Norges Bank

2 | 2026

Staff Memo

Exchange rate probability distributions derived from option prices

23 March 2026

Vetle Øye Opheim
Jonas Tendal
Markets

Keywords
Financial markets, FX, krone exchange rate, options

Staff Memos present reports and documentation written by staff members and affiliates of Norges Bank, the central bank of Norway. Views and conclusions expressed in Staff Memos should not be taken to represent the views of Norges Bank.

© 2026 Norges Bank
The text may be quoted or referred to, provided that due acknowledgement is given to source.

Staff memo inneholder utredninger og dokumentasjon skrevet av Norges Banks ansatte og andre forfattere tilknyttet Norges Bank. Synspunkter og konklusjoner i arbeidene er ikke nødvendigvis representative for Norges Banks.

© 2026 Norges Bank
Det kan siteres fra eller henvises til dette arbeid, gitt at forfatter og Norges Bank oppgis som kilde.

ISSN 1504-2596 (online)
ISBN 978-82-8379-390-1 (online)

Table of Contents

Option prices can contain information regarding the market's assessment of future asset price movements. In this *Memo*, we use option prices to estimate probability distributions of the underlying instrument using the same method as the Atlanta Fed uses in its Market Probability Tracker. We have adapted the method to use FX option prices in order to show how the market prices expectations and risk into exchange rates.

1. Introduction[1][2]

Options are derivative contracts that give the buyer the right, but not the obligation, to buy or sell an underlying asset at a set price at some point in the future. Options are used by both financial and non-financial institutions to hedge against different types of risk or to generate profit based on the investors own expectation of future price movements. An investor expecting NOK appreciation may buy a USD/NOK put (equivalently, a NOK call), giving the right to sell USD at a fixed strike. Another example is an airline company that buys oil options to hedge fuel costs against changes in oil prices.

The value of an option depends on the price underlying asset. An option to buy NOK would be worth more if NOK appreciates. At the same time, there would be uncertainty surrounding future NOK developments. Before reaching maturity, there is always a probability that the option expires worthless at maturity. On the other hand, there is also a certain probability that the option has a positive value at maturity, even if it would not have any value when redeemed today.

The current value of an option is therefore related to the probability that the option has a positive value in the future. Furthermore, as options give only the right, not the obligation, to buy or sell an underlying instrument, all options have a positive value in principle. If an option has an unfavourable strike price at maturity, the owner can choose not to exercise the option. Different options with different strike prices will therefore contain information that can be used to create a probability distribution for the market implied price of the underlying asset.

In this *Memo*, we apply an adaptation of the method used by the Atlanta Fed for its Market Probability Tracker[3]. The Market Probability Tracker is used foremost to show how US policy rate expectations are priced into the options market. Our

---

[1] The Staff Memo has been approved by Gaute Langeland, Director of Markets.

[2] Thanks to Johannes Atle Skjeltorp, Per Gunnar Bull Karlsen, Karsten Gerdrup, Alexander Flatner, Q. Farooq Akram, Sigurd Ugland, Sigvald Hansen and Erik Steinsbu Wasberg for comments and feedback.

[3] See [Market Probability Tracker - Federal Reserve Bank of Atlanta]

adaptation allows the method to be applied to FX options with the aim to uncover market participants' expectations and uncertainty priced into exchange rates through the options market.

Section 2 of this *Memo* gives a more detailed overview of options and how information embedded in option prices can be used to give an indication of market expectations. Section 3 describes the model specifications and estimation. Section 4 shows some results using the method on FX options.

2. Options

There are many different types of options. The two most common types are European and American options. European options can only be exercised at maturity, while American options can be exercised at any time between procurement and maturity. Both European and American options have a set redemption rate. The simplest type of option to value is European options due to the fact that they can only be exercised at maturity.

Variable notation and definitions
S_t: underlying spot price
F_t: underlying forward price
K: strike price
T: time to maturity in years
r_d: risk-free rate in local currency
r_f: risk-free rate in foreign currency
C, P: price of call and put options respectively

A call option gives the right to buy, and a put option gives the right to sell. Options can also be classified based on the relationship between the current price of the underlying instrument and the strike price. A call option is "in the money" (ITM) when $F_t > K$ and a put option is ITM when $F_t < K$. A call option is "out of the money" (OTM) when $F_t < K$ and a put option is OTM when $F_t > K$. If $F_t = K$, an option is said to be "at the money" (ATM).

Options bear similarities with forward contracts. Forwards and options are both agreements to either buy or sell an asset in the future. The main difference is that forward contracts contain an obligation to fulfil the transaction, whereas an option owner can always choose to not exercise any option that does not result in any profit. This means that an option that is exercised will always have a positive value at maturity. However, a forward contract will not usually have any financial value when entered into but can have both a positive and negative value later. At maturity, a forward contract can give the buyer a financial profit or loss depending on the difference between the agreed and market price. By contrast, an option owner only risks losing the option premium paid at the time of purchase.

Chart 1: Forward contract values

As a function of the underlying instrument price

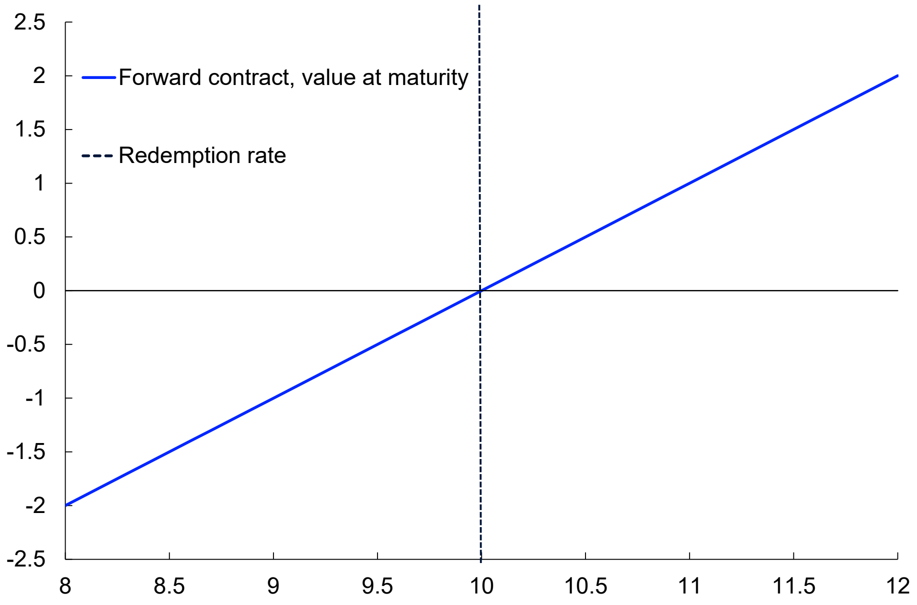

After it is entered into, the value of a forward contract will vary linearly with the price of its underlying asset (Chart 1). This means we observe only *one* forward rate at a given time. If a forward contract was entered into at any other rate than the current forward rate, the value would be positive or negative, calculated by simply comparing it with the market forward rate.

The value of an option will also change with the price of its underlying asset, but the relationship is more complicated. The value of an option can be divided into two components: intrinsic value and time value. The intrinsic value is the profit that is realised if the option is exercised today and the option is ITM. For call and put options, intrinsic value can be expressed using the two following functions:

$$C_I = \max(0, F_t - K)$$

$$P_I = \max(0, K - F_t)$$

Where C_I and P_I are the intrinsic values of call and put options respectively. Chart 2 illustrates the intrinsic value of a call option. The intrinsic value of a call option is equal to zero while the market price of the underlying instrument is lower than the redemption price. When the market price exceeds the redemption price, the intrinsic value will increase linearly with the price of the underlying instrument.

