Modern Modeling and Pricing of Interest Rates Derivatives
Day 1 - Session 4: Volatility Modelling

Vladimir Sankovich [1]
Qinghua Zhu [2]

London, June 6-7, 2018

[1] vsankovich@drwholdings.com
[2] qzhu@drwholdings.com

Caps and floors - payout

Cap/floorlets are the simplest instruments to access the market of IR volatility.

- Given a libor flow where the libor rate $L(t_s, T_1, T_2)$ setting at t_s for the period $[T_1, T_2]$ is paid at T_2, the related caplet expires at t_s and pays

  $$[L(t_s, T_1, T_2) - K]^+$$

  at T_2, where K is the *strike*.
- When one is short a libor flow, buying a caplet would reduce the exposure by limiting the maximum effective rate being paid to the strike.
- For a floorlet, the payout is $[K - L(t_s, T_1, T_2)]^+$, and it can be bought to "floor" one's downside exposure to a libor flow being received.

Standard quotes in the market

- The market typically quotes prices for caps/floors, which are strips of cap/floorlets on a set of spanning libors. Single caplets ("single-looks") are also quoted.
- Nowadays, the assumption for the quoted prices is for the instruments to be fully collateralised in the domestic currency, i.e. funded at the domestic OIS.

Caps and floors - quotes

Even though the most liquid quotes are for instruments "ATM", prices for a wide range of strikes are available through brokers screens.

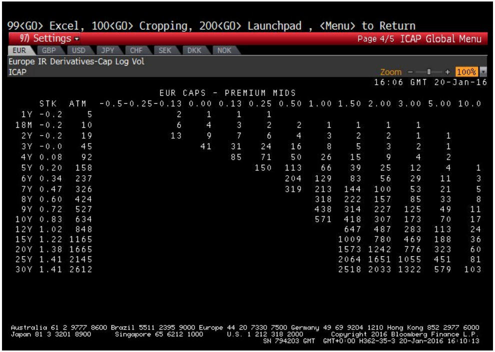
99<GO> Excel, 100<GO> Cropping, 200<GO> Launchpad , <Menu> to Return

97) Settings | Page 4/5 ICAP Global Menu

EUR | GBP | USD | JPY | CHF | SEK | DKK | NOK

Europe IR Derivatives-Cap Log Vol

ICAP | Zoom 100%

16:06 GMT 20-Jan-16

EUR CAPS - PREMIUM MIDS

	STK	ATM	-0.5	-0.25	-0.13	0.00	0.13	0.25	0.50	1.00	1.50	2.00	3.00	5.00	10.0
1Y	-0.2	5			2	1	1	1							
18M	-0.2	10			6	4	3	2	2	1	1	1	1		
2Y	-0.2	19			13	9	7	6	4	3	2	2	1	1	
3Y	-0.0	45				41	31	24	16	8	5	3	2	1	
4Y	0.08	92					85	71	50	26	15	9	4	2	
5Y	0.20	158						150	113	66	39	25	12	4	1
6Y	0.34	237							204	129	83	56	29	11	3
7Y	0.47	326							319	213	144	100	53	21	5
8Y	0.60	424								318	222	157	85	33	8
9Y	0.72	527								438	314	227	125	49	11
10Y	0.83	634								571	418	307	173	70	17
12Y	1.02	848									647	487	283	113	24
15Y	1.22	1165									1009	780	469	188	36
20Y	1.38	1665									1573	1242	776	323	60
25Y	1.41	2145									2064	1651	1055	451	81
30Y	1.41	2612									2518	2033	1322	579	103

Australia 61 2 9777 8600 Brazil 5511 2395 9000 Europe 44 20 7330 7500 Germany 49 69 9204 1210 Hong Kong 852 2977 6000
Japan 81 3 3201 8900 Singapore 65 6212 1000 U.S. 1 212 318 2000 Copyright 2016 Bloomberg Finance L.P.
SN 794203 GMT GMT+0:00 H362-35-3 20-Jan-2016 16:10:13

In the BBG grab above, you only see half of the available strikes, as the other strikes are quoted for floors. The prices of a floor and cap at the same strike are related by a so called "put-call" parity relationship, which can be derived by observing that

$$(L-K)^+ - (K-L)^+ = L - K$$

so from the floor price one can derive the cap price, and vice-versa, and thus people only quote near ATM and OTM instruments.

