First published as Commerzbank Rates Radar, 25 August 2021

Jessica James, Michael Leister, Christoph Rieger

The price of inflation uncertainty

Inflation and inflation risks - why are they relevant?

Inflation has rarely held the market's attention so closely as it does today. We seem to stand on the cusp of a regime change, from the lowflation of the past decade to an unknown higher range. Though there is talk of a 'return to normal', lowflation coupled with unprecedented liquidity injections by central banks has in fact become the 'new normal' and markets have become accustomed to life support. Should central banks react according to their mandates, a new, higher inflation regime would hit institutions who have been relying on low cost credit, and could cause many of the so called 'zombie' companies to go under - and also challenge the sustainability of ever rising public debt stocks. Savers, who are facing an erosion of the purchasing power of their savings would rejoice longer-term when rates normalise, and retail banking would have a chance to return to a sustainable 'borrow and lend' model.

The likely level of this higher inflation regime and inflation expectations are thus critically important, and widely discussed[1]. If the spike in inflation does not trigger higher inflation expectations and second-round effects, policy makers and other economic agents can afford to look through it. Things will look very differentl if rising inflation feeds through to longer-dated expectations, changing the underlying inflation dynamics.

In view of the 2% inflation targets by most major central banks, it will make a huge difference if underlying inflation settles again at levels closer to, say, 1% or 3%. And at times when inflation views diverge, so-called linear indications from **inflation-linked bonds or swaps** only tell half the story if the expected value is, say, at 2% while most people think the actual outcome is either below 1% or above 3%.

So indications of future levels are eagerly sought, and one popular method of obtaining probabilities of a rise to various levels has been the **inflation options** market. By using methods developed for liquidly traded option markets like equities or interest rates, it becomes possible to derive implied probabilities of inflation changes in the future from the prices of inflation options. We assess the value and robustness of these indications derived from inflation options, and show that while some parameters are stable and statistically valid, this is not the case for all.

Inflation risks from bonds, swaps and options

Underlying concepts

The first inflation linked bonds were issued in 1780 by the Commonwealth of Massachusetts to pay soldiers in the Revolutionary War. Owing to over-enthusiastic issuance of paper money during preceding decades, mostly to pay for various other conflicts, commonwealth government money traded at a steep discount to sterling, and was not highly regarded. Thus soldiers were paid with bonds whose coupon and redemption were linked to the price of a basket of goods including bushels of corn and specified weights of beef, sheep's wool and sole leather. The variation in price of this basket formed an effective price index, as shown below. The Commonwealth of Massachusetts clearly had a severe inflation problem!

Commonwealth of Massachusetts historical inflation

Value of a basket of goods, scaled so that basket value = 1 at time of issue, paper issued to soldiers in lieu of cash

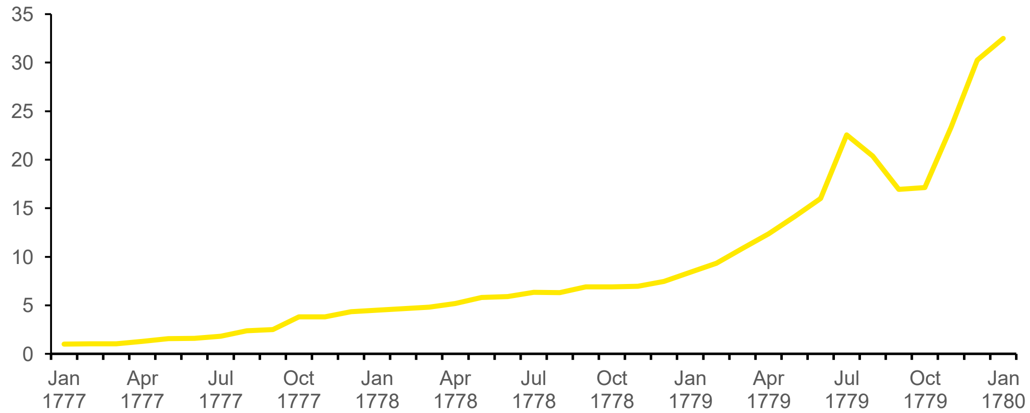

Source: Commerzbank Research, https://www.nber.org/system/files/working_papers/w10183/w10183.pdf

Fast forward a few centuries, and high inflation in some Emerging Market economies drove the creation of inflation-index linked bonds in the 1960s. The UK's high inflation levels in the 1980s prompted the creation of inflation-linked gilts, and many other developed markets followed. In 1997 the US Treasury joined the party and began to issue Treasury Inflation Protected Securities (TIPS).

All of these bonds work in similar ways. They are issued with a principal amount and coupon. The principal and coupons are adjusted according to the evolution of a chosen inflation index. The principal is usually floored at the original amount - thus if inflation is negative, the holder will still receive at least the original amount invested, though the coupons are not protected other than being floored at zero.

Focussing on the European market for inflation linked instruments, since its inception it has evolved into a mature and liquid market. In the Eurozone, the benchmark index is the non-seasonally adjusted HICP ex tobacco (Harmonised Indices of Consumer Prices, HICPx). Inflation swaps referenced to HICPx are actively traded across maturities and inflation options are also traded. Inflation linked bonds are usually known as 'linkers' or ILBs.

Index Mechanics

In the Eurozone, consumer price indices (CPIs) are only published with a time lag to allow time to gather the relevant data and compute the indices. The (unrevised) CPIs are therefore applied with a three month time lag to compute cash-flows in the Eurozone linker market.

An example: The February 2021 HICPx index will be used for May. To be precise, this index will be used for calculating the reference index for the first day of the month, in our example 1 May. Accordingly, the March 2021 HICPx index will be used as the reference index for 1 June. For every trading day in between those dates, a daily reference index is calculated by linear interpolation.

This is done via the following formula:

$$RefCPI_{t,m} = CPI_{m-3} + \frac{(d-1)}{D_m} * (CPI_{m-2} - CPI_{m-3})$$

where CPI_m denotes the respective lagged price index level, D_m the number of days in the settlement month and d the respective day of the settlement month. For example, say the February HICPx (used for 1 May) is 115.87, and the March index (applied to 1 June) is 117.20. The reference index for 14 May is therefore 116.42774. The reference index for the first settlement day of an index-linked bond is called the base index.

The question of how much inflation has taken place since a bond was launched is measured by the progression of the reference index relative to the bond’s base index. In order to do so, the Index Ratio (IR) is calculated. This ratio is the key concept linking the real and the nominal world.

The index ratio and inflation compensation

All cash-flows will be adjusted by the index ratio on the respective payment date. For example, the full invoice price of a linker trading at clean price of 110, accrued interest of 5 and an IR of 1.1 is equal to (110+5)*1.1 = 126.5. The inflation adjustment can be highly significant and influence total return considerations, in particular for short-dated or seasoned ILBs.

Assuming positive inflation throughout the life of the bond (i.e. a positive inflation accrual), coupon payments and the principal thus increase in line with prices over time (see charts below). In this way, investors get compensated ex-post for the loss of purchasing power. Accordingly, while the future cash-flows and thus the return of an index-linked bond is uncertain in nominal terms as future inflation is unknown, the return is certain in real terms, being equal to the coupon of the linker (assuming no default of the issuer). Accordingly, nominal bonds offer a certain return in nominal terms, while real returns are uncertain.

Note also that for EMU linkers the inflation-uplift on the principal will never be below one – even in the case of deflation. In other words, the principal amount is floored at par. This does not apply to the coupon, however.