Chart 2: Intrinsic value of a call option

As a function of the underlying instrument price

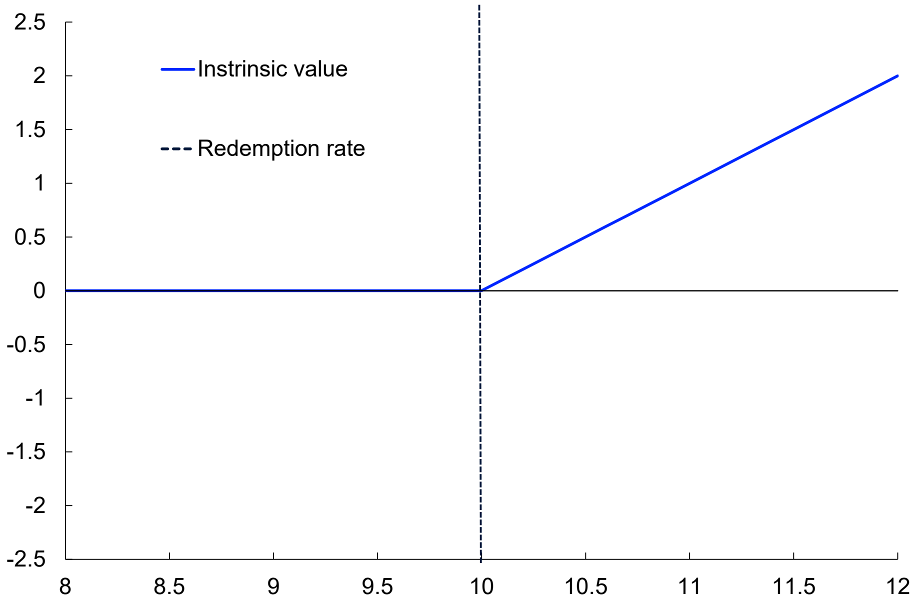

The time value results from favourable changes in the underlying price before maturity. Before maturity, an option will always have a positive time value. At maturity, the time value is zero. The time value is also positive and increases with time until maturity, even if it is equally probable that the price of the underlying asset will increase or decrease. This is due to asymmetry in the function that determines the option's intrinsic value. The option can never have a lower intrinsic value than zero, but there is no equivalent limitation on how high the intrinsic value of an option can be. If the price of the underlying instrument moves against the owner's favour, the option can, in the worst case, have no intrinsic value.

The time value is related to the *probability* that the option will have a positive value at maturity. For example, assume that we buy a call option in USD by selling NOK in 1 year and that the current USD/NOK exchange rate is 9.5, while the option strike price is NOK 10. In other words, the option gives us the right to buy USD at a higher rate than the current exchange rate. This means it is OTM and has no intrinsic value. However, there is a certain probability that the krone exchange rate weakens and is higher than 10 in one year, thus giving the option time value. The longer the maturity, the larger the time value.

Chart 3: Value of a call option

As a function of the underlying instrument price

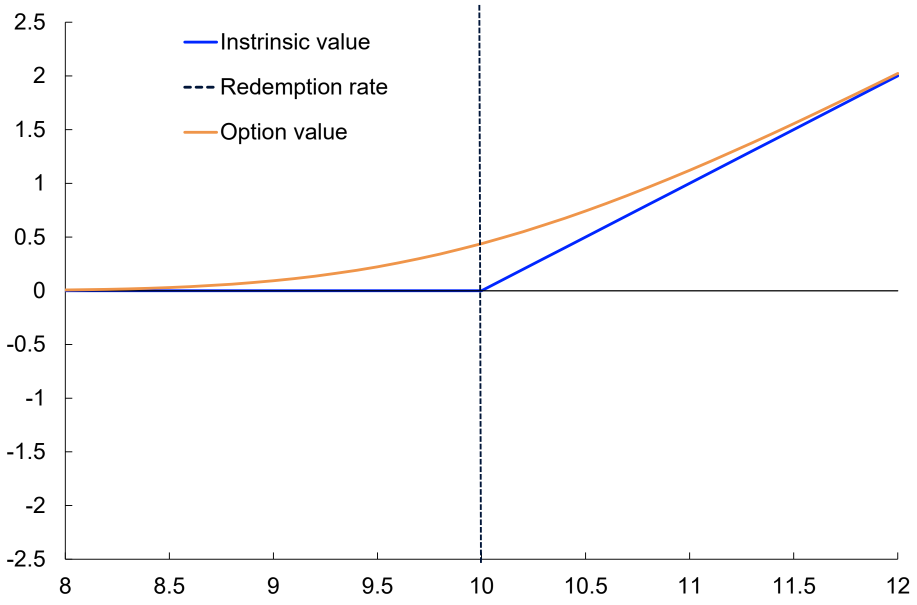

If we look at a currency pair with different strike prices, but the same maturity, each individual option will have a different probability of being ITM at maturity[4]. A higher probability always means that the option has a higher value, which will be reflected in the option's market value. We can take advantage of this to construct an implicit *probability distribution* for the price of the underlying asset using option prices.

3. Model specifications

To derive a probability distribution from option prices, we start with a completely general definition of the value of a European option. The theoretical value of a call option is the current value of the risk-neutral expected value of the option at maturity. A risk-neutral value can contain risk premiums or liquidity premiums, and the actual expectation priced into the market can deviate from calculated values[5].

The risk-neutral expected value is defined as:

$$C_i = D\mathbb{E}^{\mathbb{Q}}[\max(0, S - K_i)]$$

4 This will also be the case if the option is ITM. An option that is currently ITM will also have a certain probability of being OTM at maturity. Alternatively, we can look at the probability that a call or put option that is OTM becomes ITM instead of a put or call option with the same strike price. An OTM call option with a given strike price will initially give the same information as an ITM put option with the same strike price. This is due to the so-called put-call parity for call and put options with the same strike price: $C = P + x - e^{-rT}K$. If put-call parity does not hold, it will be possible to make a self-financing arbitrage strategy with positive gains.

5 Risk-neutral expectations and probabilities can be more simply understood using the following example: Let us assume that you own a house that is worth NOK 10m and you are willing to pay NOK 50 000 each year for fire insurance. The risk-neutral probability that your house burns down within a year will then be $\frac{50\,000}{10\,000\,000} 0.5\%$. However, this may deviate considerably from the actual (much lower) probability of a house fire because most homeowners are risk averse and therefore willing to pay more than the actual probability multiplied by the house value. This willingness to pay above actual calculated expected values is a risk premium. In the same way, option prices can also contain risk premiums because option prices in the market can be affected by investors risk preferences.

Where D is a discount factor and S is the spot price at maturity. If we assume that the underlying instrument has a continuous probability distribution, we get the following expression for the value of a call option:

$$C_i = D \int_K^{\infty} (\mathrm{S} - \mathrm{K_i}) f^{\mathbb{Q}}(S) dS$$

Where $f^{\mathbb{Q}}(S)$ is the risk-neutral distribution of the underlying instrument, and $(S - K_i)$ is the profit on a call option at maturity[6]. When we derive a probability distribution using options, it is the $f^{\mathbb{Q}}(S)$ distribution that we want to find.

