Interest rate derivatives in the negative-rate environment

Pricing with a shift

26 February 2016

Contents

The motivation behind negative rates

The rationale is to encourage investors to borrow money and invest into the economy

The recent financial crisis that started in August 2007 unfolded the trustworthiness among counterparties as one of the key concerns in financial transactions. In the height of the crisis large financial institutions collapsed while the interdependencies of one institution on another led to a widespread propagation of the default risk.

The credit quality of the counterparty become thereafter an integral part of the market risk. For many, trading became either too risky or too expensive and complex (under the pricing of the credit risk). Since then, the financial crisis infected countries of the hard core of the Eurozone.

In order to avoid that this new environment dominated by the credit-quality brings a halt into the economy, the Central Banks and particularly ECB applied some exceptional measures. First, in a series of ECB decisions from 2008 to 2011 interest rates were gradually lowered and therefore borrowing cash became cheaper. The rationale behind this is to encourage investors to borrow money and invest into the economy, which would therefore find the funds and grow. In June 5th 2014 however, ECB took this a step further by setting a key interest rate to minus 10 basis points. A negative rate implies that leaving money at rest in a bank would result in a loss. Therefore ECB would, in fact, punish investors for holding their cash. With this move the central bank aims to inspire investors further to bring in new money in the economy to help activity surge. The use of negative rates is an unconventional tool of economic policy but not unprecedented. In recent times the central banks of Switzerland, Denmark and Sweden have also taken the decision to set some of their key interest rates in a negative territory.

In the figure below we see two forward spot curves, which represent lending rates. We find the CHF 6M curve and the EUR 3M curve as of 31st Dec 2015. We see that the spot rates are negative in these curves.

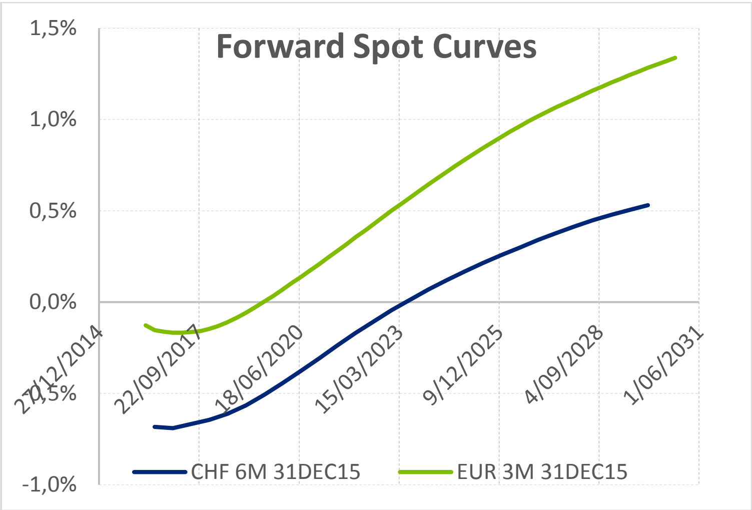

Valuation challenges in the negative rate environment

The existence of an interest-rate is linked to the fact that a lender requires a premium for undertaking the risk of lending money, hence it is logical that an interest rate is modelled to be positive. Traditionally, the occurrence of a negative rate as an outcome of a pricing model was seen as a weakness of that model.

The user must find a way to extrapolate the market-quoted volatilities into the negative domain

Examples of models in this category include the Hull-White model which assumes that the underlying (the short-rate) follows a mean-reverting process and it is a Gaussian-distributed random variable. As the domain of a Gaussian random variable is the entire real axis ranging from minus to plus infinity, all interest rates values would be in principle possible under this model. This might seem slightly unrealistic however, as the market has never experienced an interest rate set below minus 30 basis points. Therefore, assigning probabilities below this mark would be hard to justify. In the classic textbook of Brigo-Mercurio, in the section describing the short-rate dynamic of the Hull-White model, one reads “the theoretical possibility of *r* going below zero is a clear drawback of the model“[1]. However, on a counter-argument, such probabilities exist only in the tails of the Gaussian distribution and they would only play a minor role in the pricing of derivatives. As the Hull-White model leads to closed-form formulas for simple vanilla derivatives, it has been one of the most popular models. A similar model that assumes that interest-rates can be negative is the so-called “Bachelier” model, conceived by one of the early developers of option pricing theory Louis Bachelier in 1900.