Coupon on a 10y inflation linked bond
Nominal coupon = 1.5%, inflation assumed 2% p.a.

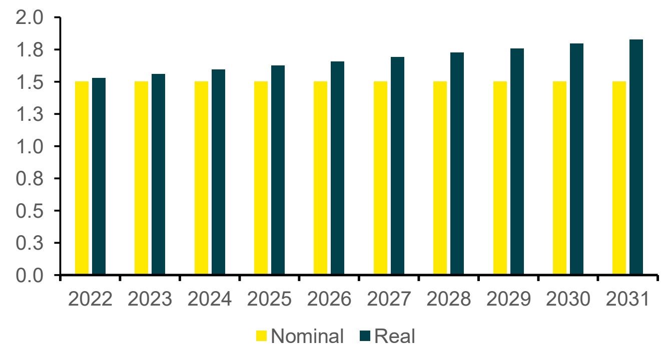

Source: Commerzbank Research

Principal repayment on an inflation linked bond
Coupon 1.5%, inflation 2% p.a.

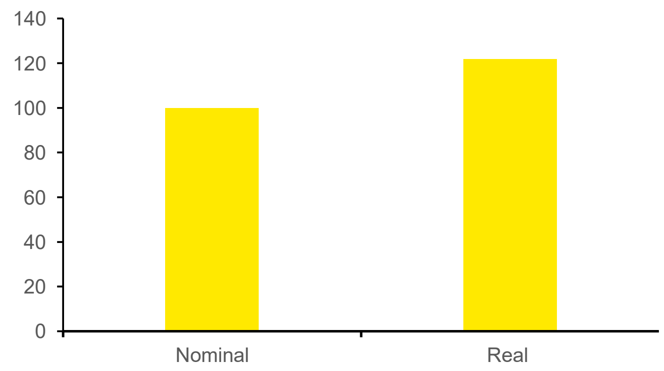

Source: Commerzbank Research

Care has to be applied when interpreting real yields of index-linked bonds. As for nominal bonds, real yields are typically derived from clean prices. Hence, they do not reflect the inflation accrual and can therefore provide a misleading yield to maturity – or internal rate of return (IRR) – in nominal terms. This is particularly relevant when comparing the yield/IRR of a nominal and an index-linked bond, as ignoring the inflation accrual can provide a false picture of the relative (expected) total returns of the bonds in nominal space; particularly for seasoned ILBs with a short maturity.

To obtain the nominal yield/IRR (r_N), the NPV of the inflation-uplifted cash flows of the ILB have to equate its all-in (inflation-uplifted) dirty price, just like for the yield of a nominal bond.

Technically this is achieved via the following formula:

$$P_{All-in} = \sum_{t=1}^{T} \frac{CPN * IR_t}{(1+r_N)^t} + \frac{100 * max\{IR_T, 1\}}{(1+r_N)^T}$$

Note that to compute the above formula, CPI forecasts for the cash flow dates are required. Market practice is to focus on the *real* yield to maturity of a linker r_{RY} . The latter does not require CPI forecasts and is computed as follows:

$$P_{All-in}/IR_0 = \sum_{t=1}^{T} \frac{CPN}{(1+r_{RY})^t} + \frac{100}{(1+r_{RY})^T}$$

This is also known as the Canadian model and market prices and yields are typically quoted this way (though without accrued interest, similar to nominal bonds). When an investor buys an ILB they nonetheless have to pay the inflation-uplifted price.

The concept of break-even inflation

Break-even inflation represents the link between the nominal and the real world. While linkers shield investors ex-post from inflation, nominal bonds also provide a certain degree of ex-ante protection against inflation. Below, a slightly simplified version of the famous Fisher equation, which highlights the fundamental link between real and nominal yields:

nominal yield = real yield + inflation expectation + (liquidity premium + risk premium)

nominal yield = real yield + break-even inflation

break-even inflation = inflation expectation + (liquidity premium + risk premium)

This concept helps us to get a handle on the following question: how much inflation is needed for investors to be indifferent between holding index-linked and nominal bonds, i.e. to break even? Break-evens can be interpreted in a number of ways:

- **Market-implied inflation expectations:** Break-evens indicate how much inflation the market expects over a given period of time. For example, a rise in break-evens generally indicates that the market adjusts its inflation expectation to the upside.

- **Inflation benchmark:** Break-evens are a valuable yardstick to put inflation views into context. If an investor views the markets' inflation expectations as too low, linkers are more attractive than their nominal peers.

- **Tradable product:** Interpreting break-evens as simply a spread between two bonds allows for direct trading of the markets' implied inflation expectation.

In market practice, break-even inflation is simply calculated as the spread between the nominal and real yield of instruments with similar maturity and credit quality. Which nominal bond should be used to compute the break-even rate? There is no mechanical rule to determine the nominal counterpart of an ILB, although it should typically be from the same issuer or of similar credit quality and match the ILB's maturity as closely as possible.

The left-hand chart below shows the break-even rate between the DBRei Apr 30 and a linear interpolation of two nominal bonds on the German curve, i.e. the on-the-run DBR Feb 30 and DBR Jan 31. The interpolation aims to overcome the maturity mismatch of the bonds.

Break-even curves are typically upward-sloping (see right-hand-chart below), reflecting higher term and inflation risk premia with increasing maturity.

Break-even inflation as a spread
Interpolated yield Bund Feb 30 and Bund Jan 31 minus (real) yield of DBRei Apr 30, in %

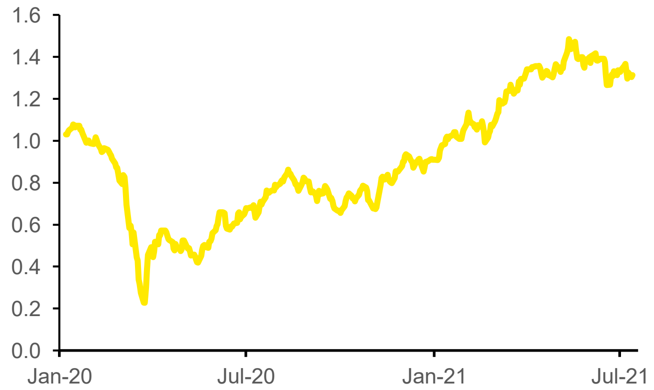

Source: Bloomberg, Commerzbank Research

Seasonality adjusted euro area break-evens
HICPx linker break-evens, in %

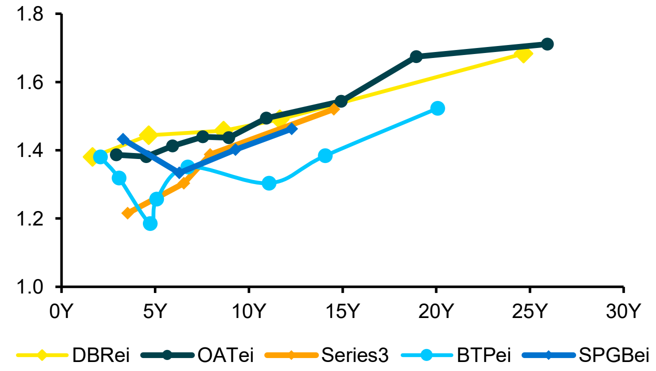

Source: Bloomberg, Commerzbank Research

The deflation floor

By construction, all cash-flows of index-linked bonds are adjusted by the index ratio (IR). The IR measures the cumulative inflation accrual since the bond was issued. Technically it equals the reference index for any particular date (e.g. HICPx for the Eurozone) relative to the reference index at the issue date of the linker.

