rate process, and $W(t)$ and $Z(t)$ are Wiener processes with

$$\mathsf{E}\left[dW(t)\, dZ(t)\right] = \rho dt,$$

where the correlation ρ is assumed constant. The diffusion coefficient $C(F)$ is assumed to be of the CEV type:

$$C(F) = F^{\beta}. \tag{34}$$

Note that we assume that a suitable numeraire has been chosen so that $F(t)$ is a martingale. The process $\sigma(t)$ is the stochastic component of the volatility of F_t, and α is the volatility of $\sigma(t)$ (the *volvol*) which is also assumed to be constant. As usual, we supplement the dynamics with the initial condition

$$\begin{aligned} F(0) &= F_0, \\ \sigma(0) &= \sigma_0, \end{aligned} \tag{35}$$

where F_0 is the current value of the forward, and σ_0 is the current value of the volatility parameter.

Except for the special case of $\beta = 0$, no explicit solution to this model is known. The general case can be solved approximately by means of a perturbation expansion in the parameter $\varepsilon = T\alpha^2$, where T is the maturity of the option. As it happens, this parameter is typically small and the approximate solution is actually quite accurate. Also significantly, this solution is very easy to implement in computer code, and lends itself well to risk management of large portfolios of options in real time.

An analysis of the model dynamics shows that the implied normal volatility is approximately given by:

$$\begin{aligned} \sigma_{\mathrm{n}}(T, K, F_0, \sigma_0, \alpha, \beta, \rho) = \alpha\, \frac{F_0 - K}{\delta(K, F_0, \sigma_0, \alpha, \beta)} \times \\ \Bigg\{ 1 + \left[\frac{2\gamma_2 - \gamma_1^2}{24} \left(\frac{\sigma_0 C(F_{\mathrm{mid}})}{\alpha} \right)^2 + \frac{\rho\gamma_1}{4}\, \frac{\sigma_0 C(F_{\mathrm{mid}})}{\alpha} + \frac{2 - 3\rho^2}{24} \right] \varepsilon \\ + \dots \Bigg\}, \end{aligned} \tag{36}$$

where F_{mid} denotes a conveniently chosen midpoint between F_0 and K (such as $\sqrt{F_0 K}$ or $(F_0 + K)/2$), and

$$\gamma_1 = \frac{C'(F_{\text{mid}})}{C(F_{\text{mid}})} ,$$
$$\gamma_2 = \frac{C''(F_{\text{mid}})}{C(F_{\text{mid}})} .$$

The "distance function" entering the formula above is given by:

$$\delta(K, F_0, \sigma_0, \alpha, \beta) = \log\left(\frac{\sqrt{1 - 2\rho\zeta + \zeta^2} + \zeta - \rho}{1 - \rho}\right),$$

where

$$\begin{aligned}\zeta &= \frac{\alpha}{\sigma_0} \int_K^{F_0} \frac{dx}{C(x)} \\ &= \frac{\alpha}{\sigma_0 (1 - \beta)} \left(F_0^{1-\beta} - K^{1-\beta}\right).\end{aligned} \tag{37}$$

A similar asymptotic formula exists for the implied lognormal volatility σ_{ln}.

4.2 Calibration of SABR

For each option maturity and underlying we have to specify 4 model parameters: $\sigma_0, \alpha, \beta, \rho$. In order to do it we need, of course, market implied volatilities for several different strikes. Given this, the calibration poses no problem: one can use, for example, Excel's Solver utility.

It turns out that there is a bit of redundancy between the parameters β and ρ. As a result, one usually calibrates the model by fixing one of these parameters:

(a) Fix β, say $\beta = 0.5$, and calibrate σ_0, α, ρ.

(b) Fix $\rho = 0$, and calibrate σ_0, α, β.

Calibration results show interesting term structure of the model parameters as functions of the maturity and underlying. Typical is the shape of the parameter α which start out high for short dated options and then declines monotonically as the option maturity increases. This indicates presumably that modeling short dated options should include a jump diffusion component.