Pricing of caplets in the old days

When considering the pricing of a caplet, a few years ago the argument would have been the following. The Libor rate setting at t_s for the period $[T_1, T_2]$ is

$$L(t_s, T_1, T_2) = \frac{1}{\tau}\left(\frac{P(t_s, T_1)}{P(t_s, T_2)} - 1\right) . \tag{1}$$

We work in the T_2-forward measure, where the value $P(t_s, T_2)$ of the zc-bond paying at T_2 is the numeraire. In this measure, the ratio of the values of the two bonds $\frac{P(t,T_1)}{P(t,T_2)}$, is a martingale, i.e.

$$\frac{P(t, T_1)}{P(t, T_2)} = \mathbb{E}_t^{T_2}\left[\frac{P(s, T_1)}{P(s, T_2)}\right] \tag{2}$$

and thus, $L(t, T_1, T_2)$ will also be a martingale

$$L(t, T_1, T_2) = \mathbb{E}_t^{T_2}\left[L(t_s, T_1, T_2)\right] . \tag{3}$$

The above implies that the expected value for the Libor rate setting at t_s is exactly the Libor forward calculated off the YC at t, and that a generic drift-less dynamics

$$\mathrm{d}L(t, T_1, T_2) = \sigma(t, L)\cdot \mathrm{d}W_t^{T_2} \tag{4}$$

can be assigned to the libor rate. Ultimately, pricing formulae can be derived for

$$Cplt(t, K; T_1, T_2) = P(t, T_2)\mathbb{E}_t^{T_2}\left[(L(t_s, T_1, T_2) - K)^+\right] . \tag{5}$$

Pricing of caplets with full collateralisation at OIS

Collateralisation and the widening of the spreads between libor and OIS have required a more revision of the argument just presented.

Nowadays, OIS swaps give us information to bootstrap a curve for OIS discount factors $D(t, T)$, and then OIS-collateralised IRS swaps can be used to build a curve for objects of the type $F(t, T; \tau) = \mathbb{E}_t^{T_c}[L(t_s, T-\tau, T)]$, where T_c indicates the T-forward measure associated to the OIS-collateralised zc-bond $D(t, T)$. This $F(t, T; \tau)$ is simply the FRA for a fully OIS-collateralised contract paying the libor rate.

Pricing a caplet resorts to

$$Cplt(t, K; T-\tau, T) = D(t, T)\mathbb{E}_t^{T_c}\left[(L(t_s, T-\tau, T) - K)^+\right], \tag{6}$$

but now libor is not martingale in the pricing measure, where its dynamics would be more complicated than one would like.

Luckily, at the resetting time t_s the FRA matches the related libor rate, so

$$Cplt(t, K; T-\tau, T) = D(t, T)\mathbb{E}_t^{T_c}\left[(F(t_s, T-\tau, T) - K)^+\right]. \tag{7}$$

Now, by tower law, F is a martingale in the T_c-forward measure, so we'll just forget about libor, and model the associated FRA directly instead!

One way or another, we get once again to model something with a drift-less dynamics

$$dF(t, T; \tau) = \sigma(t, F) \cdot dW_t^{T_c}. \tag{8}$$

Swaptions - payout

The other basic building block for the IR volatility market is represented by the vanilla European swaptions, which grant the right to enter a pre-specified fixed-floating IR swap at an expiry time t_s .

Payer/receiver swaptions give, respectively, the right to enter a swap that pays/receives a specified fixed rate K.