An alternative model to Hull-White that assumes that the forward rate is strictly non-negative is the Black model. According to the Black model the forward rate is a lognormally-distributed random variable. By the very principles of this model, interest-rates can never attain a negative value. This model also leads to closed-form formulas for the pricing of simple derivatives and thus is ranks highly among the models preferred by practitioners.

In the current negative rate environment there is a number of challenges in the use of some of the traditional models. For example, according to a Black model the price of a simple cap option depends, among various other factors, on the logarithm of the forward rate. However, if the market-quoted forward rate is negative then the logarithm is undefined. As a result, this model cannot give an answer. This is a series drawback of this model.

A related challenge is the fact that market quotes for volatilities of negative strikes do not always exist. This means that the user must find a way to extrapolate the market-quoted volatilities into the negative domain.

[1] See D Brigo and F Mercurio, “Interest-rate models: Theory and Practise” Springer 2001

Market standard models for negative rates

The Hull-White, Bachelier and Black model owe their popularity to the existence of a closed-form formula for the pricing of vanilla interest-rate derivatives. This implies that pricing of simple products can be done quickly and accurately. In fact, the Black model has some highly desired features that make practitioners seek ways to remedy its breakdown for negative rates. One such remedy is the inclusion of a "shift" to the forward rate. Although there is no clear consensus as to the exact value of the shift, it should however be such that the logarithm of the forward rate plus the shift is well-defined.

In the following table we find a summary of their advantages versus disadvantages.

Model	Probability density	Advantages	Disadvantages
Hull-White, Bachelier	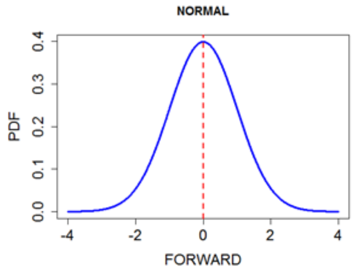	Closed-form formula exists. Model can be readily used.	Extreme negative values are possible (with small probability)
Shifted Black	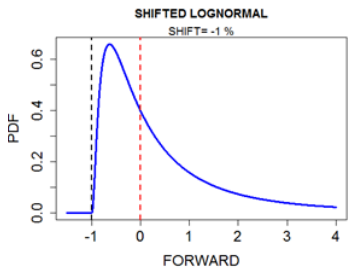	Closed-form formula exists. Negative values below the shift are not possible.	Requires as input an appropriate volatility that may not be quoted in the market.

The shifted Black model is often also called displaced diffusion, owing to the fact that it can be described as a diffusion (Geometric Brownian motion) whereby the main trend is displaced by a shift. The disadvantage of this model is the fact that one of its key inputs, namely the volatility, is not readily available from the market. Instead, it needs to be constructed.

The construction (extrapolation) of the market volatility surface to the negative-strike domain can be done using the SABR model. The SABR model is a stochastic volatility model and it is the market standard tool for interpolating on the volatility surface. Notice that in this strategy the shifted Black model is merely a quoting device and is no longer used to model the underlying. In order to allow that the output volatilities of the SABR model can be used by the Black model, one has to apply a similar shift to the SABR model. In this context the (shifted) Black model becomes a mere quotation device, rather than a pricing model.

The main strategy then for using the Black model in the negative-rate environment is:

- Calibrate the shifted SABR model using the market-quotes (on positive strikes)
- Extend the surface to the negative domain using SABR
- Retrieve a volatility from the shifted SABR model and use it in the shifted Black model

In the next section we outline the main characteristics of the shifted-SABR model.