In order to find the distribution, we start by assuming that it has a general shape with some parameters that need to be estimated. As in Fisher (2016), we assume that $f^{\mathbb{Q}}(S)$ is a beta-normal mixture distribution. A mixture distribution is the weighted sum of basis distributions, where the weights decide the final shape. Each basis distribution has the same standard deviation, calculated as the average implied volatility of the options used to estimate the distribution. This type of distribution allows for a very flexible form of $f^{\mathbb{Q}}(S)$, and is able to take higher order statistical moments into consideration, such as kurtosis and skewness. In addition, the distribution can be multimodal, ie it can have more than one peak.

The mixture distribution weights are estimated using option prices observed in the market by calculating the theoretical value of each option for each basis distribution. In other words, we derive the options' expected values for all basis distributions as if each basis distribution was the entire probability distribution for the underlying instrument. For example, if we had a dataset of 20 options, we would derive 20 potential option values for each basis distribution.

For each basis distribution, the 20 derived option values need to be compared with observed market prices. If the derived values correlate well with market prices, the distribution will be assigned a higher weight. Values that deviate significantly from market prices result in a lower weight assigned to the corresponding distribution. In practice, weights are assigned using a Bayesian regression, described in more detail in Appendix 1.

Chart 4 shows an example of how a beta-normal mixture distribution would look if all underlying basis distributions were weighted equally. The distribution parameters are set to k=10, $\sigma = 1$ and $\mu = 0{,}5$. Each basis distribution is scaled with a weight equal to $\frac{1}{k}$ in this example. When all weights are equal, the mixture distribution is a normal distribution with standard deviation σ and expected value μ. When the distribution is used for options, σ is calculated as the average implied volatility of the options used in the estimation.

---

6 The profit function of a call option is given by $\max(0, S - K_i)$, but because we only integrate over the values for S, which are higher than K_i, we can write the profit function as $(\mathrm{S} - \mathrm{K_i})$. The function does not take into consideration that the buyer needs to pay a premium, thus the overall profit or loss for the buyer at maturity is lower.

Chart 4: Example of a beta-normal distribution with equal weights

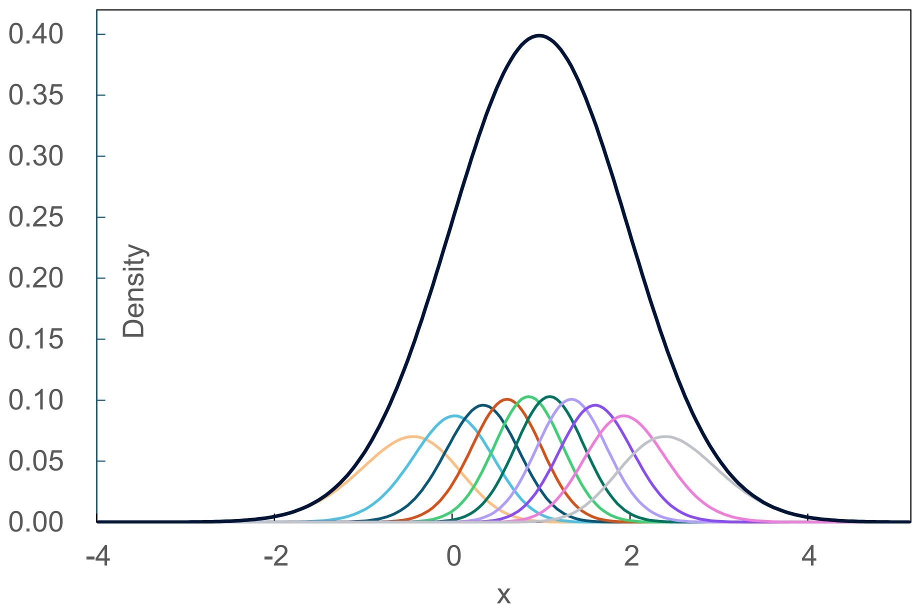

However, the weights will almost always be unequal such that the aggregated distribution deviates from a normal distribution. Chart 5 shows how the distribution looks with randomly chosen weights. This chart illustrates how the mixture distribution can have multiple distinct peaks or modal values. The distribution can also have heavy tails (kurtosis) and be skewed, depending on whether the distributions far from μ are assigned high weights.

Chart 5: Example of a beta-normal distribution with randomly chosen weights

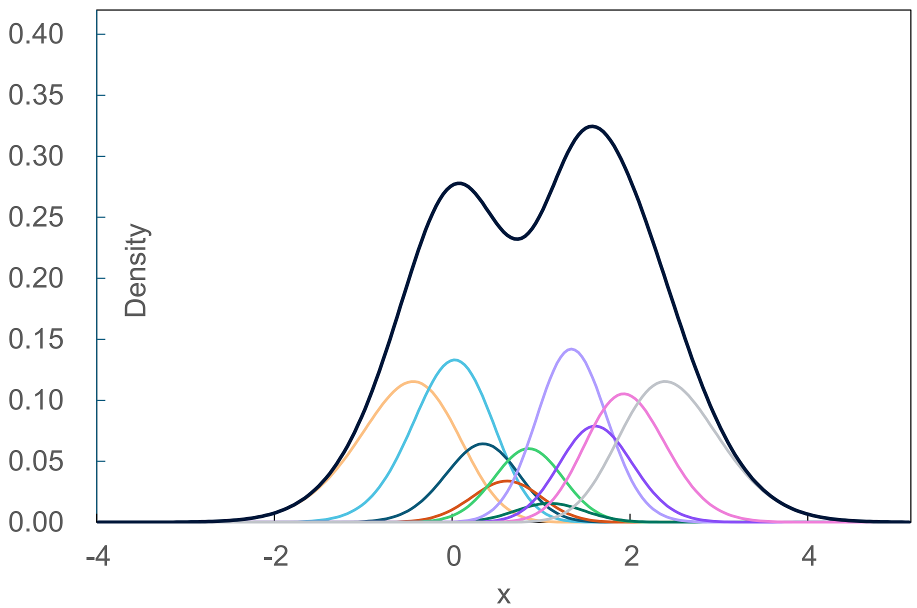

3.1 Data

The data used in the estimation are time series of implied volatility between 2007 and 2025 produced by London Stock Exchange Group (LSEG), which is initially given in terms of delta values instead of redemption prices[7]. Implied volatility and delta values therefore need to be converted to option prices and redemption prices. The delta value is a measure of the sensitivity of the market price of an option to movements in the underlying asset price, ie if a call option has a delta value of 50% and the price of the underlying instrument increases by NOK 1, the option price will increase by NOK 0.5.

We use an adapted version of the Garman and Kohlhagen (1983) formula to convert delta values to redemption rates[8], resulting in the following formula for the redemption rate K:

$$\ln(K) = \ln(F_t) + T\frac{\sigma^2}{2} - d_1\sigma\sqrt{T}$$

$$d_1 = N^{-1}(\Delta_C)$$

$$F_t = S_t e^{(r_d - r_f)T}$$

F_t is the forward rate of the currency pair, which can be calculated from the spot rate S_t. r_d and r_f are the risk-free rate in local currency and foreign currency respectively, Δ_C is the call option's delta value and N^{-1} is the inverse function of a standard normal distribution.

As we have values for implied volatility, when the redemption rate is derived from delta values, the option price can be calculated using the Garman-Kohlhagen formula. The dataset has 18 different time series for different delta values, from 10 to 50 delta for both call and put options. This means that for each date, we can construct 20 synthetic options that are either OTM or ATM.[9] ITM option prices can be calculated by using the put-call parity:

$$C - P = e^{-r_d T}(F_t - K)$$

This results in 18 pairs of call and put options, where each option pair has the same redemption rate.