Assuming positive inflation throughout the life of the bond, coupon payments and the redemption value therefore increase in line with prices. In this way, investors get compensated ex-post for the loss of purchasing power. Likewise, the IR drops in the case of falling prices. Thus, an investor entering a long position in linkers in a deflationary environment would face a downward adjustment on the cash-flows of the bond and therefore incur lower returns than a comparable investment in nominal bonds. In today's environment of ever more negative real yields, linker total returns can even be negative, as they are issued above par and redeemed at par.

To protect investors from redemption below par, inflation-linked government bonds in the Eurozone as well as in Sweden and the United States contain a par deflation floor. The floor ensures that investors receive par if the IR is less than one at maturity. In other words, the principal will not fall below par even in the case of cumulative deflation over the life of the bond. The deflation floor thus represents a put option on the price index with a strike price equal to the base index of the bond, or alternatively to an IR of one. The coupon is protected in the sense that it is floored at zero, and will not go negative.

Getting a handle on the deflation floor embedded in EMU linkers is important when pricing new issues or assessing the relative value of index-linked bonds.

Inflation Options

All of these index details above relate only to the inflation curve; the 'spot' and the 'forward' analogue equivalents in the inflation market. The curve which is actively traded and updated is the zero-coupon inflation curve - the 'zero-coupon' part referring to the fact that cashflows occur only on the maturity date - and are thus entirely a function of the HICPx level at this maturity date. So, we can get a good handle on the market's opinion of the forward path of inflation. We can even, by looking at the value of the deflation floor, derive a likelihood for deflation. But to derive probabilities for any other inflation scenario, we need option products.

Fortunately, these exist. We discuss the data availability in more detail below, but over the last decade or so, caps and floors which depend on the CPI indices of different countries have evolved. The deepest and most liquid market is probably the EUR market, which uses the HICPx (Harmonised Index of Consumer Prices ex tobacco) as its underlying index. Swaps are traded on the zero coupon curve derived from this index, and caps and floors are also somewhat liquid and active. In general these come in two varieties; YoY (year-on-year) which look at the change in inflation from one point in the future to a point a year further on, and zero coupon (ZC) which look at the change in inflation from now until the maturity date.

Commerzbank's in-house SABR model

In order to extract the implied probability distribution from quoted prices for HICPx caps and floors, a model is required.

Commerzbank is fortunate in having a sophisticated set of modelling tools for option valuation, which also yield information about implied probabilities for the future. The SABR model[3] is a very useful and flexible model with the flexibility to accommodate skewness and fat tails. It is defined by a linked set of equations, as below:

$$dX_t = \sigma_t (X_t + d)^\beta dW_t$$
$$d\sigma_t = \alpha\sigma_t dZ_t$$
$$dW_t dZ_t = \rho dt$$

[1]

Here X_t is the variable being modelled, W and Z are random processes, σ is the measure of volatility, α is the volatility of volatility (vol-vol), β is a tail parameter, and ρ is the correlation between the level of the modelled variable and its volatility. d is a displacement variable which shifts the whole distribution.

Although the interaction among them is complex, each parameter has a clear link to the property it measures. β controls tail width; a normal distribution has a β of zero, and a lognormal has a β of 1. α accounts for varying volatility levels, and ρ indicates the degree of parameter correlation which can drive some unexpectedly large market moves.

The model can be used to calculate premia for caps, floors or swaptions in the inflation market. The modelled variable X_t can be the index ratio for zero coupon options (ratio of inflation at maturity to inflation today) or the period-on-period ratio (ratio of inflation at one future date to inflation at another future date) which are useful for zero-coupon or YoY inflation options respectively.

The inputs to the model are the premia for a set of cap or floor options, which can be zero coupon or YoY, and also the zero coupon inflation swap curve. The outputs are flexible, and can include option premia for any reasonable strike, fitted parameter values, and the probability distribution function (pdf). The pdfs produced by SABR are highly useful in that they can express not only mean and standard deviation, but also the level of skewness and kurtosis implied by the input option prices. Skewness is a measure of the asymmetry of a distribution - distributions with positive skewness have more weight to the higher values, while negative skewness have more weight to the lower. Kurtosis measures the 'fatness' of the tails - its most common definition sets the kurtosis of a normal distribution to be zero, positive values indicate fatter tails than a normal, with negative values indicating thinner or more truncated tails than a normal.

In other words, we can benchmark the distribution extracted from inflation-options to other probability distributions (e.g. a simple normal distribution) and thus assess the reliability of option-implied inflation expectations - ie. if the resulting distribution captures actual market dynamics or is primarily a fitting/modelling exercise. Particularly in the context of overshooting (or deflation) concerns, the higher moments of the distribution and their robustness are central.

Such checks are also important as they address the fact that market-based measures of inflation expectations are subject to risk- and liquidity premia and hence do not only reflect genuine inflation expectations.

The Minneapolis Fed data set

The Federal Reserve Bank of Minneapolis generously makes available an extensive set of time series about US inflation probabilities. For each day since June 2009, it provides a value for the following parameters for the inflation pdf:

- Mean
- Standard deviation
- Skew
- Kurtosis
- 10% probability move
- 50% probability move (=median)
- 90% probability more

- Probability of inflation going to -1% or below (large decrease)
- Probability of inflation going to +3% or above (large increase)

This data set is available for tenors of 1,2, and 5 years, though 5 is the longest and best data set, and is the one we focus on here.

While this is a rich and interesting data set, sadly, it does not pin down the distribution! One can imagine that a double peaked distribution could mimic the parameters of a single peak, or that a degree of 'lumpiness' could distort the values. It is however possible to plot the pdf on their website for any selected day and maturity, which is useful.

How have they gone about creating the pdfs which they use to generate the above moments of the distribution? They base their method on a 2012 paper by Kitsul and Wright[4], which itself is based on a 1979 paper by Breeden and Litzenberger[5]. The method uses daily quotes on zero coupon inflation caps at strike prices of 1%, 2%, 3%, 4%, 5% and 6%, and on zero-coupon inflation floors at strike prices of -2%, -1%, 0%, 1%, 2% and 3%. In the Kitsul and Wright paper, the authors use these quotes to form implied pdfs for inflation, but without assuming normality. Instead, they initially assume that inflation will only have integer values, which is not as bad as it sounds, if one simply assumes that one will use these points to create a curve. Then, consider a butterfly portfolio that involves buying caps with strikes of k − 1% and k + 1% while shorting two caps with a strike of k%. This portfolio is a pure Arrow-Debreu security with a payoff of $1 if inflation is k% and zero otherwise, for any integer k. A risk-neutral investor will pay

$$e^{-r_n n} p_{k,n}$$

for this security, where r_n denotes the continuously compounded interest rate for *n* years and $p_{k,n}$ is the probability that inflation reaches a level of *k* after n years. They discount by the nominal government term structure. This mechanically yields risk-neutral probabilities of inflation being -2%, -1%, 0%, ... 5%.

For the tails of the pdf, if an investor buys an inflation cap at 5% and shorts one at 6%, then this investor receives $1 if inflation is 6% or more and 0% otherwise. This gives the probability of inflation being 6% or higher.

The same process works in the left tail as well, for example yielding the probability of inflation being -2% or lower. Thus they create pdfs over different maturities. For this process they use only zero coupon, not YoY, as mathematically the latter is more tricky, but it is a pity as the YoY data is generally more liquidly traded.