Shifted SABR model

The SABR model introduced by Hagan and collaborators in 2002[2] is a stochastic volatility model that couples the forward rate and its volatility according to the following processes:

$$dF_t = \alpha \cdot {F_t}^{\beta} \cdot dW_t$$
$$d\alpha = v \cdot \alpha \cdot dZ_t$$
$$E^{Q_T}[dW_t \cdot dZ_t] = \varrho dt$$

The first equation describes the evolution of the forward rate. It contains no drift and the volatility is equal to α. The parameter β allows the model to switch between a lognormal-like process with $\beta=1$ and a normal-like process with $\beta=0$. The volatility of the forward rate α is itself a stochastic variable, driftless and with volatility equal to v. The two stochastic processes are driven by the Gaussian variables W_t and Z_t which are coupled by a correlation parameter ρ. The volatility of the forward rate has a starting value equal to α_0.This model is therefore described by four parameters: α_0, β, ρ and v. From now, and with a slight abuse of notation, we will refer for simplicity to the parameter α_0 as α.

There are a number of extensions to this model. For example, the ZABR model[3] assumes that the volatility of the volatility (here denoted by v) is a "local" volatility, i.e. a function of time and α.

The SABR model owes its popularity to the fact that it can lead to a closed-form expression of the Black implied volatility, as a function of the four parameters. This means that the SABR volatility can be used into the Black formula in order to give the price of a caplet. This is a particularly desirable feature for trading or risk-management systems, as it allows quick and accurate pricing.

The formula for the SABR implied (Black) volatility can be found in Hagan's paper and it is the following:

$$\sigma_B(K,f) = \frac{\alpha}{(f\beta)^{\frac{(1-\beta)}{2}}\left\{1+\frac{(1-\beta)^2}{24}log^2\frac{f}{K}+\frac{(1-\beta)^4}{1920}log^4\frac{f}{K}+\cdots\right\}} \cdot \left(\frac{z}{\chi(z)}\right)$$
$$\cdot\left\{1+\left[\frac{(1-\beta)^2}{24}\frac{\alpha^2}{(fK)^{1-\beta}}+\frac{1}{4}\frac{\beta\rho\alpha v}{(fK)^{\frac{(1-\beta)}{2}}}+\frac{2-3\varrho^2}{24}v^2\right]\right\}\cdot T+\cdots$$

where z and $\chi(z)$ are abbreviations of

$$z = \frac{v}{\alpha}(fK)^{\frac{(1-\beta)}{2}}log\frac{f}{K}$$
$$\chi(z) = log\left(\frac{\sqrt{1-2\rho z+z^2}+z-\varrho}{1-\varrho}\right)$$

This formula albeit complex is easy to code and can give instantaneously implied volatilities using as input the four SABR parameters, the forward rate f, the strike K and the time to maturity T.

The derivation of this formula is based on certain truncations and therefore this formula is not exact. It is a good approximation for small values of the variance v^2T. However, due to these approximations, this formula leads to important errors. These errors become apparent close to the zero-strike limit. Close to this limit the probability density function of the forward rate becomes negative, which is unnatural. An equivalent way to see the breakdown of the SABR model is to price butterfly spreads which due to the positivity of the convexity of the cap payoff should remain positive.

[2] P Hagan, D Kumar, AS Lesniewski and DE Woodward, "Managing Smile Risk", Wilmott Magazine, July 2002
[3] J Andreasen and BN Huge, "ZABR- Expansion for the masses", Available at SSRN 1980726, 2011

As an illustration of this model failure close to the zero strike axis we present in the figure below the probability density function of the SABR for a particular choice of parameters.

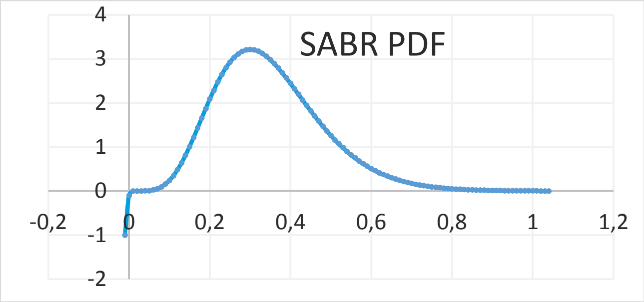

Despite this serious problem of obtaining a distribution function that is not well-defined, the SABR model is still widely used. One reason for this is the fact that it is able to accommodate a wide range of shapes of the volatility surface. Its four parameters can control the convexity, the location of the ATM point, and the skewness of the smile. There is some degree of overlap in the role that the four parameters play. For this reason it is a common market practice during the calibration of the SABR model that one of the four parameters is kept fixed and the calibration is run across the remaining three parameters.