Thus, for example, the payoff of a payer swaption at the expiry time t_s is

$$[S(t_s) - K]^+ A(t_s) \tag{9}$$

where, assuming the underlying swap to be fully collateralised, one has

$$S(t) = \frac{\sum_{j=1}^{n^{flt}} \tau_j^{flt} D(t, T_j^{flt}) F(t, T_j^{flt})}{A(t)} , \tag{10}$$

$$A(t) = \sum_{i=1}^{n^{fix}} \tau_i^{fix} D(t, T_i^{fix}) . \tag{11}$$

Swaptions - physical/cash settle, forward premium

Similarly to caplets,

- swaptions can be used to limit one's exposure to a swap (in the same way as caplets reduce the exposure to a libor).
- the assumption for the quoted prices is for both the underlying swap and the swaption itself to be fully collateralised in the domestic currency, i.e. funded at the domestic OIS.

To ease the settlement between counterparties, some swaptions are "cash-settled" (more specifically "ParYieldUnadjusted" in the ISDA 2006) – for these particular products (which are the standard for markets like EUR), the annuity at payout time is determined from the swap rate fixing as

$$A(S(t_s)) = \sum_{i=1}^{n^{flt}} \frac{1/f}{(1 + S(t_s)/f)^i}\,, \tag{12}$$

where f is the fixed leg frequency. We will see how these products can be priced similarly to their "Physical" counterparts, with some approximations and some caveats.

Swaptions - quotes

To ease quoting by reducing the dependency on the curve construction used by the two trading parties from value date to settlement date, instruments are often quoted "forward premium", i.e. for a premium exchanged at the settlement date, together with the swap or its cash-settled amount (well, if the option is exercised).

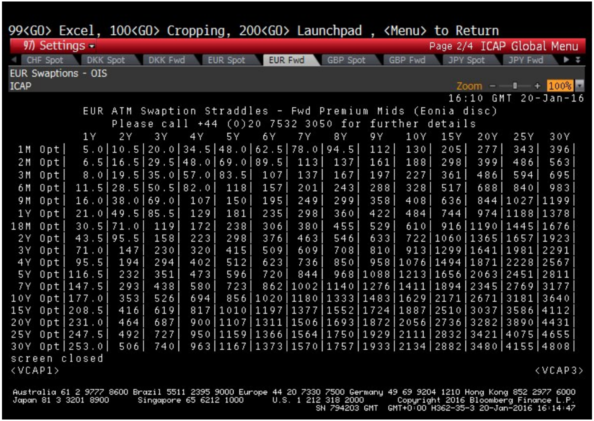

99<GO> Excel, 100<GO> Cropping, 200<GO> Launchpad , <Menu> to Return

97) Settings Page 2/4 ICAP Global Menu

CHF Spot | DKK Spot | DKK Fwd | EUR Spot | EUR Fwd | GBP Spot | GBP Fwd | JPY Spot | JPY Fwd

EUR Swaptions - OIS

ICAP

Zoom 100%

16:10 GMT 20-Jan-16

EUR ATM Swaption Straddles - Fwd Premium Mids (Eonia disc)

Please call +44 (0)20 7532 3050 for further details

	1Y	2Y	3Y	4Y	5Y	6Y	7Y	8Y	9Y	10Y	15Y	20Y	25Y	30Y
1M Opt	5.0	10.5	20.0	34.5	48.0	62.5	78.0	94.5	112	130	205	277	343	396
2M Opt	6.5	16.5	29.5	48.0	69.0	89.5	113	137	161	188	298	399	486	563
3M Opt	8.0	19.5	35.0	57.0	83.5	107	137	167	197	227	361	486	594	695
6M Opt	11.5	28.5	50.5	82.0	118	157	201	243	288	328	517	688	840	983
9M Opt	16.0	38.0	69.0	107	150	195	249	299	358	408	636	844	1027	1199
1Y Opt	21.0	49.5	85.5	129	181	235	298	360	422	484	744	974	1188	1378
18M Opt	30.5	71.0	119	172	238	306	380	455	529	610	916	1190	1445	1676
2Y Opt	43.5	95.5	158	223	298	376	463	546	633	722	1060	1365	1657	1923
3Y Opt	71.0	147	230	320	415	509	609	708	810	913	1299	1641	1981	2291
4Y Opt	95.5	194	294	402	512	623	736	850	958	1076	1494	1871	2228	2567
5Y Opt	116.5	232	351	473	596	720	844	968	1088	1213	1656	2063	2451	2811
7Y Opt	147.5	293	438	580	723	862	1002	1140	1276	1411	1894	2345	2769	3177
10Y Opt	177.0	353	526	694	856	1020	1180	1333	1483	1629	2171	2671	3181	3640
15Y Opt	208.5	416	619	817	1010	1197	1377	1552	1724	1887	2510	3037	3586	4112
20Y Opt	231.0	464	687	900	1107	1311	1506	1693	1872	2056	2736	3282	3890	4431
25Y Opt	247.5	492	727	950	1159	1366	1564	1750	1929	2111	2832	3421	4075	4655
30Y Opt	253.0	506	740	963	1167	1373	1570	1757	1933	2134	2882	3480	4155	4808