5 Building the vol cube

Market implied volatilities are usually organized by:

(a) Option maturity.

(b) Tenor of the underlying instrument.

(c) Strike on the option.

This three dimensional object is called the *volatility cube*. The markets provide information for certain benchmark maturities (1 month, 3 months, 6 months, 1 year, ...), underlyings (1 year, 2 years, ...), and strikes (ATM, ± 50 bp, ...) only, and the process of volatility cube construction requires performing intelligent interpolations.

5.1 ATM swaption volatilities

The market quotes swaption volatilities for certain standard maturities and underlyings. Matrix of at the money volatilities may look like this:

mat\tenor	0.25	1	2	3	4	5	7	10	15	20
0.25	6.7	13.3	15.5	15.7	15.6	15.5	15.0	14.2	13.5	13.1
0.5	11.9	14.8	16.2	16.2	16.1	15.9	15.3	14.5	13.8	13.3
1	16.7	17.1	17.2	17.0	16.8	16.6	16.0	15.2	14.4	13.9
2	18.5	18.2	17.90	17.7	17.4	17.2	16.7	15.9	15.0	14.5
3	18.9	18.4	18.2	18.0	17.7	17.5	17.0	16.3	15.3	14.8
4	18.9	18.3	18.1	17.9	17.6	17.5	16.9	16.2	15.2	14.7
5	18.8	18.1	17.9	17.6	17.4	17.3	16.7	16.0	15.0	14.5
7	18.0	17.4	17.1	16.8	16.6	16.4	15.9	15.3	14.2	13.8
10	16.2	16.1	15.8	15.6	15.4	15.2	14.8	14.2	13.0	12.6

5.2 Stripping cap volatility

A cap is a basket of options of different maturities and different moneynesses. For simplicity, the market quotes cap / floor prices in terms of a single number, the *flat volatility*. This is the single volatility which, when substituted into the valuation formula (for all caplets / floorlets!), reproduces the correct price of the instrument. Clearly, flat volatility is a dubious concept: since a single caplet may be part of different caps it gets assigned different flat volatilities. The process of constructing actual implied caplet volatility from market quotes is called *stripping* cap volatility. The result of stripping is a sequence of ATM caplet volatilities for maturities all maturities ranging from one day to, say, 30 years. Convenient benchmarks are 3 months, 6 months, 9 months, ... The market data usually include Eurodollar options and OTC caps and floors.

There are various methods of stripping cap volatility. Among them we list:

- *Bootstrap*. One starts at the short end and moves further trying to match the prices of Eurodollar options and spot starting caps / floors.
- *Optimization*. Use a two step approach: in the first step fit the caplet volatilities to the *hump function*:

$$H(t) = (\alpha + \beta t)\, e^{-\lambda t} + \mu. \tag{38}$$

  Generally, the hump function gives a qualitatively correct shape of the cap volatility. Quantitatively, the fit is insufficient for accurate pricing and we should refine it. An good approach is to use smoothing B-splines. Once α, β, λ, and μ have been calibrated, we use cubic B-splines in a way similar to the method explained in Lecture 1 in order to nail down the details of the caplet volatility curve.

5.3 Adding the third dimension

It is convenient to specify the strike dependence of volatility in terms of the set of parameters of a smile model (such as a local volatility model or a stochastic volatility model). This way, (a) we can calculate on the fly the implied volatility for any strike, (b) the dependence of the volatility on the strike is smooth.

6 Sensitivities and hedging of options

6.1 The greeks

Traditional risk measures of options are the *greeks*: delta, gamma, vega, theta, etc.[2], see [2]. Recall, for example, that the delta of an option is the derivative of the premium with respect the underlying. This poses a bit of a problem in the world of interest rate derivatives, as the interest rates play a dual role in the option valuation formulas: (a) as the underlyings, and (b) as the discounting rates. One has thus to differentiate both the underlying *and* the discount factor when calculating the delta of a swaption!

