Option Implied Probability Density Functions: Methodology and Use in Understanding Investor Sentiment

Seamus O'Donnell & Mary O'Keeffe, Financial Markets Division[1]

Abstract

The following paper presents a modelling approach which estimates risk neutral option-implied probability density functions (PDF) from market traded options. An option-implied PDF depicts estimates of future movements of an asset price, as priced by investors. Such an approach is of benefit as it incorporates information from the full distribution of investor beliefs, rather than relying just on the mean expectation. Based on market prices, PDFs indicate the distribution around the mean at various probability bands in graphical form, which can in turn provide insight into how markets are pricing in future movements in the price of an asset and the volatility around these expectations. The purpose of this article is to demonstrate the benefits of such models and provide guidance on how they may be constructed and interpreted.

The article presents an overview of the methodology as well as somewhat more intuitive explanations of the process and outputs. Guidance is also given on the interpretation of model outputs, using options written on Brent Crude Oil as sample PDFs. Finally, a more in-depth study is provided on market expectations for the EUR/USD exchange rate out to March 2017, as at 15 August 2016, with outputs from this tool informing the analysis.

1 The views expressed in this article are those of the authors and are not necessarily those held by the Central Bank of Ireland. The authors would like to acknowledge, with thanks, the helpful comments and suggestions of William Molloy, Brian Gallagher, Patrick Haran, David Doran, Thomas Brophy, Maurice McGuire and other colleagues in the Financial Markets Division of the Central Bank of Ireland.

Introduction

One of the primary ways in which market participants hedge against the risk of movements in asset prices is through the use of derivatives, such as options contracts. Options are financial instruments that give the buyer the right (but not the obligation) to buy or sell an asset at a point in time in the future, at a price agreed at the time of purchase. Through their pricing, information can be extracted which can provide insight into market expectations of the future movements in the price of an asset, in particular, the mean estimate at a point in time and the distribution around the mean within which the asset is expected to fluctuate at a given confidence interval.

In recent years, there has been a great deal of research dedicated to the extraction of information from the options market. One development through this research is the methodology used to construct option implied risk-neutral PDFs. An option-implied PDF provides the expected value of an asset price, as perceived by investors, normally presented in graphical form; a series of PDFs can be linked together to form a fan chart, which illustrates expected values priced in by investors at future points in time. Through these two graphs, a large amount of market information embedded in option pricing can be depicted in two relatively intuitive forms.

The following paper outlines a modelling approach which constructs these two outputs. The method applied follows that described by Bliss and Panigirtzoglou (2000), and the approaches taken by the Bank of England's Clews, Panigirtzoglou and Proudman (2000), the European Central Bank's (ECB) Vincent-Humphreys and Guttierez (2010) and Wright (2016).

The purpose of this article is to demonstrate the effectiveness of such a modelling approach and provide guidance on how they may be constructed and interpreted. In order to demonstrate the benefits associated with the development of an option implied PDF approach, two studies on Brent Crude Oil prices and the EUR/USD exchange rate are also presented.

Section 1 provides an overview of the methodology applied, Section 2 outlines the outputs from the model and the manner in which they may be interpreted using a study on Brent Crude Oil, and Section 3 provides a study on market expectations for the value of the euro versus the dollar as implied by market options, demonstrating the benefits of constructing option implied PDFs.

1. Model Overview

A market based forward rate is the market's aggregate expectation of the future value of a variable, e.g. the Euribor forward curve is the market's mean expectation for the future rates of Euribor at different time horizons. Other derivatives, such as options, can provide additional information regarding investors' expectation of the future movements of assets prices, such as the probability distribution of different future prices around the mean. As previously mentioned, an option gives the buyer the right (but not the obligation) to buy (call options) or sell (put options) an asset at a point in time in the future, at a price agreed at the time of purchase. The agreed price is known as the strike price. The price of an option is also called the option premium.

One of the main factors driving this cost is the 'moneyness' of the option, i.e. the strike price relative to the current price of the underlying asset. For example if the EURUSD is currently at 1.10, and you are seeking to purchase a call option with a strike price of 1.09 this holds value to the option buyer, given that if it were exercised now the option buyer will have the right to buy the asset at less than its current value, and thus must pay for this privilege. This option is therefore considered to be 'in-the-money'. In contrast if the option's strike price were 1.11, it would be of less value given that it does not hold any immediate value, other than the speculative value of potential future price changes.