The shifted SABR model is similar to the classic SABR model apart from the fact that a shift parameter "*b*" is introduced in the stochastic process of the forward rate:

$$dF_t = \alpha(F_t + b)^{\beta} \cdot dW_t$$
$$d\alpha = v \cdot \alpha \cdot dZ_t$$
$$E^{Q_T}[dW_t \cdot dZ_t] = \varrho dt$$

Conveniently, it turns out that including the shift in the formula of the Black SABR volatility is not too difficult: It is always coupled in an additive way to the forward rate and to the strike in the expression of the implied vol. As the variable *b* shifts the forward rate and the strike the problem of the negative probabilities is also shifted further along the negative rate axis. Because of this, we can now invoke the adapted SABR formula in order to obtain a volatility for a negative strike.

Caplet stripping

By its design, the SABR model outputs the price of the most basic vanilla option, which is a caplet. It does not output the value of a cap. However, the market is not quoting caplets but caps. Therefore, a conversion is needed from the quoted cap volatilities to caplet volatilities.

Since a cap consists of more than one caplet, there is some freedom in the choice of the caplet volatilities that can collectively reproduce the price of a cap. The science (and, to some extent, art) of generating caplet volatilities from cap volatilities is the so-called caplet stripping. A more complete investigation into this subject can be found in Hagan[4].

Here we follow the methodology described by Bloomberg[5] which is best explained in an example.

Let us consider that the underlying tenor is the 6M forward rate and we are interested in obtaining the various caplet vols from the cap vols. Our first assumption concerns the values of the caplet vols that compose the same cap. For example, for a 1Y cap which would consist of two caplets, we can assume either that (i) the 6M and 1Y caplet vols are equal, or (ii) the 6M and 1Y caplet vols show a term structure. There is no “correct” assumption here, as the caplet volatilities are not traded instruments. In this case, we may hypothesize that the 6M and 1Y caplet volatilities are equal. Let us denote this by σ_{1Y}. Then this value can be obtained by solving the following equation

$$CAP_{1Y}(K, \Sigma_{1Y}) = CAPL_{6M}(K, \sigma_{1Y}) + CAPL_{6M}(K, \sigma_{1Y})$$

where we have denoted by Σ_{1Y} the 1Y cap volatility and by σ_{1Y} the 6M and 1Y caplet volatility. In this equation K represents a fixed strike of the vol table. One can continue bootstrapping the caplet volatilities in this fashion for further expiries. For example, the next series of caplet vols can be obtained by solving

$$CAP_{2Y}(K, \Sigma_{2Y}) = CAP_{1Y}(K, \Sigma_{1Y}) + \ CAPL_{18M}(K, \sigma_{2Y}) + CAPL_{2Y}(K, \sigma_{2Y})$$

In this equation we see that the “forward” cap (cap at 2Y minus cap at 1Y) determines fully the caplet volatility.

The situation is slightly more complex if we consider the ATM strike, instead of a strike at a fixed percentage. This is because the location of the ATM strike, which is the forward rate, changes at every expiry. For example, in order to obtain the 2Y caplet volatility at the 2Y ATM cap strike we need to solve the following equation:

$$\begin{aligned} CAP_{2Y}(K_{ATM}^{2Y}, \Sigma_{ATM}) = CAPL_{6M}\left(K_{ATM}^{2Y}, \sigma_{1Y}(K_{ATM})\right) + CAPL_{1Y}\left(K_{ATM}^{2Y}, \sigma_{1Y}(K_{ATM})\right) \\ + CAPL_{18M}\left(K_{ATM}^{2Y}, \sigma_{2Y}(K_{ATM})\right) + CAPL_{2Y}\left(K_{ATM}^{2Y}, \sigma_{2Y}(K_{ATM})\right) \end{aligned}$$

The difficulty in solving this equation for the 2Y caplet volatility is that the 1Y caplet volatility at the 2Y ATM cap strike will not be known, unless this particular strike is exactly one of the fixed-percentage strikes, which is unlikely.