4. FX option results

In theory, the method described above can be used for options with many different types of underlying financial instruments. We use it on FX options to show how probability distributions can be estimated from market prices. First, we present the results for USD/NOK and EUR/NOK, then USD/EUR and finally EUR/CHF.

In contrast to other types of option, all FX options are both a call and put option. For example, a call option to buy USD against NOK is simultaneously a put option to sell NOK against USD. When discussing exchange rates, the first

---

[7] FX option prices are often given in terms of delta values rather than redemption rates. Prices given in terms of delta values simplify the comparison of implied volatility on different dates and between different exchange rates. The delta value is a measure of how fast the price of an option moves when the price of the underlying instrument changes. It is formally defined as the derivative of the option price with respect to the price of the underlying asset. Delta values are between 0 and 100, where a low or high delta value indicates that the price of the underlying instrument is far from the option's redemption rate. A high delta value indicates that the option is deep ITM. ATM options will have a delta value close to 50. The volatility smile is a plot of implied volatility against different delta values. It is usual that ATM options have lower implied volatility than OTM options, leading to the plot resembling a smile.

[8] The Garman-Kohlhagen formula is a modified version of the Black-Scholes formula adapted to FX options.

[9] ATM options will have a delta value close to 50.

three letters of a currency pair denote the base currency, while the last three letters denote the quoted currency, eg USD/NOK is the amount of NOK needed to buy USD 1, while EUR/USD is the amount of USD needed to buy EUR 1.

We use options with three or 12 months until maturity, ie the implied expectations three or 12 months ahead. For example, a three-month distribution calculated 3 January 2025 will describe exchange rate expectations for 3 April 2025.

4.1 NOK options

To calculate NOK expectations based on the options market, we use currency pairs USD/NOK and EUR/NOK. Charts 6 and 7 show distributions for these two currency pairs for a three-month and one-year horizon as of 13 October 2025.

Chart 6: USD/NOK distributions

As of 13 October 2025.

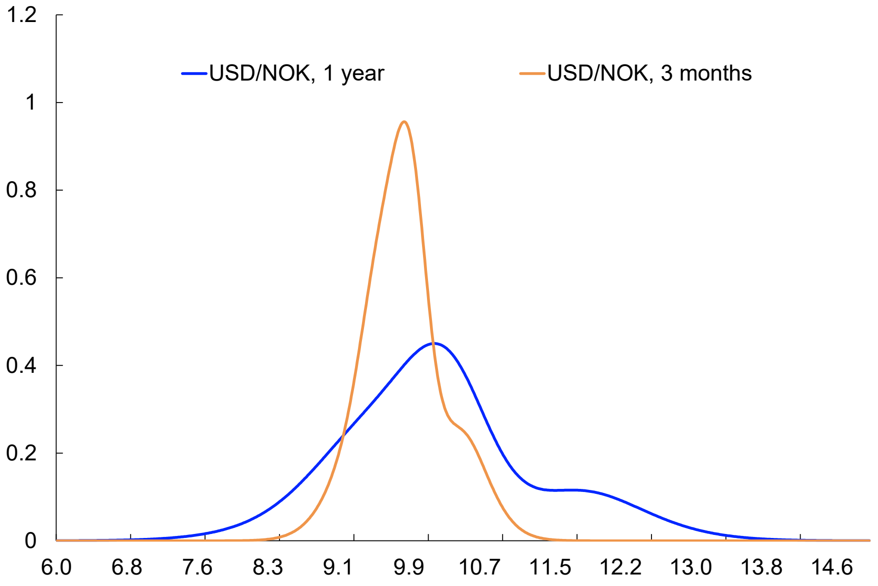

Source: LSEG and Norges Bank

Distributions based on NOK options against large international currencies like USD and EUR will usually have a heavy right-side tail, which is clearly visible in both charts. This means that the options market has priced in a higher likelihood of a significant NOK depreciation than a correspondingly significant NOK appreciation. The standard deviation is also higher for USD/NOK than for EUR/NOK, which likely reflects expectations that exchange rate movements will be more volatile for the USD currency pair than the EUR currency pair.

Chart 7: EUR/NOK distributions

As of 13 October 2025.

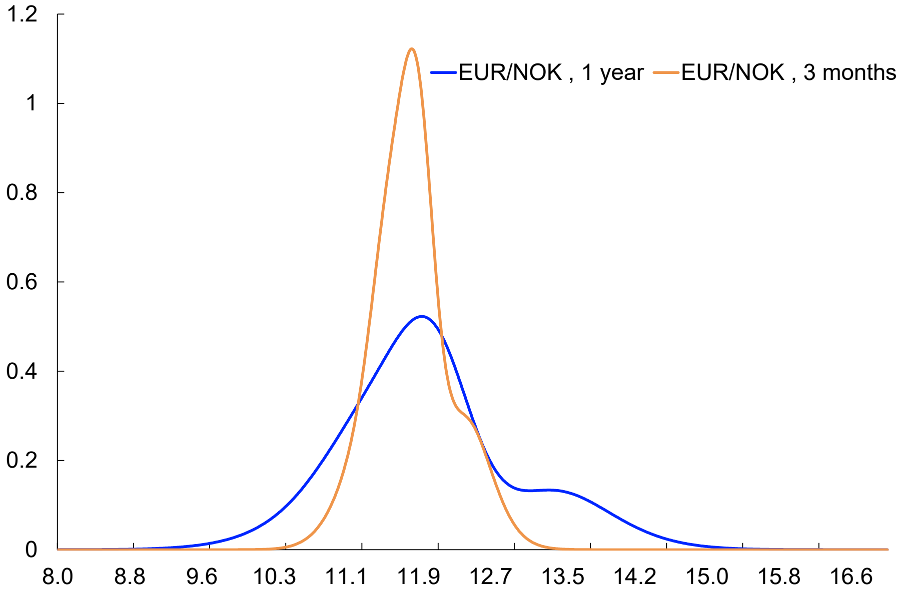

Source: Norges Bank

The shapes of the distributions are similar, but the one-year distributions are wider. The standard deviation increases with time horizon by the same mechanism that causes the exchange rate to deviate more from the current forward rate with a longer time horizon. The right-side tail of the krone exchange rate distribution also tends to be heavier for a one-year horizon compared with a three-month horizon, indicating that the pricing in of risk of a significant NOK depreciation is higher on a one-year horizon than a three-month horizon.

In many cases, it could be more informative to express the distributions in terms of percentage change from the current forward rate than in terms of exchange rate level. The exchange rate level can fluctuate significantly over time. A transformation to percentage change from the current forward rate is the kind of normalisation that simplifies the comparison of distributions for different dates[10]. Skewness in the distribution is also easier to identify, eg by comparing the calculated probability of a 5% depreciation with a 5% appreciation.

Chart 8 shows examples of two such distributions for USD/NOK as of 3 July 2007 and 2 July 2008. NOK appreciated in this period from around 5.8 to 5.1. On the earlier date, the USD/NOK pricing was fairly symmetrical; calculated probability of a 10% NOK depreciation against USD was approximately equal to the probability of a 10% appreciation. In addition, the probability of a 10% movement in either direction was very low. One year later, market conditions had changed considerably. The implied volatility priced into the krone exchange rate was higher, resulting in a larger range of possible outcomes and

10 The following formula can be used to transform a probability distribution from one variable to another: $f_R(R) = f_P\big(P_0(1+R)\big).\frac{d}{dR}\big(P_0(1+R)\big) = f_P(P_0(1+R))P_0$.

uncertainty was priced in more asymmetrically. The priced-in probability of a 10% depreciation during the next three months was larger than the corresponding probability of an appreciation.