In risk managing a portfolio of interest rate options, we use the concepts (explained in Lecture 1) of partial sensitivities to particular curve segments. They

[2] Rho, vanna, volga,... .

can be calculated either by perturbing selected inputs to the curve construction or by perturbing a segment of the forward curve, and calculating the impact of this perturbation on the value of the portfolio.

Vega risk is the sensitivity of the portfolio to volatility and is traditionally measured as the derivative of the option price with respect to the implied volatility. Choice of volatility model impacts not only the prices of (out of the money) options but also, at least equally significantly, their risk sensitivities. One has to think about the following issues:

(a) What is vega: sensitivity to lognormal volatility, normal volatility, another volatility parameter?

(b) What is delta: which volatility parameter should be kept constant?

6.2 Risk measures under SABR

Let us have a closer look at these issues in case of the SABR model. The delta risk is calculated by shifting the current value of the underlying while keeping the current value of implied volatility σ fixed:

$$\begin{aligned} F_0 &\to F_0 + \Delta F_0, \\ \sigma &\to \sigma, \end{aligned} \tag{39}$$

where ΔF_0 is, say, -1 bp. This scenario leads to the option delta:

$$\Delta = \frac{\partial V}{\partial F_0} + \frac{\partial V}{\partial \sigma} \frac{\partial \sigma}{\partial F_0} \,. \tag{40}$$

The first term on the right hand side in the formula above is the original Black delta, and the second arises from the systematic change in the implied volatility as the underlying changes. This formula shows that, in stochastic volatility models, there is an interaction between classic Black-Scholes style greeks! In the case at hand, the classic delta and vega contribute both to the smile adjusted delta.

Similarly, the vega risk is calculated from

$$\begin{aligned} F_0 &\to F_0, \\ \sigma_0 &\to \sigma_0 + \Delta\sigma, \end{aligned} \tag{41}$$

to be

$$\Lambda = \frac{\partial V}{\partial \sigma} \frac{\partial \sigma}{\partial \sigma_0} \,. \tag{42}$$

These formulas are the classic SABR greeks.

Modified SABR greeks below attempt to make a better use of the model dynamics. Since σ and F are correlated, whenever F changes, on average σ changes as well. A realistic scenario is thus

$$\begin{aligned} F_0 &\to F_0 + \Delta F_0, \\ \sigma_0 &\to \sigma_0 + \delta_f \sigma_0. \end{aligned} \tag{43}$$

Here $\delta_F \sigma_0$ is the average change in σ_0 caused by the change in the underlying forward. The new delta risk is given by

$$\Delta = \frac{\partial V}{\partial F_0} + \frac{\partial V}{\partial \sigma} \left(\frac{\partial \sigma}{\partial F_0} + \frac{\partial \sigma}{\partial \sigma_0} \frac{\rho \alpha}{F_0^\beta} \right). \tag{44}$$

This risk incorporates the average change in volatility caused by changes in the underlying.

Similarly, the vega risk should be calculated from the scenario:

$$\begin{aligned} F_0 &\to F_0 + \delta_\alpha F_0, \\ \sigma_0 &\to \sigma_0 + \Delta \sigma_0, \end{aligned} \tag{45}$$

where $\delta_\sigma F_0$ is the average change in F_0 caused by the change in SABR vol. This leads to the modified vega risk

$$\Lambda = \frac{\partial V}{\partial \sigma} \left(\frac{\partial \sigma}{\partial \sigma_0} + \frac{\partial \sigma}{\partial F_0} \frac{\rho F_0^\beta}{\alpha} \right). \tag{46}$$

References

[1] Brigo, D., and Mercurio, F.: *Interest Rate Models - Theory and Practice*, Springer Verlag (2006).

[2] Hull, J.: Hull, J.: *Options, Futures and Other Derivatives* Prentice Hall (2005).