It seems likely that the Minneapolis Fed study used a similar method, possibly with different granularity for the option strikes dependent upon what data was available.

It is worth noting that the Minneapolis Fed is not the only Central Bank to extract inflation pdfs from inflation option prices, though as far as we are aware it is the only one to make such a generous data set available. One of the others is the Bank of England[6], which uses an approach more similar to that of Commerzbank. They use the same initial method by Breedon and Litzinger to link option prices with an implied pdf, but they interpolate the option prices to give a denser initial data set, which assists with the fitting process, though it does not in itself add information. Finally they allow the volatility smile (the graph of implied volatility levels vs moneyness) to have flexible properties modelled with the same SABR technique as that used by Commerzbank.

Empirical Observations and data limitations

Underlying data

It is clear from the methods outlined above that the essential element for deriving the market implied future probability distribution of inflation is a set of inflation option premia at different strikes. It is enlightening to go to the raw data and see what goes into the various elegant calculations! To understand what is available, we looked at the EUR inflation option market, which is the most liquid and frequently traded. There are also inflation option prices available for the US, and some individual European countries like the UK, but all of these are likely to have lower liquidity and data availability. We used Bloomberg as our data source, which is generally held to contain all the liquidly traded inflation option products in the Eurozone.

Initially, there seems to be a good coverage of strikes, if one searches by listed tickers. There appear to be inflation cap premia with strikes ranging from -1% to +6%, in 0.25% steps. These would pay out when inflation rises above the specified strike. Inflation floor premia vary similarly from -3% to +6%. They are available for multiple tenors up to 30 years. So at a cursory glance one might think that there was a rich data set available. It is worth noting that the ticker names for the various strikes follow no known pattern - the section of the name which indicates the strike goes 'A, 3, L' for cap premia of strikes '2.5%, 3.0%, 3.25%'. Having gone to the trouble of discovering the ticker names of inflation caps and floors for all the above strikes, it is then possible to look at the data coverage and quality.

Between 6% and -1%, for the caps, there are 24 tickers (not all steps of 0.25% exist). But 24 points is enough to plot a nice smooth curve. However, when one investigates the data feeds, sadly only 8 of these contain data. It is a similar story for the floors, again with only 8 populated strikes. So every pdf extracted, no matter how sophisticated the method used, has come from only 16 points, or 8 if only cap or floor data is used.

Data limitations

At this stage it is worth a quick look at how easy it is to extract parameters like mean and standard deviation, skew and kurtosis, from distributions, and see what the effect of a small number of data points might be.

As an example, we take the Weibull distribution, a useful analytical form which has parameters which may be adjusted to yield distributions with varying tail thickness, asymmetry, etc. The expression for the distribution is

$$f(x) = \begin{cases} \frac{k}{\lambda}\left(\frac{x}{\lambda}\right)^{k-1} e^{-\left(\frac{x}{\lambda}\right)^k}, & x \geq 0 \\ 0, & x < 0 \end{cases}$$

where k is the shape parameter and lamda the scale parameter of the distribution. For our own simple test we can set

$$\text{scale parameter } \lambda = 1$$

to simplify the various expressions, thus although it is mentioned here for completeness, it will not appear again. For this simplified case, you can see below what the distribution looks like for k=1.5 and k=2.0.

Weibull distribution

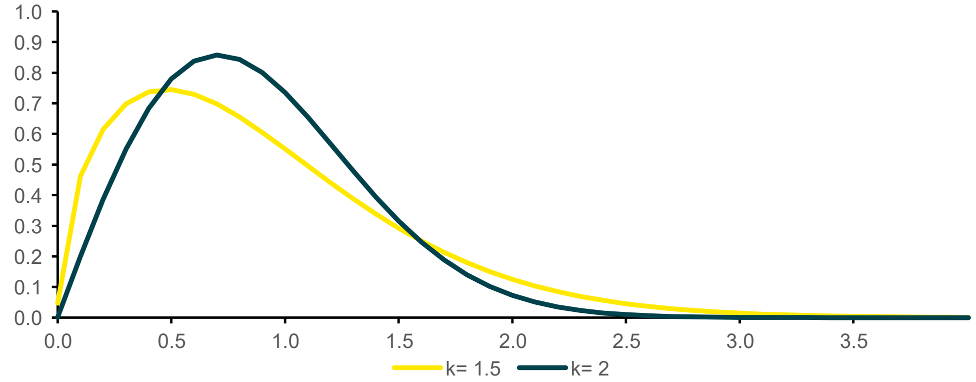

Source: Commerzbank Research

It is a very popular distributional form used in many applications - engineering, insurance, extreme value theory, weather forecasting and hydrology are just a few examples[7]. Interestingly, it was used by one of the authors to help describe the density variation at the edge of heavy atomic nuclei, in her PhD thesis[8]. One of its many useful features is that is has analytical forms for the values and standard errors of most moments, as we show below.

$$\text{mean } \mu = \Gamma\left(1 + \frac{1}{k}\right)$$

$$\text{Standard Deviation } \sigma = \left[\Gamma\left(1 + \frac{2}{k}\right) - \Gamma\left(1 + \frac{1}{k}\right)^2\right]^{1/2}$$

$$\text{skewness } \gamma_1 = \frac{\Gamma\left(1 + \frac{3}{k}\right) - 3\mu\sigma^2 - \mu^3}{\sigma^3}$$

$$\text{Kurstosis } \gamma_2 = \frac{\Gamma\left(1 + \frac{4}{k}\right) - 4\gamma_1\sigma^3\mu - 6\mu^2\sigma^2 - \mu^4}{\sigma^4}$$

$$\text{standard error for skewness } \sigma(\gamma_1) = \left[\frac{6N(N-1)}{(N-1)(N+1)(N+3)}\right]^{1/2}$$

$$\text{Standard error for kurtosis } \sigma(\gamma_2) = \left[\frac{4(N^2-1)\sigma(\gamma_1)}{(N-3)(N+5)}\right]^{1/2}$$

where *N* is the number of points in the data set. Note that the standard errors assume that there is a numerical data set for the distribution whereas the moments are analytical.

It can easily be seen that the standard errors depend strongly on the number of points in the fit. To check how many points are needed for a reliable estimate of mean, standard deviation, skewness and kurtosis, we generated numerical data for the distributions and the moments, with varying numbers of points, for k=1.5 and k=2, and then tabulated the results in the table below.

Numerical and analytical moments for k=1.5 and k=2 Weibull distribution

Numerical results for varying number of data points N

k=1.5	Mean	Standard Deviation	Skewness	SD(Skewness)	Kurtosis	SD(Kurtosis)
N=12	0.895	0.566	1.633	0.406	2.795	1.518
N=60	0.903	0.609	1.107	0.095	1.466	0.370
N=120	0.903	0.612	1.081	0.049	1.407	0.192
N=600	0.903	0.613	1.071	0.010	1.388	0.040
Analytic	0.903	0.613	1.072		1.390	
k=2	**Mean**	**Standard Deviation**	**Skewness**	**SD(Skewness)**	**Kurtosis**	**SD(Kurtosis)**
N=12	0.886	0.427	1.138	0.406	0.864	1.518
N=60	0.886	0.462	0.648	0.095	0.262	0.370
N=120	0.886	0.463	0.635	0.049	0.249	0.192
N=600	0.886	0.463	0.631	0.010	0.245	0.040
Analytic	0.886	0.463	0.631		0.245	

Source: Commerzbank Research

To understand the tables, first look at the analytic results on the bottom row of each table. These are precise values for the moments calculated with the expressions given above. The numerical results are derived from data sets of the moments generated with the number of points N in the table. Thus, for N=12, we are attempting to calculate the mean, sd, skewness and kurtosis using only 12 points, and so on.