Chart 8: USD/NOK distributions

As of 3 July 2007 and 2 July 2008. Percentage change from same day forward rate. Three-month horizon

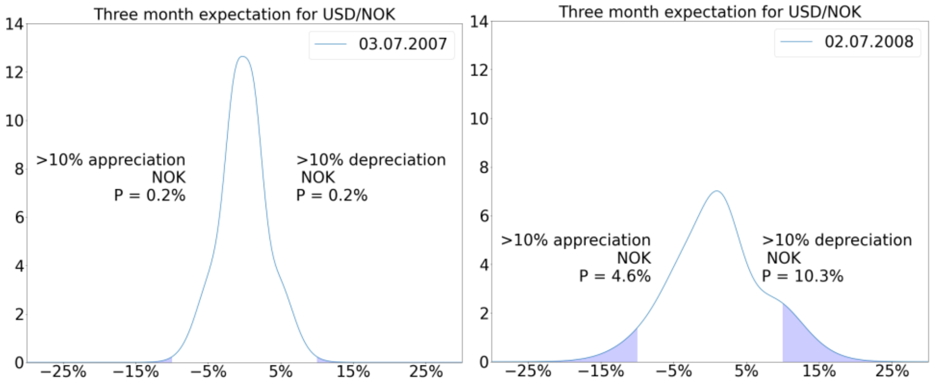

Source: LSEG and Norges Bank

The distributions shown in terms of percentage change can also be used to make a standard distribution over a longer time horizon by taking the average of a range of distributions over a given time interval. By using such distributions, we can investigate how a "typical" distribution will look for a given time period. Chart 9 shows two such distributions for the periods 2007-2025 and 2020-2025.

Chart 9: USD/NOK distributions

Average USD/NOK distributions calculated for two different periods Percentage change from same day forward rate. Three-month horizon

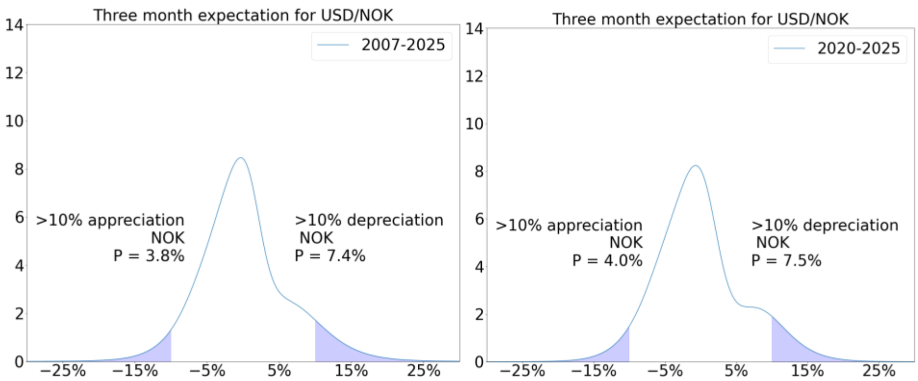

Source: LSEG and Norges Bank

The average distribution for the entire period shows a tendency for larger uncertainty of a significant NOK depreciation against USD to be priced in relative to a corresponding appreciation. The chart showing the average distribution since 2020 resembles the distribution for the entire period. Therefore, the risk priced into the krone exchange rate the past four to five years has largely been the same as the average risk priced in historically.

4.2 EUR/USD options

Asymmetry in an options market-based distribution of the EUR/USD exchange rate may show how risk between EUR and USD is priced. Chart 10 shows two distributions for the EUR/USD exchange rate with a three-month horizon, calculated as of 3 March and 26 May 2025. In March, the distribution had a marked left-side tail, suggesting that the market priced in elevated risk of a larger EUR depreciation against USD.

Three months later, the situation changed. USD depreciated against EUR, causing a shift in the entire distribution and a heavy right-side tail, suggesting a shift in depreciation risk priced into the options market during this period from EUR to USD.

Chart 10: EUR/NOK distributions

As of 3 March and 26 May 2025. Three-month horizon

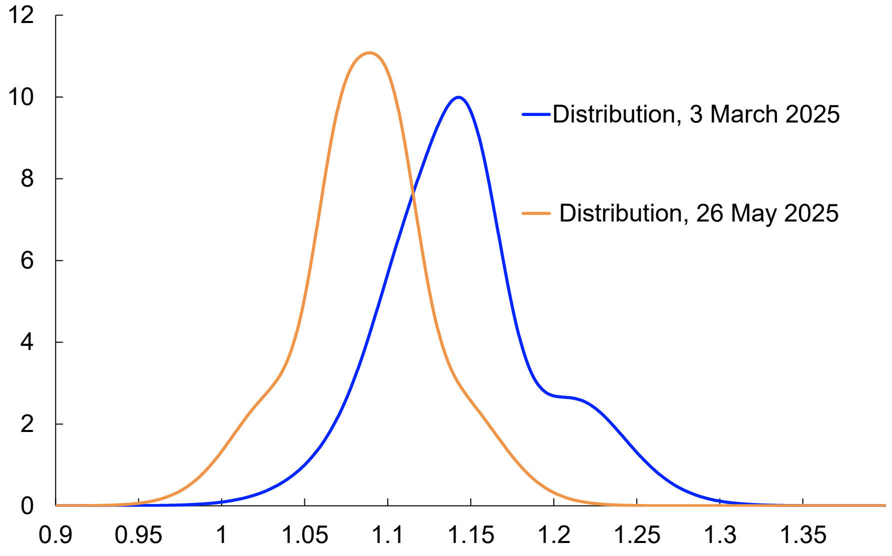

Source: LSEG and Norges Bank

The change occurred primarily during April 2025; a period of heightened financial market volatility owing to new US tariffs on imported goods. The US dollar usually appreciates during periods of financial market stress, which causes the relative value of options realising a profit following USD appreciation to increase. In this case, we see the opposite. Options that hedge against USD depreciation became relatively more expensive, reflecting the heavy right-side tail on the May distribution.

4.3 CHF options in 2015

The examples covered so far have been based on currencies using floating exchange rate regimes. The method covered in this Memo can also be used to investigate how the options market prices risk of collapse in fixed (pegged) exchange rate regimes. More specifically, we can investigate whether the market prices in the collapse of a pegged exchange rate regime or collapse of an exchange rate ceiling or floor. An example is the exchange rate floor between CHF and EUR from September 2011 to January 2015.

During the first half of 2011, CHF appreciated considerably against EUR due to the Euro area debt crisis (Chart 11). On 6 September 2011, the Swiss central bank announced market interventions to keep the EUR/CHF exchange rate over 1.20. The exchange rate quickly depreciated following this, and the floor was kept in place until 15 January 2015. The floor was then abandoned and CHF appreciated immediately by over 20% against EUR.