screen closed

<VCAP1> <VCAP3>

Australia 61 2 9777 8600 Brazil 5511 2395 9000 Europe 44 20 7330 7500 Germany 49 69 9204 1210 Hong Kong 852 2977 6000
Japan 81 3 3201 8900 Singapore 65 6212 1000 U.S. 1 212 318 2000 Copyright 2016 Bloomberg Finance L.P.
SN 794203 GMT GMT+0:00 H362-35-3 20-Jan-2016 16:14:47

More about availability quotes in general

- The need for banks to provide an "Independent Price Verification" (IPV) to the regulator has paved the way to the birth of services like Markit's Totem: quotes for a vast set of instruments are collected from all the major players at the end of each month, and the consensus price is returned along with some statistics on the distribution of the contributed prices.
- The results are often used also to "calibrate" the models to the market.
- The downside of these services is that the players not to be rejected tend to stick around the average, thereby gradually creating the impression of a consensus, also for more complex instruments that by their own nature, aren't that liquid. This can ultimately become a source of systemic risk.

Pricing of swaptions

To price a swaption, one can elect to work in the swap measure where the numeraire is the annuity $A(t)$. Doing so, the price of a payer swaption becomes

$$Payer(t, K) = A(t)\,\mathbb{E}_t^A\left[\frac{(S(t_s) - K)^+ A(t_s)}{A(t_s)}\right] = A(t)\,\mathbb{E}_t^A[(S(t_s) - K)^+]\,. \tag{13}$$

The value of the swap rate $S(t)$ is the ratio between the value of the floating leg and the value of the numeraire $A(t)$, thus it's a martingale in the chosen measure. Once again, we managed to find something amenable of being modeled with a drift-less dynamics

$$\mathrm{d}S(t) = \sigma(t, S) \cdot \mathrm{d}W_t^{A_c}\,. \tag{14}$$

Normal / Lognormal

Amongst the simplest models dynamics one could come up with we have

- Bachelier's model, with constant volatility σ

$$\mathrm{d}S(t) = \sigma \cdot \mathrm{d}W_t \implies S(t_s) \sim \mathcal{N}\left(S(t), \sigma\sqrt{\Delta t}\right) , \tag{15}$$

$$\begin{aligned}\mathbb{E}_t[(S(t_s) - K)^+] &= \int (s-K)^+ \varphi\left(\frac{s - S(t)}{\sigma\sqrt{\Delta t}}\right) \frac{\mathrm{d}s}{\sigma\sqrt{\Delta t}} = \\ &= (S(t) - K)\Phi\left(\frac{s - S(t)}{\sigma\sqrt{\Delta t}}\right) + \sigma\sqrt{\Delta t}\varphi\left(\frac{s - S(t)}{\sigma\sqrt{\Delta t}}\right) .\end{aligned} \tag{16}$$