The premise of any model which constructs option implied PDFs is that the relationship between the expected value of a variable and the amount investors are willing to pay for this option is indicative of the probability assigned to the price matching the strike price.

While all models will be underpinned by this assumption, there exists a range of possibilities regarding the conversion of options data into probabilities. The approach followed here uses the non-parametric method for estimating fixed date expiry PDFs described in Bliss and Panigirtzoglou (2000) and Cooper (2000), with this approach following the result derived by Breeden and Litzenberger (1978). Presented as follows is the modelling approach and technical model overview. A glossary of terms is also provided in Annex 1 which defines a number of terms used throughout.

Modelling approach

The relationship between the value of the underlying asset and the options strike price and premium is used to extract the market assigned probability to a given outcome. The starting point in the modelling process is the collation and sorting of these variables into respective premium-strike pairs.[2] Once these pairs are identified it is possible to calculate the market implied probabilities for each pair from this information.

However, the methodology followed does not infer probabilities directly from the available premium-strike pairs. This has been found to lead to unwanted instability in results given small changes in data inputs i.e. a small change in the cost of one option may lead to a large distortion in the entire distribution, as described by Bliss and Panigirtzoglou (2000). In order to construct a more robust distribution, the Black-Scholes option pricing model is used to transform the collected premium-strike pairs into implied volatility-strike pairs. Implied volatility is the expected degree of variation of the price of the underlying asset, derived from the price of the option, as described by Black and Scholes (1973).

Following this transformation, a discrete set of implied volatility-strike pairs is available for each maturity. In order to depict a full range, and not just this set of discrete points, it is required to firstly interpolate between these data points; this is done through a cubic smoothing spline, as described by Campa et al. (1997). The benefit of using a smoothing spline is that it is a non-parametric approach which produces robust outputs. Subsequently, in order to extend this range beyond the range of traded options (i.e. the last available data points at either side); a quadratic curve is used to extrapolate the data beyond the available data points. This full range now represents what is known as a ‘volatility smile’ given the changing slope of the curve as options move either further in or out of the money at the extremities of the distribution.

The interpolated implied volatility-strike values (volatility smile) are then transformed back into premium-strike pairs. This is done by using the inverse of the Black-Scholes equation previously discussed to convert the implied volatilities into premiums. Finally, using partial differentiation (as outlined in more detail below), probabilities can be implied based on these strike-premium pairs. The results are then graphed in order to produce a series of PDFs which are then used to create a fan chart depicting the distribution of probabilities assigned by the market to a given variable over a range of maturities.

Technical modelling approach overview

For this paper, the non-parametric method for estimating fixed date expiry PDFs described in Bliss and Panigirtzoglou (2000) and Cooper (2000) is used. This method is based upon the result derived by Breeden and Litzenberger (1978), whereby the PDF can be extracted by calculating the second partial derivative of the call price function, extracted from contemporaneous option prices, with respect to the strike price. Unlike other models, this method does not make any assumptions about the probability distribution of the data.

2 European style options are used as data inputs throughout our analysis. European style options can only be exercised on their expiration date (i.e. at a single predefined point in time in the future).

Using the Cox and Ross (1976) stochastic pricing model, the call option price at time t, C_t is defined as the risk neutral expected payoff of the option at the time of the option's maturity, T, discounted back using the risk-free rate;

(1)

$$C(S, K, \tau) = e^{-r\tau} \int_K^\infty (S_T - K) g(S_T) dS_T$$

Where, S_T is the price of the underlying asset at time T, $g(S_T)$ is the risk neutral PDF, K is the strike price of the option, r is the risk-free rate and $\tau = T - t$. A put option price is defined as;