Chart 11: EUR/CHF exchange rate

Euros per Swiss Franc. March 2011-October 2015

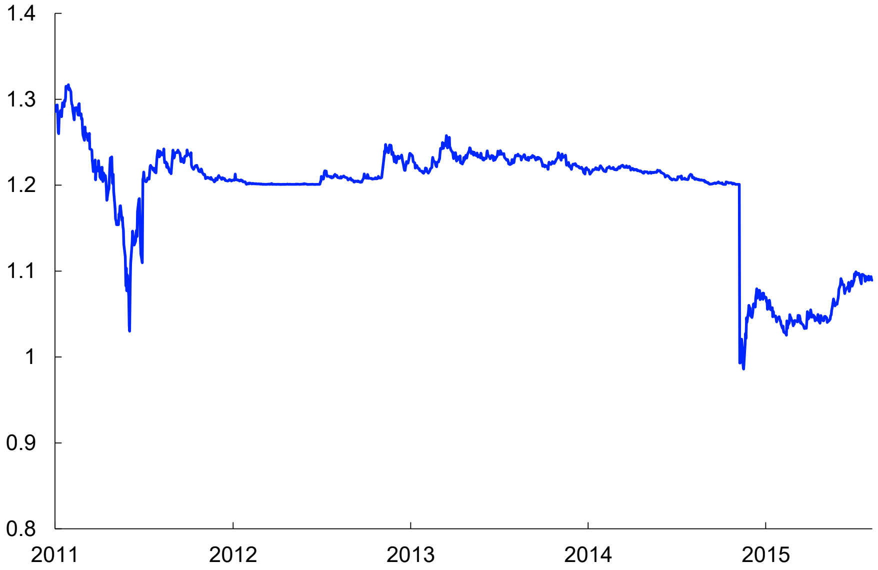

Source: Norges Bank[11]

Chart 12 shows the options market-based EUR/CHF distributions with a three-month horizon as of 8 and 15 January 2015. The distribution as of 8 January shows higher probability for a rate of 1.20 or higher, however, some of the left-side tail is under 1.20. Based on this distribution, the probability of a break in the exchange rate floor within the next three months can be estimated to be around 40%. Breaks in the floor that were priced in were, however, small.

The distributions show that the removal of the exchange rate floor not only led to a shift in the distribution, but exchange rate volatility also increased considerably. The standard deviations were 2.3% and 9.1% for the 8 January 2015 and 15 January 2015 distributions respectively. The implied volatility of the exchange rate then normalised, and was typically between 5% and 6% in 2019.

[11] The exchange rate is calculated using the EUR/NOK and CHF/NOK exchange rates.

Chart 12: CHF/EUR distributions

Exchange rate (x-axis) and probability density (y-axis). As of 8 and 15 January 2015

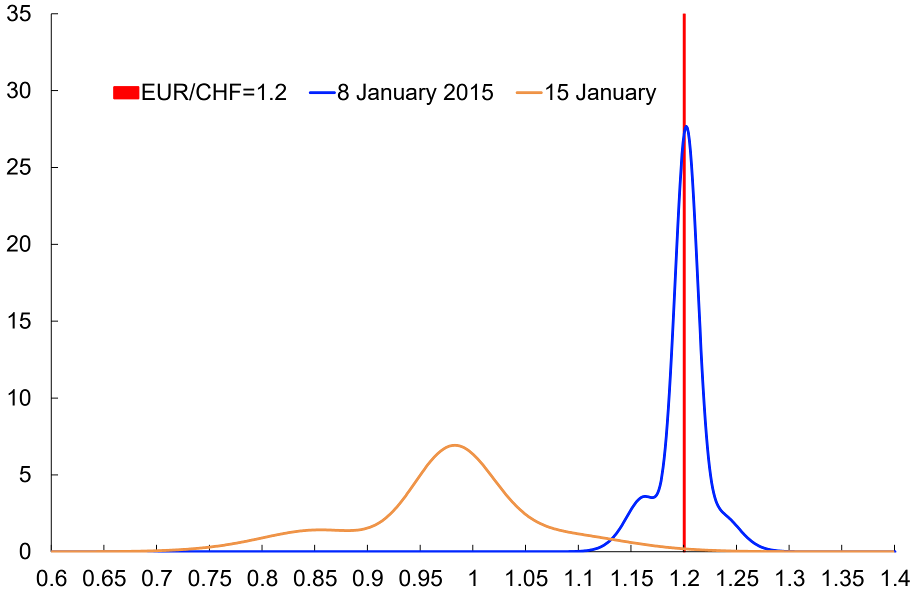

Source: LSEG and Norges Bank

This may indicate that the market priced in a certain probability of floor breaks, but that the exchange rate movements in such an event would be smaller than what actually occured . The exchange rate floor removal led to a considerable increase in uncertainty surrounding the exchange rate level that remained for a long period. Even though the likelihood of a complete abandonment of the exchange rate floor may have been considered low, options with redemption rates above the floor would always have a positive value, precisely because such a scenario can never be ruled out.

4.4 Expectations or risk premiums?

In this *Memo*, "expectations" refers to market expectations based on a risk-neutral expectation. This can deviate from market participants' "actual" expectations. The conversion from market pricing to a probability distribution builds on the assumption that the willingness to pay for an option is related to the probability that an option has value at maturity, ie observed prices are assumed to be set by market participants that buy or sell options based on assessments of the likelihood of various exchange rate movements ahead.

For many market participants, option trades are fuelled by a need or a preference to hedge against risk. In other words, option buyers may be risk averse: an airline company may buy call options based on oil contracts to hedge against higher fuel costs, a portfolio manager may buy put options based on an equity index to limit losses due to market collapse, and a manufacturing company may buy FX options to hedge foreign trade against exchange rate fluctuations.

The motive to buy options for insurance purposes may be an important part of options pricing but it does not give a complete picture. In practice, option prices are set both by market participants that want insurance and those with future market expectations that can be expressed through option contracts. It is difficult to differentiate between risk premiums and market expectations. The distributions calculated in this *Memo* must therefore be assumed to potentially contain both risk premiums and future exchange rate expectations.

As with other types of insurance, options require the buyer to pay a premium, with the promise that they will receive a larger sum under certain conditions, such as a movement in the underlying asset's market price in the option's favour. On the other side, option sellers function as sellers of insurance. Option sellers receive option premiums from the buyer and take on the risk that the option has a positive value at maturity and that the option is then exercised. In contrast to an insurance seller, the options seller can hedge the position by taking an offsetting position in the underlying instrument (delta hedging). A call option seller can buy the underlying instrument in an amount that is proportional to the exposure incurred by the option. This purchase will have a financing cost, and this cost can in turn affect the price of the option.

Bank of England (2004) argues that the expected value of risk-neutral distributions can often deviate from the expected value of the "true" distribution. The expected value of the risk-neutral distribution will be very close to the forward price of the underlying instrument with the same maturity. This is due to an assumption that arbitrage between the forward market and the options market is not possible[12] but also restricts where the distribution will be centred. For other statistical moments, such as variance, skewness and kurtosis, there may be a higher correspondence between the risk-neutral distribution and the "true" distribution. Rompolis and Tzavalis (2010) shows that the risk premium in the risk neutral distribution for stock market indices, can largely be explained by skewness in the distribution. For stock indices there is normally a negative skewness, indicating a higher price on insurance against large stock market declines. For the Norwegian Krone, there is normally a positive skewness in the distribution, indicating a higher price on insurance against a weakening Krone. The positive skewness here can be interpreted as a risk premium on owning Norwegian Kroner.

5. Summary

Option prices can contain information regarding the market's pricing of the probability of price movements of an option's underlying instrument. Using options with many redemption rates, it is also possible to estimate a probability distribution for the price of the underlying asset.