- Lognormal model (e.g. BSM), with constant volatility σ

$$\mathrm{d}S(t) = \sigma\, S(t) \cdot \mathrm{d}W_t \overset{\text{Itô}}{\iff} \mathrm{d}\ln S(t) = \sigma \cdot \mathrm{d}W_t - \frac{1}{2}\sigma^2\,\mathrm{d}t , \tag{17}$$

$$\ln S(t_s) \sim \mathcal{N}\left(\ln S(t) - \frac{1}{2}\sigma^2\Delta t, \sigma\sqrt{\Delta t}\right) , \tag{18}$$

$$\mathbb{E}_t[(S(t_s) - K)^+] = \int (e^x - K)^+ \varphi\left(\frac{x - \ln S(t) + \frac{1}{2}\sigma^2\,\Delta t}{\sigma\sqrt{\Delta t}}\right) \frac{\mathrm{d}x}{\sigma\sqrt{\Delta t}} = \tag{19}$$

$$= \Phi(d_+)S(t) - \Phi(d_-)K , \tag{20}$$

with $d_\pm = \frac{1}{\sigma\sqrt{\Delta t}}[\ln(\frac{S(t)}{K} \pm \frac{\sigma^2\,\Delta t}{2}]$.

Volatility Smile

- Two models above are too stylised, i.e. there's no value of σ for which they can reproduce the prices of options with different value of the strike K.
- Still, they're used to quote prices by reporting the (log)normal volatility $\sigma(K)$ that would match the forward price of the option struck at K.
- For a given underlying, $\sigma(K)$ is the "volatility smile", either normal or lognormal, depending on which of the two models is used as reference.

Hedging

- A model should reproduce the prices seen in the market, but in that respect a volatility smile might do just as well.
- The most important role of a model, though, is that of giving indications about how to hedge one's exposure to the value of an option by trading an appropriate amount of the underlying. For example, under the Bachelier and the BSM models one would buy an amount corresponding to the first derivative wrt to today's price of the underlying – *delta*. This way, *within the model*, the 1st order variation of the value of the portfolio built would be minimal (gamma remains).
- The problem with marking a different volatility $\sigma(K)$ for each strike K, and using the hedging resulting from this assumption, is that the deltas for options at different strikes effectively come from different and inconsistent models, so they can't really be aggregated.
- Analogously, we know that the value of a portfolio of options on an underlying changes with the smile, and we also know that the smile changes day on day, but our model is not predicting that, let alone it is saying how the smile changes, in particular in relation to the value of the underlying.
- This clearly can't lead to good hedges of day-to-day portfolio price variations.

Local volatility

A first answer to the quest of a model capable of matching in a consistent way all the European options' prices on an underlying came in the form of the *local volatility* models [Dup94, DK94]. The dynamics is

$$dS(t) = \sigma(t, S(t)) \cdot dW_t \,, \tag{21}$$

where $\sigma(t, S(t))$ is a deterministic function of the underlying and time.

Intuitively, if we neglect of a moment the variability in t, the above represents a diffusion in a medium where the diffusion coefficient varies with the position. By modulating appropriately the diffusion coefficient, we can hope to be able to generate whatever diffusion profile.

Back to finance, we hope to be able to find a local volatility that can reproduce the prices for the Europeans we see in the market.

Local volatility - Dupire's formula

This is indeed possible:

$$\partial_{S_t}(S_t - K)^+ = \mathbb{1}_{S_t > K}\,, \quad \partial^2_{S_t}(S_t - K)^+ = \delta(S_t - K)\,, \tag{22}$$

$$\mathrm{d}(S_t - K)^+ = \mathbb{1}_{S_t > K}\mathrm{d}S_t + \frac{1}{2}\delta(S_t - K)\sigma_t^2\mathrm{d}t\,, \tag{23}$$

$$\partial^2_K Call(t,K) = \partial^2_K \mathbb{E}[(S_t - K)^+] = \mathbb{E}[\delta(S_t - K)] = f_{S_t}(K) \tag{24}$$

Taking expectation of the second equation,

$$\mathrm{d}\mathbb{E}[(S_t - K)^+] = \partial_t Call(t,K)\,\mathrm{d}t = \mathbb{E}[\mathbb{1}_{S_t > K}\mathrm{d}S_t] + \mathbb{E}[\frac{1}{2}\delta(S_t - K)\sigma_t^2\mathrm{d}t] = \tag{25}$$

$$= 0 + \frac{1}{2}f_{S_t}(K)\mathbb{E}[\sigma_t^2|S_t = K]\,\mathrm{d}t = \frac{1}{2}\partial^2_K Call(t,K)\sigma^2(t,K)\,\mathrm{d}t \tag{26}$$

so that the local vol matching the European prices can be calculated as [Dup94]

$$\mathbb{E}[\sigma_t^2|S_t = K] = \sigma^2(t,K) = \frac{\partial_t Call(t,K)}{\frac{1}{2}\partial^2_K Call(t,K)}\,. \tag{27}$$

Notice how the equation above also describes the relationship between a stoch vol and the prices it generates, as well what local vol to use to match up the European prices of a stoch vol model.