The strategy to go forward, as indicated by Bloomberg, is to use the SABR model of the previous expiry (in our example the 1Y expiry) in order to interpolate across the 1Y caplet vol surface and obtain the 1Y caplet vol at the 2Y ATM cap strike, namely $\sigma_{1Y}(K_{ATM}^{2Y})$. Once this quantity is evaluated, then one can bisect the above formula to obtain the 2Y caplet vol $\sigma_{2Y}(K_{ATM}^{2Y})$.

---

[4] P Hagan and M Konikov, “Interest Rate Volatility Cube: Construction And Use”, available at http://www.javaquant.net/papers/PandorasCube.pdf

[5] Joshua X. Zhang, Zhenyu Wu “Bloomberg Volatility Cube”, Bloomberg white paper, 2015

Particularly in the current negative-rate environment the ATM point is one of the most important quotations in the volatility surface as it is the closest to the area of the negative strikes. In contrast, the fixed strikes are usually quoted as of 1% onwards which is far from the forward rates, in the current standards. For this reason, the consideration of the ATM quote in the above bootstrap method is crucial in the success of the calibration.

Notice also that, through the inclusion of the ATM point in the bootstrapping process, every previous expiry plays a significant role to the calibration quality of every next expiry. This implies that calibration errors will accumulate across expiries. It is therefore crucial that calibration is very accurate at least across the first

Deloitte valuation tool

Deloitte uses data from Bloomberg's BVOL CUBE to calibrate a shifted SABR model along the lines of the previous sections. Using this methodology one can price caps, floors, swaptions and CMS options quoted at zero or at negative strikes.

The calibration of the SABR model is done using numerical routines that search for the global solution in the space of parameters, such as simulated annealing[6]. Furthermore, the landscape of the error-surface is examined using heat maps, such as the following:

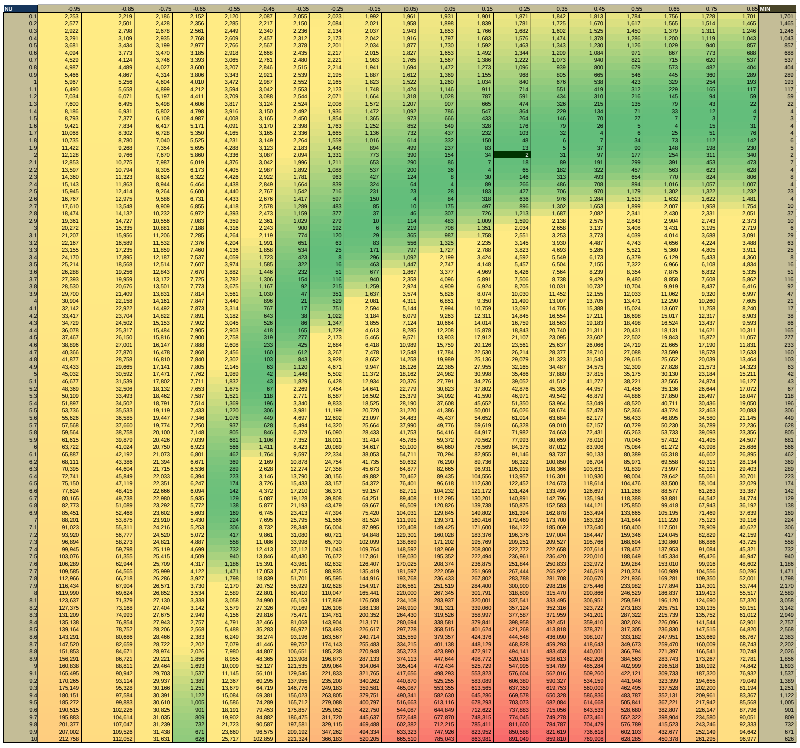

Columns in this table correspond to values of the parameter ρ which ranges from [-1,1], while rows correspond to the values of the parameter $\nu \in (0, \infty)$. The surface in the above figure is a cross-section for $\alpha = 1/2$. Green regions indicate areas of low calibration error, while red regions indicate areas of large error. The global minimum is located at the highlighted cell. This heat map allows us to have a visual inspection of the structure of the error surface. Local-search algorithms will not always find the optimal solution if regions of low error are separated by large error-barriers. Indeed this is the case in the above figure where we see that a green area emerges at the bottom-left corner and it is not connected to the optimal solution.