In this *Memo*, our calculations are based on the method used by the Atlanta Federal Reserve Bank in its Market Probability Tracker, adapted to be used with FX options. Based on the estimations in this *Memo*, it appears that a larger

[12] if the expected value of the options market deviates from forward market prices, arbitrage becomes possible as combinations of option contracts can be used to make synthetic forward contracts. Assume that the expected value of USD/NOK options gives an exchange rate in one year of 10, while the one-year forward rate is 11. Assume also that the one-year rate in Norway and the US is equal. It will then be profitable to sell USD in the forward market in one year at a forward rate of 11. To secure this position, we can buy a call option and sell a put option with a redemption rate of 10. Such a combination of options resembles a forward contract in which we are obliged to buy USD against NOK in one year at an exchange rate of 10. If the options market prices in an expectation value equal to 10, the two option contracts have the same price, meaning that the transaction entails no expense. Since we have entered into a forward contract in the forward market to sell USD against NOK at a rate of 11 in one year, we are left with a risk-free profit of NOK 1 regardless of market movements.

probability of a considerable NOK depreciation against USD and EUR is usually priced into the market compared with the probability of a corresponding appreciation.

Options market-based probability distributions also have some limitations. Options are priced to ensure that there is no considerable possibility of arbitrage between the options market and forward market so the expected value of the distributions will usually be close to the market forward rate with the same maturity. As a rule, we must also assume that the distributions contain both market expectations and risk premiums as it is currently not easy to distinguish between them.

Sources

Bahra, B. (1997). *Implied risk-neutral probability density functions from option prices: theory and application.* Bank of England working papers.

Bank of England. (2004). *Notes on the Bank of England Option Implied Probability Density Functions.* Bank of England.

Blake, A., & Rule, G. (2015). *Deriving option-implied probability densities for foreign exchange markets.* Bank of England.

Breeden, D., & Litsenberger, R. (1978). Prices of State-Contingent Claims Implicit in Option Prices. *Journal of Business, Vol. 51, No. 4*.

Figlewski, S. (2018). Risk-Neutral Densities: A Review. *Annual Review of Financial Economics*, 329-359.

Fisher, M. (2016). *Simplex Regression*. Retrieved from https://www.atlantafed.org/-/media/documents/cenfis/market-probability-tracker/market-probability-tracker_simplex-regression.pdf

Garman, M. B., & Kohlhagen, S. W. (1983). Foreign currency option values. *Journal of International Money and Finance, Volume 2, Issue 3*, 231-237.

Gelman, A., & Rubin, D. B. (1992). Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 457-511, Vol 7 No 4.

Rompolis, L. S., & Tzavalis, E. (2010). Risk Premium Effects on Implied Volatility Regressions. *The Journal of Financial Research*, Volume 33 Issue 2 125-151.

Savitzky, A., & Golay, M. J. (1964). Smoothing and Differentiation of Data by Simplied Least Squares Procedures. *Analytical Chemistry*, 1627-1639, Vol 36 Issue 8.

Appendix 1: Estimation

The risk-neutral distribution $f^{\mathbb{Q}}(x)$ is assumed to be a beta-normal mixture distribution, ie that $f^{\mathbb{Q}}(x)$ is a weighted sum of distributions with the following form:

$$f^{\mathbb{Q}}(S) = \sum_{j=1}^{k} \beta_j g_j(S)$$

$$g_j(S) = Beta\left(N\left(\frac{x-\mu}{\sigma}\right)\middle| j, K-j+1\right) N'(S|\mu, \sigma^2)$$

Where $N'()$ and $N()$ are standard density and cumulative normal distributions respectively. β_j represents the weights for each of the distributions. This means that we need two restrictions on the value of the weights: $0 \leq \beta_j \leq 1$ and $\sum_{j=1}^{k} \beta_j = 1$. This ensures that all weights have a value between 0 and 1, and that all weights sum to 1.

For a given call option, we can write the option value as:

$$C_i = D\int_{K_i}^{\infty} (S - K_i)\sum_{j=1}^{k} \beta_j g_j(S)\, dS$$

$$C_i = D\sum_{j=1}^{k} \beta_j \int_{K_i}^{\infty} (S - K_i) g_j(S) dS$$

In the next stage of the estimation, we need to calculate all theoretical option values for each individual basis distribution. Let $\boldsymbol{X^C}$ be a matrix where we calculate an expected value for each call option in the dataset under each basisdistribution j, ie that each value i $\boldsymbol{X^C}$ will be the expected payout resulting from an option if we assume each individual basis distribution is the probability distribution of the underlying instrument. This results in each element *i,j* of the matrix being defined as the expected value for option *i* for the probability distribution $g_j(S)$:

$$X_{ij}^{C} = \int_{K_i}^{\infty} (S - K_i) g_j(S) dx$$

Which means we can write the value of a call option as:

$$C_i = D\sum_{j=1}^{k} \beta_j X_{ij}$$

For put options, we calculate a matrix $\boldsymbol{X^P}$ where each element is defined as:

$$X_{ij}^{P} = \int_{-\infty}^{K_i} (K_i - \mathrm{S}) g_j(S) dx$$

Finally, all option values can be combined as a single vector $\boldsymbol{y} = [\boldsymbol{C}', \boldsymbol{P}']$, and all calculated option values are combined in a matrix $\boldsymbol{X} = [\boldsymbol{X^C}, \boldsymbol{X^P}]$. To find the weights β_j in the beta-normal distribution, we estimate the regression:

$$y_i = D\sum_{j=1}^{k} \beta_j X_{ij} + \varepsilon_t$$

Where y_i is an option price for either a call or put option. X_{ij} is calculated numerically for a given value of μ and σ. μ is set as the observed forward price of the underlying instrument. σ can be set in different ways based on observed implied volatility. Here we use the average implied volatility for options used in the estimation.

As in Fisher (2016), we use a Bayesian approach to estimate the weights[13]. Let $\boldsymbol{\beta} = (\beta_1, .., \beta_k)$, and assume that $\boldsymbol{\beta}$ has a Dirichlet distribution prior:

$$\boldsymbol{\beta} \sim Dir\left(\frac{\alpha}{k}\right)$$

The Dirichlet distribution ensures that the weights have values between 0 and 1, and that they sum to 1.

α is a concentration parameter. α decided where concentrated weights in the normal distribution will be. If α is large, this means that the spread of large β weights will be narrow. α is also a random variable with a log-normal prior distribution:

$$\alpha \sim LogNorm(\gamma, \eta^2) \qquad \text{where } \gamma, \eta^2 = (1,1)$$

Furthermore, we define a prior distribution for the discount factor. The discount factor can have a distribution that does not allow for negative values. In addition, the discount factor usually has a value near 1. We choose to use a truncated normal distribution centred around 1:

$$D \sim TruncatedNorm(\lambda, \tau^2, min = 0) \text{ where } \lambda, \tau^2 = (1,1)$$

Finally, we define a prior distribution for the standard deviation of the remainder ε_t, $\sigma_\varepsilon \sim HalfNorm(u^2)$. σ_ε is not a parameter of particular interest here, and we use a relatively uninformative prior distribution by setting $u = 5$.

Appendix 2: Probability distributions with an alternative method[14]

Another method to derive probability distributions from option prices is by using the result from Breeden and Litzenberger (1978). They show that the risk-neutral distribution can be expressed as the second derivative of the option price with respect to strike.

$$f^{\mathbb{Q}} = e^{rT} \frac{\partial^2 C}{\partial K^2}$$

A risk-neutral distribution can therefore be obtained from option prices by numerically differentiating the option price curve twice with respect to strike. Numerical differentiation can however be sensitive to noise in the dataset. We use the following method to construct a probability distribution from option prices with this method:

1. We fit a cubic spline to the implied volatilities across the 18 delta points..
2. Strike and option prices are then computed from the delta values and implied volatilities.

[13] The estimation is run using the PyMC Python package.

[14] See Bahra (1997) and Figlewski for further discussions of alternative methods.