(2)

$$P(S, K, \tau) = e^{-r\tau} \int_0^K (K - S_T) g(S_T) dS_T$$

The PDF can now be inferred directly from these elements; however, as detailed in the previous section, following the results derived from Shimko (1993), it is suggested that better results can be obtained if we use the option premia to calculate implied volatilities at each strike. In order to complete this transformation the Black-Scholes equation is employed. Firstly, implied volatilities are computed by solving numerically the implied value of sigma for each option, as all other values are observable, using the Black-Scholes equation:

(3)

$$C(S, \tau) = N(d_1)S - N(d_2)Ke^{-r\tau}$$

where;

(4)

N=*standard normal cumulative distribution function*

and;

(5)

$$d_1 = \frac{1}{\sigma\sqrt{\tau}}\left[ln\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)\tau\right]$$

and;

(6)

$$d_2 = \frac{1}{\sigma\sqrt{\tau}}\left[ln\left(\frac{S}{K}\right) + \left(r - \frac{\sigma^2}{2}\right)\tau\right]$$

Taking the range of implied volatilities, the discrete set of observations is interpolated across to obtain a larger number of observations along the volatility smile. The Black-Scholes equation is then used once again to transform the interpolated strike-sigma pairs back into strike-premium pairs.

As shown in Breeden and Litzenberger (1978), the risk neutral PDF can be calculated as the second partial derivative of the call price function (1) with respect to the strike price, K;

(7)

$$\frac{\delta^2 C}{\delta K^2} = e^{-r\tau} g(S_T)$$

While this derivation implies continuous strike-price pairs, empirically there are a finite number of observations and cubic spline interpolants are used in order to estimate the PDF across a larger number of strikes. This is done through transformation of the observed values into sigma-strike pairs to form a 'volatility smile' and then interpolating across this smile to obtain more observations for the PDF, as outlined above.

Finally, in order to construct the fan charts, the appropriate interpolated values are found for each specific probability band across each future maturity date. The respective values are joined together to form the fan charts, which can then be used to interpret the implied expectations of the evolution of the path of the underlying variable and associated volatility over time.

Key modelling assumptions

The primary assumption made relates to that of risk-neutrality. Risk neutrality is the assumption that the prices taken from options reflect investors' true expectations of future asset prices. In reality, however, it is likely that some level of non-neutral risk behaviour (risk aversion or seeking) exists in pricing. Therefore this could distort the shape of a PDF; the extent to which this is true is likely dependent on a number of factors, including the asset class considered, the time to maturity, etc. No correction is made within the described model, and so this should be noted when interpreting outputs.

Finally, clarity should be given relating to the use of the Black-Scholes equation within the model. Readers familiar with the equation will be aware that a number of assumptions are made within this quantitative process relating to the relationship between the underlying asset and market efficiency. However, these assumptions do not affect our analysis; the formula is solely used as a transformation process, which is later reversed, thereby imposing none of these assumptions on estimated PDFs.

Chart 1: Brent Crude Oil option implied PDF, Dec 2016

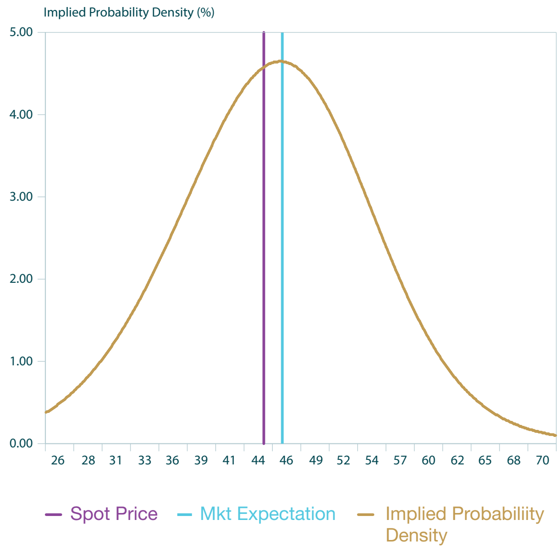

Source: CBI calculations.