Smile dynamics under Local Volatility

To make use of Eq. (27) one needs to build a continuum of prices by interpolating the available ones – this requires a lot of care, as the result is very sensitive to the interpolation procedure and the way prices are "cleaned-up".

Yet, the biggest problem is once again hedging: even though the model now replicates all prices observed on a specific day in a coherent way, unfortunately the model's smile dynamics, i.e. how the smile moves day to day, is not consistent with what happens in the market.

- (for short maturities) the following relation between implied normal vol $\sigma_n(K,S)$ and the local vol $\sigma(s)$ holds

$$\frac{S-K}{\sigma_n(K,S)} = \int_K^S \frac{1}{\sigma(u)}\mathrm{d}u \quad \Longrightarrow \quad \sigma_n(K,S) \approx \sigma(\frac{K+S}{2}) \tag{28}$$

  Intuitively, we're matching the distance between K and S measured in local vol units, with its measure in the units of a fixed volatility, we're matching the quantiles of the distributions. The last inference is done by approximating the integral with the mid-point of the integrand, in the limit of K close to S.
- When the forward S moves to S', if the local vol function stays the same, one ends up with a smile $\sigma_n(K,S')$ defined as

$$\sigma_n(K,S') \approx \sigma(\frac{K+S'+(S-S)}{2}) = \sigma(\frac{(K+S'-S)+S}{2}) \approx \sigma_n(K+(S'-S),S)\,, \tag{29}$$

  i.e. the smile translates in the opposite direction of the forward – which is not what typically happens in the market.

Stochastic Volatility

- As we've just seen, local volatility models can't reproduce the so called *smile backbone*, i.e. how the smile reacts to a move in the value of the underlying. Because of this, even though the model matches the market on day one, it surely won't match it the following day – the model will need recalibrating and the hedges will need rebalancing, which is clearly not good.
- This is one of the reasons that lead to the development of stochastic volatility models, where the volatility σ_t is itself a stochastic process.
- The SABR model [HKLW02] has been the mainstream stochastic-local volatility model for IR vanillas for many years now

$$\begin{cases} \mathrm{d}S_t = \alpha_t\, S_t^{\beta}\, \mathrm{d}W_t \\ \mathrm{d}\alpha_t = \alpha_t\, \nu\, \mathrm{d}Z_t \\ \alpha_0 = \alpha \end{cases} \tag{30}$$

where W_t and Z_t are Brownian processes in the chosen measure, $\langle \mathrm{d}W_t, \mathrm{d}Z_t \rangle = \rho \mathrm{d}t$, and α, β, ρ, ν are the parameters of the model.

SABR

Why SABR?

- The model is flexible enough to generate skew & smile.
- The dynamics of the smile (i.e. how it moves when the forward moves) is also rather flexible, and it can be calibrated to match the one observed in the market (vs a pure local volatility model). In particular, the backbone behaves like the market see for example [RPW09].
- It was the simplest stoch-vol model being homogeneous in both forward and volatility (i.e. without typical scales).
- Above all, a simple approximation for the corresponding volatility smile $\sigma(K; S_0, \alpha, \beta, \rho\nu, \tau)$ was presented together with the model in [HKLW02]. That approximation has become market standard, notwithstanding issues issues like the arbitrages it produces (above all for long maturities, and low forwards), or like not admitting negative rates. . . all things that people have dealt with one way or another.