To test the calibration quality of the tool further we examine:
The positivity of the probability density function of the SABR forward. This can be done by differentiating twice the output cap price with respect to strike, namely $(F_t) \sim \frac{\partial^2 CAP_t}{\partial K^2}$ [7]. Note that this is similar to testing the positivity of the so-called "butterfly spread".
The bootstrapping quality of the stripping of the caplet volatilities.
The matching between the input vs output cap prices.

[6] S Kirkpatrick, CD Gelatt, MP Vecchi, "Optimization by Simulated Annealing", Science 220 (1983), pp. 671
[7] DT Breeden and RH Litzenberger "Prices of state-contingent claims implicit in option prices" The Journal of Business 51 (1978) 621

The following figure illustrates the probability density function of the shifted SABR model. We see that the problematic left-tail features of the SABR PDF are pushed further after the shift (which is set at 1%).

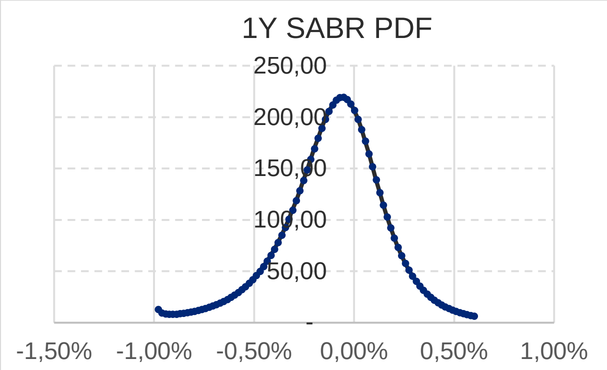

The following figure shows the result of the SABR calibration for the 1Y tenor on caplet vols. Calibration of this tenor required matching the 1%, 1.5% and the ATM volatilities, as we believe these correspond to quotes with the highest liquidity level at this tenor. We compare the SABR output volatility against the bootstrapped caplet volatility.

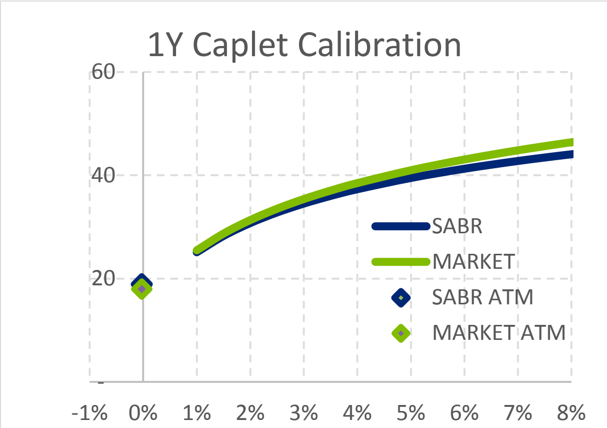

The following tables show the difference of the input market cap prices and the resulting SABR cap prices across various tenors and strikes. Numbers are expressed in EUR (valuation date 31 Aug 2015, CHF on 3M). Differences here are of the order of a few basis points (the notional is 10,000,000 EUR), indicating a good quality of calibration. Notice that the ATM strike of the 1Y tenor is already negative.