2. Model Output and Interpretation

The primary model outputs are the respective PDFs and fan charts. Chart 1 illustrates a sample PDF constructed from options written on Brent Crude Oil with a maturity of December 2016 (with the options recorded as at 15 August 2016). The PDF therefore depicts investor estimates for the price of oil in December 2016. Three parameters define this chart. First, the chart is labelled by the date at which the options expire, in this instance December 2016. Second, the chart's x-axis denotes the investor perceived price of the asset. Here the blue line equates to the mean expectation of the future value, with the central expectation for oil to equal $46 in December 2016. Finally, the chart's y-axis equates to the density at which expectations are observed. The density is measured in probability percentage, with the full area under the curve equalling 100%.

Chart 2 depicts a fan chart of option implied expectations relating to Brent Crude Oil (again with the options recorded as at 15 August 2016). While Chart 1 relates to an individual point prediction, Chart 2 is constructed from a series of PDFs, tying together market expectations across a range of time horizons (from September 2016 out to February 2017), and thus charting the implied path of the underlying asset price, surrounded by 10 per cent probability bands. At each separate time horizon depicted, there is a corresponding PDF being captured. In this instance the red dashed line is the distribution relating to the PDF above, i.e. December 2016. As the time horizon increases, the probability bands widen, which makes sense intuitively given increased uncertainty.

Here again three parameters define the graph. Firstly, the y-axis equates to the maturity of the set of options written. Secondly, the x-axis now equates to the expected price of Brent Crude Oil. Finally, the density at which expectations are observed is captured by the width between each coloured line within the chart. Each purple line, equates to a 10 per cent implied probability band, with the two outside lines equalling the 1st and 99th percentile of the distribution and the black dashed line representing the mean (i.e. 50 per cent probability band). The width between each of these lines is indicative of the level of uncertainty. Here it is evident that in the tails of the distribution, there is a greater dispersion of estimates.

Chart 2: Brent Crude Oil option implied fan chart to Feb 2017

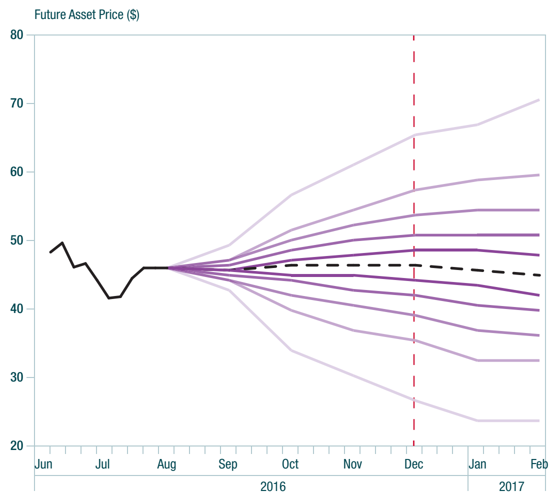

Source: CBI calculations.

A number of summary statistics relating to both the PDFs and fan charts are of particular interest when attempting to draw conclusions, including a distribution's standard deviation, which can be interpreted as the volatility expected by the market in relation to the evolution of the price. Further, the skewness of a distribution reflects the balance of risks around the aggregate expectation of the evolution of a price. If the skew is to the downside, on balance, the market has a greater expectation of downside risks to the price and vice versa.

The analysis of these summary statistics can provide analysts with a significant level of information relating to market expectations. For example, it is evident from Chart 1 above that there is an expectation for oil to remain little changed out to December 2016, with the spot lies only slightly below the mean expectation. However, equally as evident is that there is a large degree of uncertainty around this view, with the range of expectations between $26 and $70 at the respective 1st and 99th percent confidence interval. The skew of the distribution at December 2016 equates to 0.28, meaning that there is a slightly greater weight of expectation for oil to exceed this mean expectation, but not to a great extent. If we look at Chart 2, however, we can see that the expected oil price falls below the spot value beyond the last observation period, with the skew moving towards zero.

3. Market Outlook for the EUR/USD Exchange Rate as at 15 August 2016

The EUR/USD exchange rate plays an important role in economic developments. There are a range of factors which affect the exchange rate between two currencies including interest rate differentials, current account deficits, debt profiles, growth expectations, which can all contribute to exchange rate pricing. With so many contributing factors, it can be very difficult to isolate one driving factor at a particular point in time.