Effect of the SABR parameters

The parameters have an intuitive effect on the smile

- α controls the overall variance of the distribution, i.e. the level of the smile,
- ρ largely controls the skew of the distribution,
- β also controls the skew by providing a different level of volatility depending on where the forward has diffused to,
- the vol of vol ν controls more the kurtosis, i.e. it controls the smile.

The effect of the parameters becomes clear when looking at the short time the normal smile Taylor expansion (see, for example, [AH13])

$$\sigma + \frac{1}{2}\left[\rho\,\nu + \sigma\,\frac{\beta}{S}\right](K-S) + \frac{1}{12\sigma}\left[(2-3\rho^2)\nu^2 + \sigma^2\frac{\beta(\beta-2)}{S^2}\right](K-S)^2 \qquad (31)$$

where $\sigma = \alpha\, S^{\beta}$ is the ATM normal vol. From the above one can easily identify the role of the parameters in determining the smile.

Effect of SABR parameters

Here's the effect of moving the sabr params on a 2yx5y USD swaption

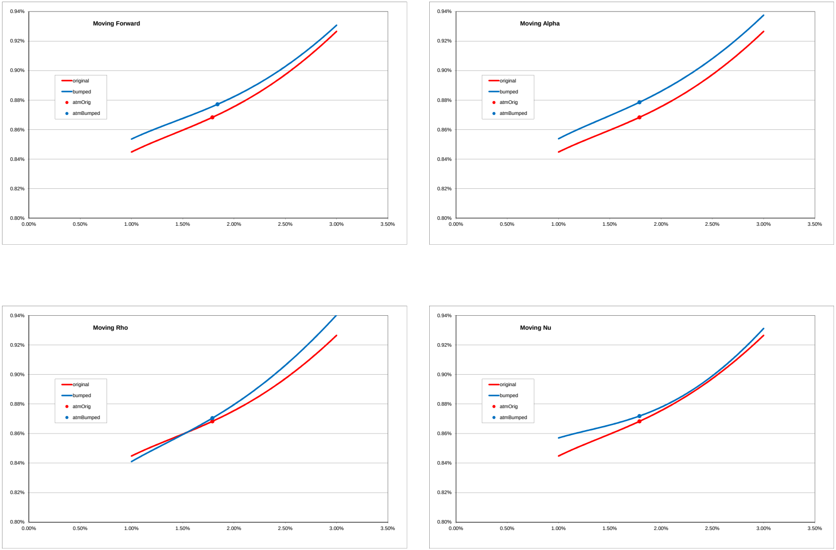

Using SABR

How is the model used?

- The β parameter is chosen according to "aesthetics", or trying to match the backbone, or trying to control the CMS prices obtained by replication on the resulting smile (more about this in the coming sessions).
- The parameter α is used to match the ATM vol, which varies day on day, whereas ρ and ν are used to control the skew and the smile, and are rather stable.
- Some people prefer to parametrise the model in terms of the resulting ATM vol (normal vol for rates these days!).
- A grid of parameters is maintained per every expiry x tenor point for which liquid quotes are available in the swaptions market.
- A similar grid is maintained for caps.
- Notice that convention for determining expiries for the standard swaptions and caplets don't match in general. Also one has to choose how to deal with the gridding along tenors, and minor adjustments of the grid points may be needed.

The meaning of the grid of parameters - interpolation

Armed with a grid of parameters calibrated to particular swaption smiles, we have to determine a way for interpolating across the points.

If interpolating across tenors doesn't pose major concerns, when it comes to interpolating across expiries, one has to be a bit more careful.

- We're implicitly assuming the market to be kind of time-homogeneous, so we expect a "real" swap to roll down the expiries, and be subject to the parameters applying to the changing expiry.
- Thus, we want to guarantee that the variance of that swap increases with time, and to do so we would like to interpolate linearly in variance between two expiries. As a proxy for the variance, one can use the $\tau \cdot \sigma_{n,ATM}^2$, or it's proxy $\tau \cdot (\alpha S^\beta)^2$, and deduce the values for the interpolated SABR parameters.
- The above procedure guarantees we're interpolating homogeneous quantities – imagine for a second β varies from 0 at 1Y expiry to 1 at the 2Y expiry (clearly nonsense, but it proves the point), if we interpolated directly on α, we'd been mixing up a lognormal volatility with a normal one.