Cap Prices SABR

Tenor	Expiry Date	Settlement Date	ATM K	ATM	1.00%	1.50%	2.00%	2.50%	3.00%	3.50%	4.00%
1Yr	8/31/2016	9/2/2016	-0.03%	4,000	1.125	0	0	0	0	0	0
2Yr	8/31/2017	9/4/2017	0.01%	14,974	570	214	98	51	29	18	12
3Yr	8/31/2018	9/4/2018	0.10%	41,770	6,739	2,871	1,353	695	383	224	137
4Yr	8/31/2019	9/3/2019	0.23%	83,279	26,151	13,882	8,166	5,233	3,587	2,590	1,949
5Yr	8/31/2020	9/2/2020	0.37%	141,604	69,355	41,475	25,836	16,659	11,071	7,563	5,301
6Yr	8/31/2021	9/2/2021	0.52%	214,672	136,804	86,938	55,791	36,142	23,690	15,758	10,665
7Yr	8/31/2022	9/2/2022	0.67%	297,584	227,695	152,863	102,937	69,587	47,335	32,476	22,514
8Yr	8/31/2023	9/4/2023	0.80%	381,791	329,614	228,706	158,670	110,234	76,870	53,918	38,099
9Yr	8/31/2024	9/3/2024	0.93%	473,168	451,234	321,717	228,460	161,723	114,298	80,776	57,157
10Yr	8/31/2025	9/2/2025	1.05%	565,622	583,315	423,909	305,898	219,287	156,347	111,012	78,593

Cap Prices BBG

Tenor	Expiry Date	Settlement Date	ATM K	ATM	1.00%	1.50%	2.00%	2.50%	3.00%	3.50%	4.00%
1Yr	8/31/2016	9/2/2016	-0.03%	3,734	1.373	0	0	0	0	0	0
2Yr	8/31/2017	9/4/2017	0.01%	15,618	522	206	106	64	43	31	23
3Yr	8/31/2018	9/4/2018	0.10%	41,540	6,190	2,611	1,254	672	393	246	163
4Yr	8/31/2019	9/3/2019	0.23%	83,834	26,593	13,856	7,852	4,801	3,130	2,149	1,539
5Yr	8/31/2020	9/2/2020	0.37%	143,725	68,584	39,464	23,958	15,413	10,452	7,415	5,456
6Yr	8/31/2021	9/2/2021	0.52%	216,734	134,771	82,769	52,280	34,424	23,700	16,975	12,580
7Yr	8/31/2022	9/2/2022	0.67%	297,644	223,552	144,678	95,000	64,376	45,287	33,040	24,889
8Yr	8/31/2023	9/4/2023	0.80%	382,146	326,640	218,964	147,514	101,469	71,926	52,648	39,668
9Yr	8/31/2024	9/3/2024	0.93%	472,196	449,229	310,603	214,573	150,305	107,861	79,605	60,386
10Yr	8/31/2025	9/2/2025	1.05%	564,106	583,755	413,569	291,693	207,604	150,667	112,093	85,575

How we can help

Our team of quants provides assistance at various levels of the pricing process, from training to design and implementation.

Deloitte's option pricer is used for Front Office purposes or as an independent validation tool for Validation or Risk teams.

Some examples of solutions tailored to your needs:
- A managed service where Deloitte provides independent valuations of vanilla interest rate produces (caps, floors, swaptions, CMS) at your request
- Expert assistance with the design and implementation of your own pricing engine
- A stand-alone tool
- Training on the SABR model, the shifted methodology, the volatility smile, stochastic modelling, Bloomberg or any other related topic tailored to your needs

The Deloitte Valuation Services for the Financial Services Industry offers a wide range of services for pricing and validation of financial instruments.

Why our clients haven chosen Deloitte for their Valuation Services:
- Tailored, flexible and pragmatic solutions
- Full transparency
- High quality documentation
- Healthy balance between speed and accuracy
- A team of experienced quantitative profiles
- Access to the large network of quants at Deloitte worldwide
- Fair pricing

Contacts

Nikos Skantzos

Director
Enterprise Risk Services
Diegem

T: +32 2 800 2421
M: + 32 474 89 52 46
E: nskantzos@deloitte.com

Nicolas Castelein

Director
Enterprise Risk Services
Diegem

T: +32 2 800 2488
M: +32 498 13 57 95
E: ncastelein@deloitte.com

Kris Van Dooren

Senior Manager
Enterprise Risk Services
Diegem

T: +32 2 800 2495
M: + 32 471 12 78 81
E: kvandooren@deloitte.com

George Garston

Consultant
Enterprise Risk Services
Diegem

T: +32 2 800 2087
M: + 32 471 82 93 18
E: ggarston@deloitte.com