Looking one line up from the analytic results, ie, N=600, it can be seen that the numerical results deliver excellent agreement with the analytical - the largest difference is in the kurtosis for k=1.5 (1.388), which differs only slightly in the final figure from the analytic result (1.390). The standard errors derived using the expressions above for the case of N=600 are very small.

However, as N gets smaller, the differences and standard errors become very significant. When N=12, the standard error for skewness is almost as large the skewness value itself, and for the kurtosis in this case, the standard error for the kurtosis is actually larger than the kurtosis value. The numerical results are similarly poor.

In all cases, the results for the mean and standard deviations are close to the true values, indicating that these first two moments are fairly robust to small numbers of points in the fit. When we know that there are at most 8 points in each of the cap and floor data sets, it is clear that the first two moments are likely to be the only reliable ones.

Output from Minneapolis Fed data set

Now that it is clear that sparse data sets can have severe limitations, let us take a look at the Minneappolis data set in detail. To gain an initial impression of the data, we can look at a single day's pdf. Below is the pdf from 9th June 2021, for CPI in 5 years' time - while the actual pdf data is unavailable for download, we have carefully reproduced this day's pdf in graphic form.

Probability density function for inflation
Data for 9th June 2021, US CPI

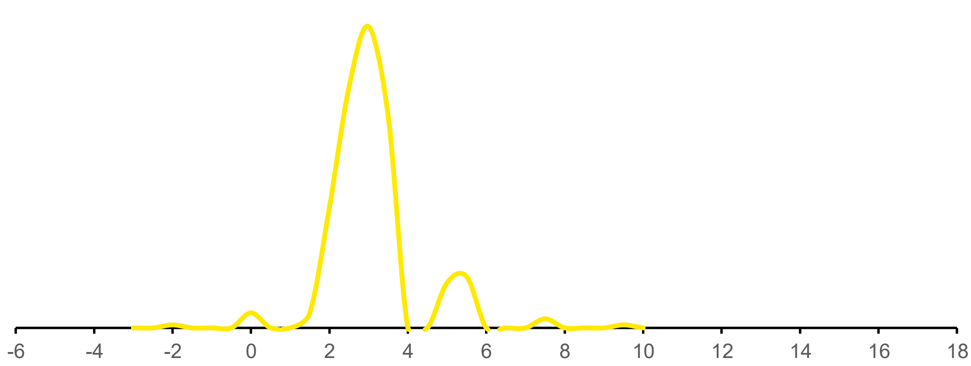

Source: Minneapolis Fed

It is immediately clear that the method they are using to extract information from the market has some problems. This pdf has three major peaks and two minor ones in the tails. No-one would suggest that the actual market implied distribution has these features - on the contrary, it is exactly what one would expect to see from this highly stratified type of method, which, while appealing in its simplicity, needs a densely populated data set to produce realistic results.

Now we take a look at the time series of the moments of the pdf over time, downloaded from the Minneapolis Fed website. The following graphs show mean and standard deviation on the left, with skewness and kurtosis on the right.

Implied Mean and Standard Deviation for inflation probability distribution
Data implied 5y ahead, weekly frequency

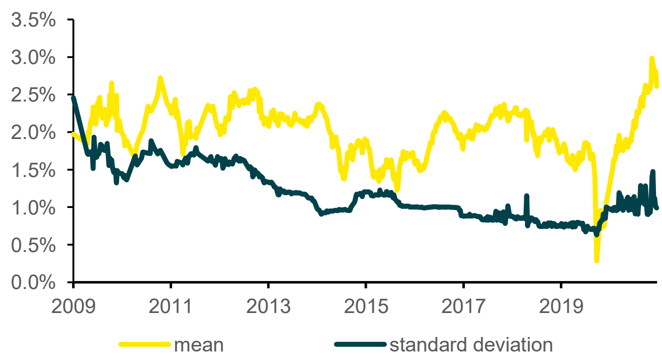

Source: Minneapolis Fed

Implied Skewness and Kurtosis for inflation probability distribution
Data implied 5y ahead

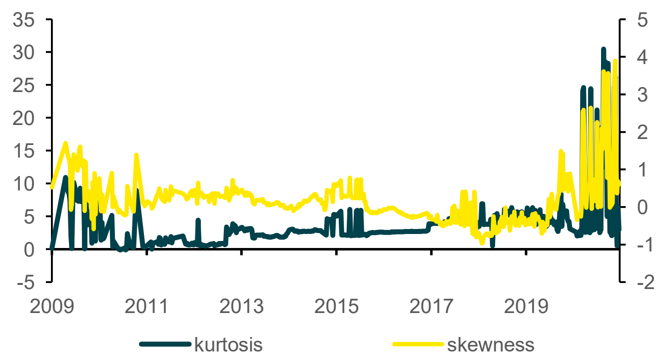

Source: Minneapolis Fed

While the mean and standard deviation look entirely reasonable, with very realistic values and fairly smooth variation, the same cannot be said for the skewness and kurtosis. The right hand graph above has kurtosis on the left hand (primary) axis, and skewness on the right hand (secondary) axis. Particularly in recent times, both moments show every sign of being heavily affected by fitting problems, with kurtosis varying wildly - for example, on 10th March 2021 it had a value of about 5, but the next data point on 17th March 2021 was more than 28. This is entirely consistent with our earlier investigation of the standard errors for these two moments which showed that they grow unmanageably large for sparse data sets.

Output from Commerzank's SABR model

The in-house SABR model can be used to sequentially create pdfs on days in history, and the parameters extracted. First, let's take a look at the pdfs which can be derived from single day's data. We have selected two different sources - HICPx inflation index and option data for the EUR market (right hand graph below), and interest rates for the EUR market (left hand graph below). We have chosen to show the interest rate graph because it is a good demonstration of the flexibility of the model, and because there is a deep and liquid market in options on Euribor, guaranteeing a well populated initial data set.

SABR generated interest rate PDF
3m Euribor, 6m ahead, 8 June 2021

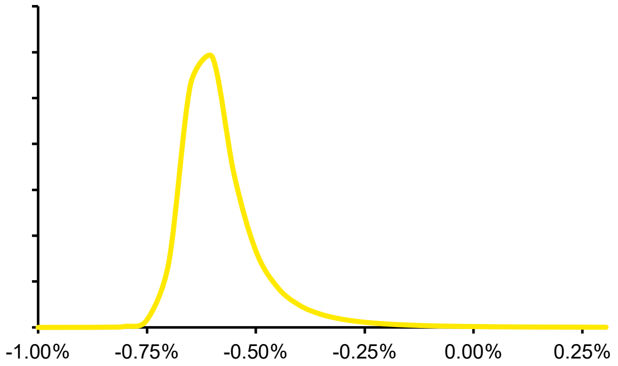

Source: Commerzbank Research

SABR generated inflation PDF
HICPx, 5y ahead, 25 jun 2021

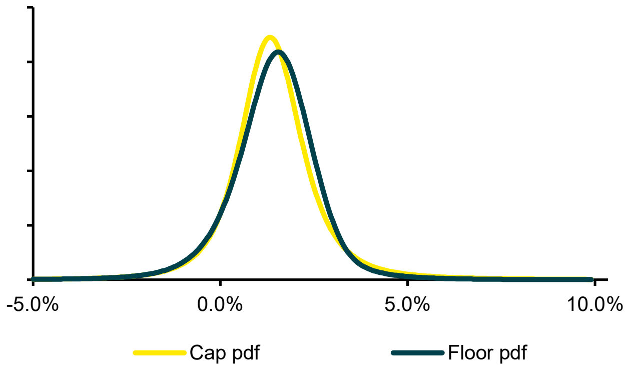

Source: Commerzbank Research

It is clear from the interest rate pdf above that the SABR model is very capable of producing distributions which are skewed, and with tails which vary from thin to fat. The inflation pdfs generated at a similar time are symmetrical - caps and floors are separately fitted as each set is likely to be more relevant to different tails of the distribution. However, we can be confident that if there were, on this date, any measurable degree of skewness or kurtosis in the underlying data, it would be well represented in the resulting PDF.