Event day model

- Not all days are made the same for the market, some recurring events (like central banks meetings etc) are known to contribute more to the total variance of a swap.
- This is particularly true for events happening within a short time-frame, for which traders feel like they can take a view on the relative value between one day and another. Given that people can predict a lot more in the short end, telling one day from another, the term structure of volatility is a lot more granular.
- This must be taken into account when interpolating across two expiries, by distributing the variance accrued between two points of the grid differently for every day in between. This is commonly referred to as "event day modeling", and it boils down to assigning appropriate weights to signal how much the clock for variance should tick on particular events relative to "standard" days.
- For example, assigning a weight w_i to each day in between two expiry grid points A and B, one can interpolate on an intermediate day t_j as

$$t_j\,\sigma^2(t_j) = \lambda\, t_A\,\sigma^2(t_A) + (1-\lambda)\, t_B\,\sigma^2(t_B)\,, \quad \text{where } \lambda = \frac{\sum_{i<=j} w_i}{\sum_i w_i}\,. \qquad (32)$$

Hedging

- In theory, the parameters ρ, β, and ν should be stable, in practice most of the time they are recalibrate on a monthly basis or so. On the other hand, the forward of the swap and α are clearly stochastic, and they liberally change all the time, even intraday. For hedging we will then have to neutralise the derivatives of the value of our portfolio wrt to S (delta) and α (vega).
- In case it were too expensive to zero out all vega and delta at the same time, one can try exploit the correlation between the moves of the forward and the volatility predicted by the model, to try to hedge "statistically" one with the other, thereby reducing the overall variability of the value of the portfolio.

$$\begin{aligned}
&\partial_S V_t \, dS_t + \partial_\alpha V_t \, d\alpha_t = \\
&= \partial_S V_t \, dS_t + \partial_\alpha \, V_t \, \nu\alpha_t \, (\rho \, dW_t + \sqrt{1-\rho^2} dZ'_t) = \\
&= \partial_S V_t \, dS_t + \partial_\alpha V_t \, \nu\alpha_t \, (\rho \frac{1}{\alpha_t S_t^\beta} dS_t + \sqrt{1-\rho^2} dZ'_t) = \\
&= \left[\partial_S V_t + \partial_\alpha V_t \, \frac{\rho\nu}{S_t^\beta} \right] dS_t + \partial_\alpha V_t \, \nu\alpha_t \, \sqrt{1-\rho^2} dZ'_t
\end{aligned} \tag{33}$$

- In the above, we have decomposed the variation of the value of our portfolio into two orthogonal components along dS_t and dZ'_t.

Hedging II

- The previous equation shows that whenever there's some residual vega in the portfolio, it's beneficial to try to hedge through delta the variation in value associated to the most likely variation in α following a given variation of the forward. The corrected delta is

$$\partial_S V_t + \partial_\alpha V_t \, \frac{\rho\nu}{S_t^\beta} \, . \tag{34}$$

- The hedging with the corrected delta is more effective the higher the correlation, in line with the intuition that the forward can capture more of the variability in the volatility (the limit being a local vol model). In fact, when $\rho \to 1$ the unhedged term in $\mathrm{d}Z_t'$ vanishes as expected.
- One can see that two different models calibrated to the same smile give the same min-var delta as long as they have the same backbone (e.g. $\mathrm{d}\sigma/\mathrm{d}S$).

- Jesper Andreasen and Brian Huge, *Expanded forward volatility*, Risk (2013).
- Emmanuel Derman and Iraj Kani, *Riding on a smile*, Risk (1994).
- Bruno Dupire, *Pricing with a smile*, Risk (1994).
- Patrick S. Hagan, Deep Kumar, Andrew S. Lesniewski, and Diana E. Woodward, *Managing smile risk*, Wilmott Magazine (2002), 84–108.
- Riccardo Rebonato, Andrey Pogudin, and Richard White, *Delta and vega hedging in the sabr and lmm-sabr models*, Risk (2009).