3. The second derivative of price with respect to strike $(\frac{\partial^2 C}{\partial K^2})$ is calculated numerically.
4. $(\frac{\partial^2 C}{\partial K^2})$ is smoothed using a Savitzky-Golay filter

Chart 13: Comparison of two methods

USDNOK. 3-month horizon. 3 October 2022

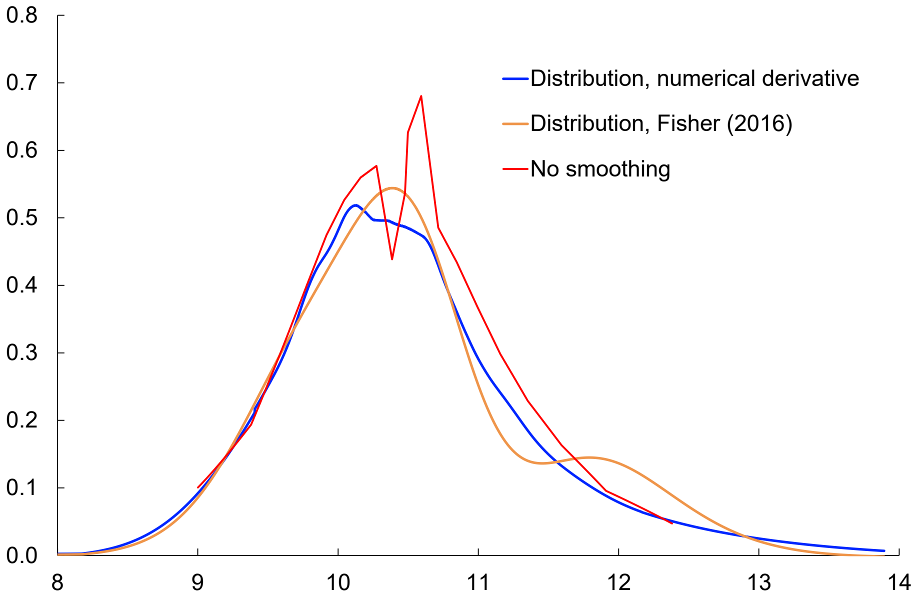

Source: LSEG and Norges Bank

Chart 13 shows a comparison between this method, and the method utilized by Fisher (2016). In addition, we show the distribution as it would be without any smoothing. The key difference between the methods is visible in the right tail of the distribution. The method of Fisher (2016) has a visible secondary mode in the right tail, while the method using numerical derivatives is smoother. The latter method, on the other hand, has a higher probability of the exchange rate exceeding 13 on a 3-month horizon. The fact that the method using numerical derivatives has no secondary modes, can be partially explained by the use of a smoothing filter. The smoothing ensures that the derivatives are less sensitive to noise in the data. At the same time, the smoothing method can also remove multimodality in the distribution.

Appendix 3: Robustness checks

One method of testing the robustness of the estimated distributions is to calculate how well they replicate observed option prices. For calls and puts, the option value can be calculated as:

$$\widehat{C}_i = D \int_K^{\infty} (\mathrm{S} - \mathrm{K_i}) f^{\mathbb{Q}}(S) dS$$

$$\widehat{P}_i = D \int_0^{K} (\mathrm{K_i} - \mathrm{S}) f^{\mathbb{Q}}(S) dS$$

Where $\widehat{C}_i$ and $\widehat{P}_i$ are estimated option prices. To ensure that option prices across different strikes are comparable, we compute implied volatilities from the estimated prices, $\widehat{\sigma_{IV}}$. To measure the accuracy of our estimate, we calculate the root mean squared error between estimated and observed implied volatilities:

$$RMSE = \frac{1}{2}\sum_{i}^{N}\sqrt{\left(\sigma_{IV,i} - \widehat{\sigma_{IV,i}}\right)^2}$$

Where N is the number of options for a given delta. Table 1 shows RMSE for three month USD/NOK options for deltas ranging from 50 to 10 for the period 2007-2025. The results indicate an RMSE of 0.33 to 0.61 percent, with ATM options having higher scores than OTM options.

Delta	50	45	40	35	30	25	20	15	10
RMSE, N=903	0.61 %	0.41 %	0.34 %	0.33 %	0.33 %	0.33 %	0.34 %	0.35 %	0.37 %

We test convergence of the MCMC-algorithm by utilizing a Gelman-Rubin test, also known as a $\hat{R}$ test (Gelman and Rubin, 1992). We estimate the distributions using four independent Markov chains. The $\hat{R}$ statistic measures convergence across the chains. The results are summarized in table 1. Convergence for parameter values across chains are tested by taking the largest $\hat{R}$ values in $\boldsymbol{\beta}$ for each estimation, and testing whether it is below 1.05. A value less than 1.05 indicates convergence.

Table 1: Gelman-Rubin statistics

Exchange rate	Number of estimations	Proportion of R > 1,05
EUR/NOK	143	7.14 %
EUR/USD	134	11.84 %
USD/NOK	128	14.67 %

Using only the largest $\hat{R}$ makes this a rather conservative test. A proportion of 14.67 percent with $\hat{R}$ is greater than 1.05 is still somewhat high however for USD/NOK. This does not necessarily indicate a problem for the robustness of the estimations. $\hat{R}$ greater than 1.05 may indicate that the number of basis distributions is relatively high compared to the number of option prices used in the dataset. In those cases, the algorithm may struggle to separate the value of one β from neighboring values. In that case, the Markov chains may struggle to converge to the same values.

Chart 14: Credible intervals for two EUR/USD distributions

As of 3 November 2021 and 1 December 2021

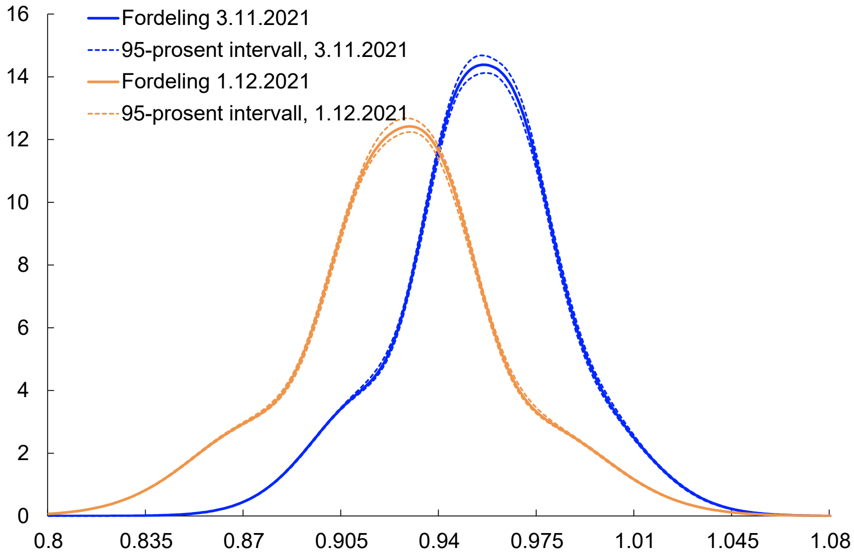

To investigate whether this is a problem more closely, we look at two distributions for EUR/USD estimated on 3 November 2021 and 1 December 2021. On the former date, the $\hat{R} = 1.18$, while on the latter date $\hat{R} = 1.01$. For those dates, we calculate 95 percent credible intervals for the distributions. This is done by calculating a distribution for each posterior draw and then finding the credible interval for the distribution. This is shown in chart 14.

For both distributions, the credible intervals are fairly close to the posterior mean distribution. It therefore indicates that lack of convergence perhaps was not a problem for the distribution as of 3 November 2021.