Having established in interest rate data, that the SABR model, given sufficient density of input data, can produce realistic PDFs, we can take a look at the parameters of this model through time.

Initially, although we do not plot them here, we find that both beta (the tail parameter) and *d* (the overall displacement) are very close to zero. We might expect *d* to be close to zero as it is there to allow negative peak values, but the fact that beta is zero indicates that overall the tails of the distribution are similar to those of a normal distribution. We also find that the sigma parameter of the pdfs derived from both cap and floor data are stable and reasonable after mid-2014, but that until mid 2018 there are not sufficient OTM cap and floor options to give stable results for rho or alpha. After this date, however, we do have stable results for these two parameters.

SABR model correlation parameter rho

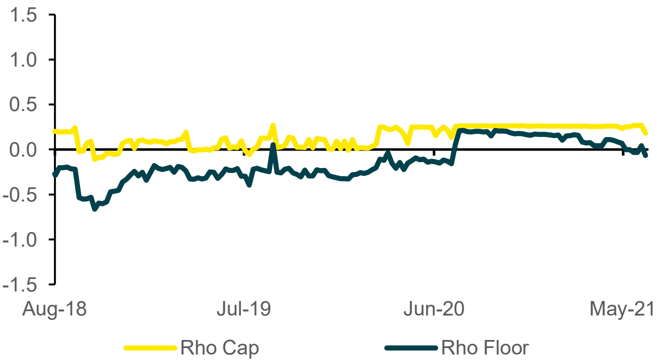

Source: Commerzbank Research

SABR model vol-vol parameter alpha

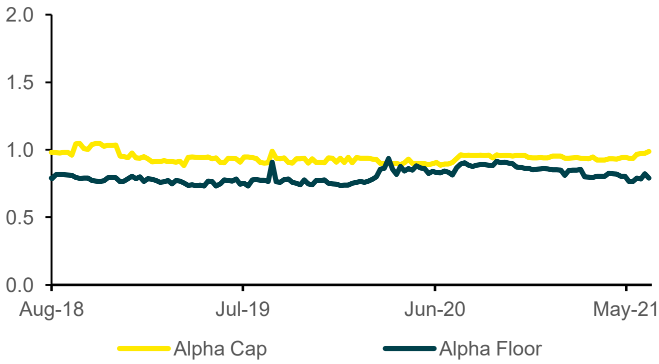

Source: Commerzbank Research

SABR model sigma parameter

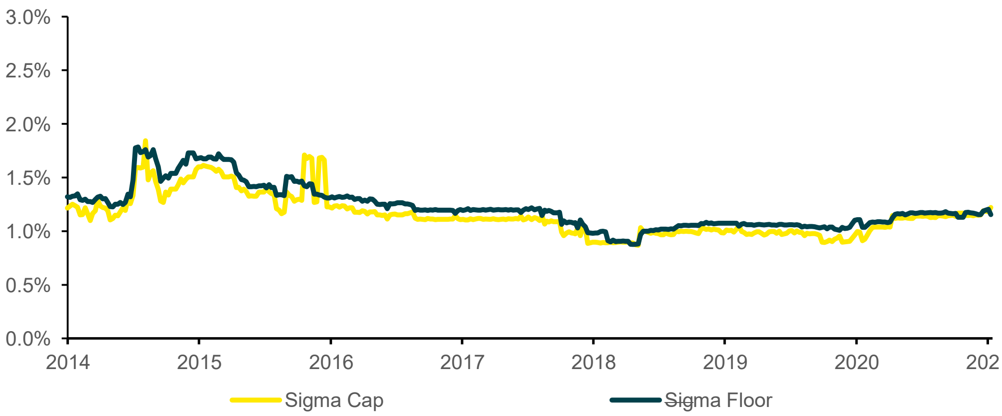

Source: Commerzbank Research

While we would not expect the cap and floor parameters to be identical, it is reassuring as to the quality of the fit that they are similar. What is notable is that rho, the correlation between volatility and rate level, is close to zero, and alpha, the volatility of volatility, is close to 1. Because if rho=0 and alpha=1, the SABR distribution collapses to a normal!

"We have normality. I repeat, we have normality. Anything you still can't cope with is therefore your own problem[9]."

We have shown that the flexible SABR approach yields results for the EUR HICPx implied inflation pdf that suggests it is close to a (well-behaved) normal distribution. Additionally, the moments of the distribution beyond those needed for the normal are poorly fitted for the case of the USA CPI implied inflation pdf.

Having realised this, we can do some further investigations, for both cases.

Usefully, in addition to the distribution moments, the Minneapolis Fed dataset additionally contains the 10%, 50% and 90% probability levels. We know, of course that the 50% probability level is the same as the median of the distribution, so to decide whether skewness is important in this distribution, we can compare the mean and the median. Distributions where these two moments are the same are overwhelmingly likely to have minimal skewness.

To assess the likelihood of significant kurtosis, we compare the 10% and 90% probability levels with those obtained from a normal distribution with the identical mean and standard deviation from the same day in the data set. The results are shown below.

Mean and median
Minneapolis Fed data set

Source: Minneapolis Fed

10% and 90% Probability Levels
Minneapolis Fed & normal distribution

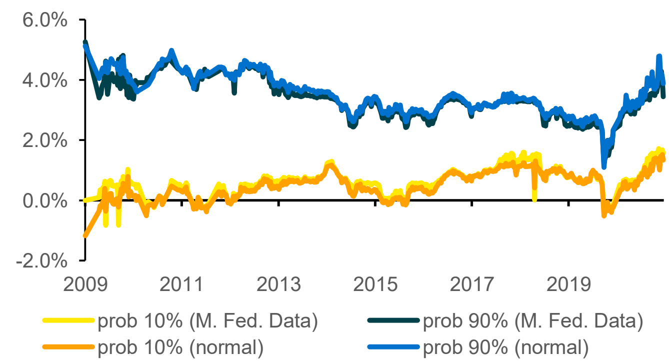

Source: Minneapolis Fed, Commerzbank Research

The two graphs above add up to a convincing case that there is no information in the Minneapolis Fed data set beyond mean and standard deviation, and that the distributions are normally distributed for all intents and purposes. Although it would be extremely useful to be able to compute skewness and kurtosis for the distributions, the sparse input data means that there is no real possibility of obtaining meaningful values for those moments.

Turning to the EUR HICPx distributions, fitted with the more flexible SABR model, the timeseries of its parameters alpha, beta and rho are already consistent with normality. Below, we plot the mean of the distribution through time, together with the SABR 10% and 90% probabilities, and those obtained from a normal distribution with identical mean and standard deviation equal to the SABR sigma.

Comparison of SABR with normal distribution
90% and 10% probabilities from both distributions

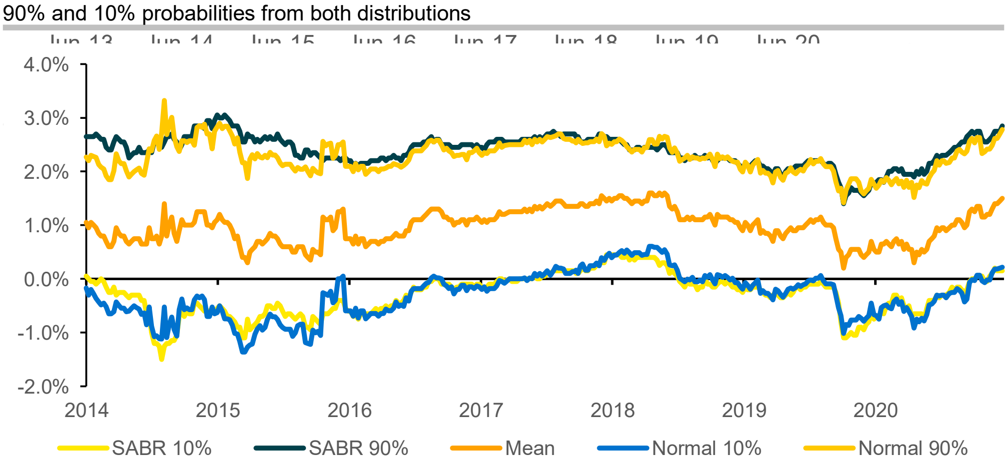

Source: Commerzbank Research

Clearly, the distribution is as close to normal as makes no difference. What we are seeing is that, although the SABR model is perfectly capable of producing output distributions with skew or fat tails, the input data is consistent with a normal distribution, and indeed its sparse nature means that it is highly unlikely that meaningful values of higher moments could be extracted.

Put differently, the quoted caps/floor prices themselves seem to be based on a normal distribution, and irrespective of the sophistication of the model/function applied it seems primarily an inter/extra-polation exercise from a limited set to active points on the grid.

Is inflation normally distributed?

Is it a problem that inflation options appear to assume that inflation is following a normal distribution? This would of course not be the case if actual inflation does follow a normal distribution.

To establish empirically whether this is the case, one needs a very long time series with annual inflation rates.

The euro area HICPx data only goes back to 1995, and it is not possible to make any firm statements on the underlying distribution, although eyeballing a simple histogram suggests a skewed distribution, as we see below in the top left hand graph.

More interesting in that regard is an analysis of inflation for the US, where CPI data goes back to 1914 (top right hand chart), or the UK, where the BoE provides data back to 1209[10] (lower chart)

It is worth a note on the 8 centuries of BoE data - this has been painstakingly compiled from many sources and represents the best estimate of inflation over the years. While one might say that data from before the Wars of the Roses is not necessarily relevant to today, the fact that we now enter uncharted territory with the current inflation environment argues that an examination of data over a huge range of different scenarios has value!

HICPx from 1995 to today
Frequency distribution, %

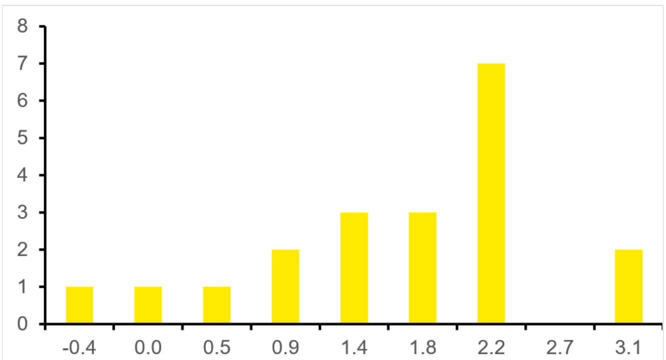

Source: Bloomberg, Commerzbank Research

US CPI data from 1914 to today
Frequency distribution, in %

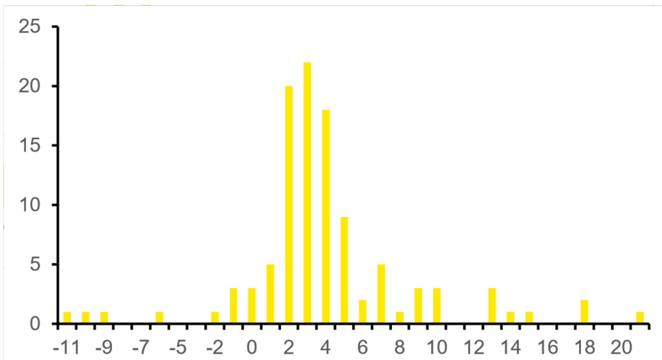

Source: Bloomberg, Commerzbank Research

Bank of England annual inflation data from 1209 to 2016
Frequency distribution, in %

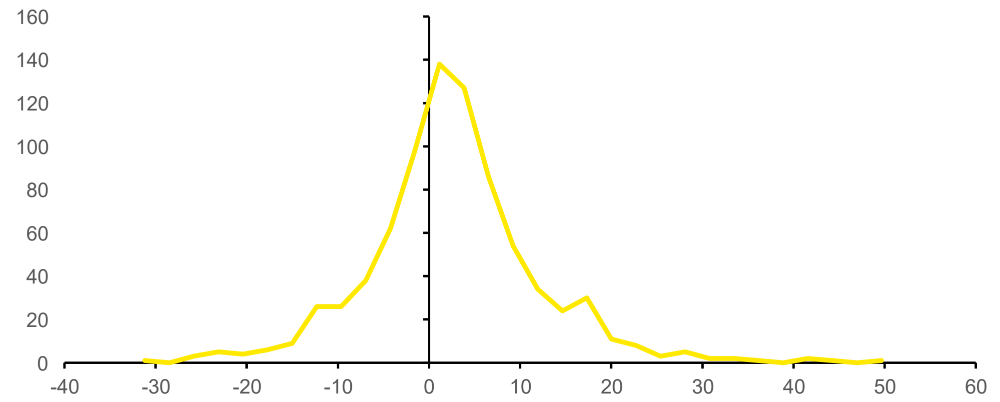

Source: Bank of England dataset, Commerzbank Research

Analysing this data, we find that for the two longer term series at least, there is enough data to find higher moments of the distributions with standard errors low enough to give us some confidence in the results.

	Mean	Standard Deviation	Skewness	SD(Skew ness)	Kurtosis	SD(Kurtosis)
HICPx, N=19	1.50%	0.92%	-0.533	0.274	-0.059	1.029
US CPI, N=106	3.20%	4.83%	0.742	0.055	3.406	0.216
BoE, N=807	1.31%	9.40%	0.503	0.007	2.600	0.030

It is extremely interesting that for both the US CPI and the BoE, their values for skew (slightly positive) and kurtosis (somewhat fat tailed) are similar to each other, and adds weight to the idea that a 'correct' inflation distribution ought to deviate from the normal distribution to an extent. This is backed up by a Working Paper from the ECB[11], who find that US SPF inflation forecasts show strong assymmetry with a positive skew.

Thus, our finding that € HICPx option valuations assume that inflation is following a normal distribution weakens the reliabilty of the implied probability distributions. Considering that in EUR the deflation floor is probably the most active point given its importance for HICPx linkers, some caution is required for higher strikes/the right tail - and in particular as ECB policy makers are increasingly citing changes in precisely the right tail of the option-implied inflation distribution (see for example **here**[12]).

Commerzbank SABR output compared with other data

It is interesting to compare the SABR output with other data - below are the mean, 10% and 90% probability levels for HICPx 5 years ahead from the ECB Survey of Professional Forecasters[13], which comes out each quarter, with a quarterly smoothing of the same SABR quantities.

Survey of Professional Forecasters data, with quarterly smoothed SABR model data, for EUR

Inflation 5y ahead, SPF data left hand axis, SABR data right hand axis

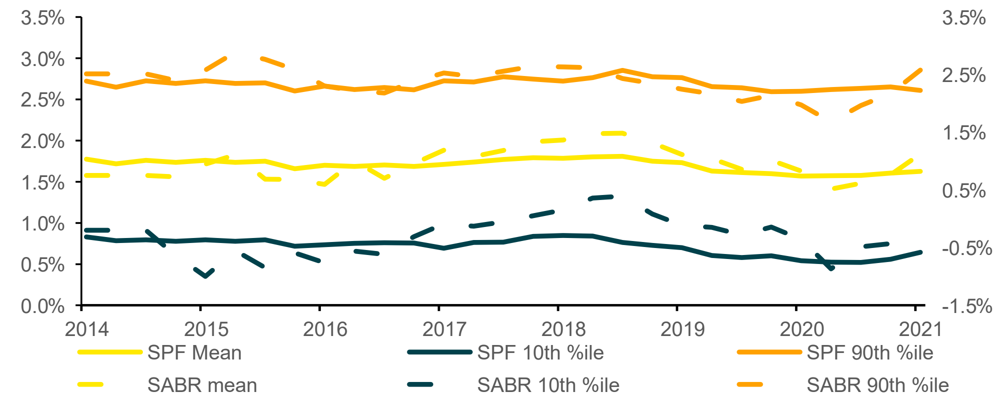

Source: ECB, Commerzbank Research

We see that there is an approximate agreement as to overall level, though the SABR results have a slightly wider range of anticipated outcomes. These differences are not surprising as options comprise risk-neutral probabilities and therefore do not reflect the true distribution given the existence of risk premia. They also refer to HICP ex tobacco, which is typically 7bp lower than headline HICP.

Finally we can do a similar exercise with the US data. This time we take the quarterly CPI forecasts for 5y ahead, from the US Survey of Professional Forecasters, compiled by the Philadelphia Fed. They give the mean, 25th and 75th percentile probabilities, and we perform a similar comparison to that in the previous graph, comparing their figures with a quarterly smoothing of the option-derived Minneapolis Fed data. The SPF data is on the left axis below, with the option-derived data on the right axis.

Survey of Professional Forecasters data, with quarterly smoothed option-derived data, for the USA
Inflation 5y ahead

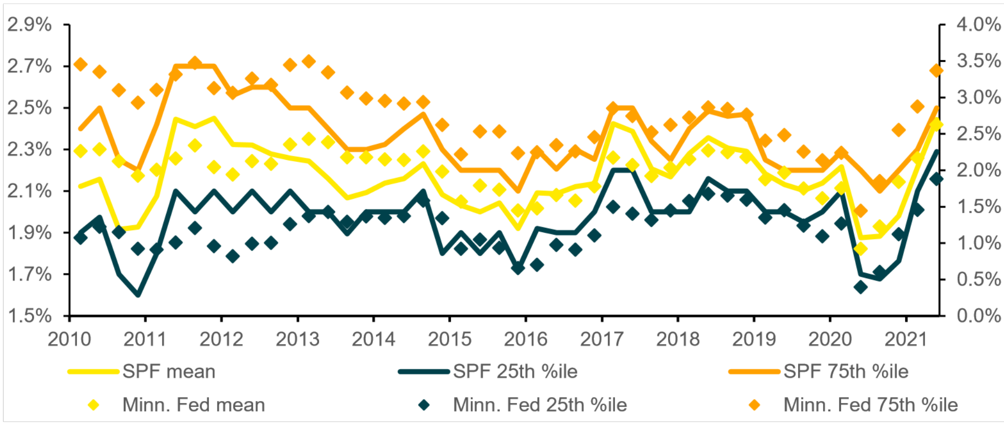

Source: Philadelphia Fed, Minneapolis Fed, Commerzbank Research

As before, there is broad agreement as to the mean of the distributions, and more activity and dispersion for the option-derived data. It is also notable that for both the survey- and the market based USD measures, the ranges are much lower compared to EUR, ie. the degree of inflation uncertainty seems notably lower in the US than in the euro area.

Liquidity issues: normal in - normal out!

It is apparent that the sparse availability of input option data means that there is little chance of calculating valid estimates of higher moments than mean and standard deviation. However, even the data which is available has significant fluctuations in liquidity and can often be very stale. The deflation floor is actively traded, but quotes on strikes away from this level are often mechanically generated levels from brokers, rather than actual trades. This being the case, it begs the question, how do these quotes get calculated? Hazarding a guess, it is unlikely that the brokers do more than use a simple model to generate them - and they probably use a normal distribution. This is absolutely consistent with all our findings above - if the SABR model for example is given data which is clearly skewed or fat tailed, it will represent this in its final output. But, if all it has to work on was originally generated from a normal distribution... well, normal in, normal out!

Conclusion

With the global inflation regime on the cusp of possible change from lowflation to a more dynamic and much higher zone, the markets are minutely focussed on the probability of large moves. However, we have shown that the inflation option market is not liquid or granular enough to extract reliable tail or skew probability parameters beyond the normal distribution. So it would be wise to treat estimates of large moves derived in this way with a good deal of caution.

[1] ECB blog - Inflation dynamics during a pandemic (**here**).
[2] Shiller - The Invention of Inflation-Indexed Bonds in early America (**here**).
[3] SABR stands for Stochastic Alpha, Beta, Rho. See PS Hagan, D Kumar, A Lesniewski, DE Woodward (2002) Managing smile risk, Wilmott, 84-108.
[4] The economics of options-implied inflation probability density functions, *Yuriy Kitsul* and *Jonathan H. Wright* ***Journal of Financial Economics***, 2013, vol. 110, issue 3, 696-711.
[5] Prices of State-contingent Claims Implicit in Option Prices, *Douglas T Breeden* and *Robert H Litzenberger,* ***The Journal of Business***, 1978, vol. 51, issue 4, 621-51.
[6] Bank of England Working Paper (**here**).
[7] https://en.wikipedia.org/wiki/Weibull_distribution is a good general reference for the properties and uses of the Weibull distribution.
[8] Jessica James Nuclear Structure Effects in Atomic Parity Non-Conservation, D.Phil thesis, submitted 1994 at the University of Oxford.
[9] Douglas Adams, 'The Hitch Hiker's Guide to the Galaxy', as the starship Heart of Gold re-enters normal space andtime.
[10] Bank of England long term data set, "A Millennium of Macroeconomic Data", https://www.bankofengland.co.uk/statistics/research-datasets
[11] ECB Working Paper, page 19, (**here**).
[12] Isabel Schnabel, Member of the ECB Executive Board Petersberger Sommerdialog, presentation given 3 July, 2021 (**here**).
[13] ECB Survey of Professional Forecasters (**here**).