Over the past two and a half years, the euro has seen significant depreciation in value against the dollar. In mid-2014, the euro traded at close to $1.40. However, beginning in May 2014 the euro fell sharply for ten months, depreciating by close to 25 per cent against the dollar from peak to trough (Chart 3). Monetary policy expectations were one likely catalyst for these moves, with a divergence in policy occurring between the Federal Reserve and ECB, as rates were expected to be raised in the US while policy easing was expected to continue in the euro area. More recently, the euro has tended to remain relatively range bound, having recovered somewhat from its March 2015 low.

Chart 3: EUR/USD Spot Rate

Source: Bloomberg.

In order to assess the likely path of the currency pair, as seen by market participants, we are able to look beyond the EUR/USD spot price, and take a forward looking perspective on the cross from the currency options market. This forward looking perspective is likely to encompass a range of variables which would weigh upon the exchange rate out to each respective maturity date. Chart 4 below depicts the fan chart for the EUR/USD pair out to March 2017, as constructed using the approach presented, using data from 15 August 2016. Here it is evident that the central market expectation, marked by the black dashed line, is for the euro to appreciate slightly against the dollar out to the last maturity date. With regards to the balance of risks, across the majority of maturities the distribution is relatively balanced. Each distribution is only marginally positively skewed, indicating that the market sees a slightly greater chance of the euro appreciating beyond its central expectation, rather than falling short of it.

With this outlook in mind, it appears that, all other things being equal, the market expected the euro to increase slightly against the dollar over the coming months. However, in the recent past there has been significant re-evaluation in expectations reflecting a range of factors, including expectations of monetary policy action, illustrating that this type of analysis is very much focused on a point in time.

Chart 4: Euro option implied fan chart to March 2017

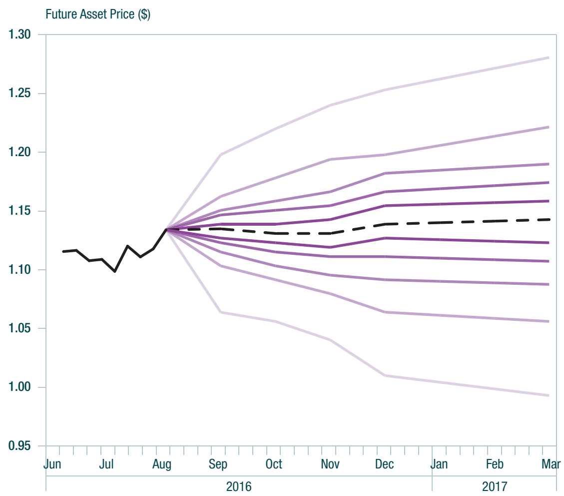

Source: CBI calculations.

4. Summary

Presented in this paper is an approach that constructs PDFs and fan charts from market traded options. These are intuitive graphical representations of information extracted from market options, which provide insight into market expectations of the future movements of financial asset prices, and will complement current market analysis conducted. The purpose of this note is to demonstrate the effectiveness of such models, and provide guidance on the manner in which they are constructed and may be interpreted.

Examples of applying the method were also provided. Firstly, options written on Brent Crude Oil, as at 15 August 2016, were taken to demonstrate the manner in which the model outputs can be interpreted, and how information in the probability bands of a distribution can be extracted in order to complement information regarding the market's mean expectation. A second study of the market's expectation of the EUR/USD exchange rate, as at 15 August 2016, is also presented which demonstrates how outputs from the model can be used in practical application.

Annex 1: Glossary of Terms

Underlying Asset: is the asset which is to be bought or sold on the future date, against which the option is written.

Strike price: a strike price is the price at which an option can be exercised, i.e. the agreed price at which the underlying asset can be bought/sold.

European options: can be exercised only upon an agreed exercise date. Options which grant the right to sell the underlying asset are known as “puts”, with options which grant the right to buy the underlying asset called “calls”.

Non-parametric: means that no assumption has been made in relation to the parameters of the frequency distribution in the fitting process. Inversely a parametric approach makes some assumptions in regards the structure of the probability distribution of the data, and imposes this structure when fitting data.

Skewness: is the third moment of a distribution measuring the asymmetry of the probability distribution. A normal distribution has a skewness of 0; negative scale indicates that there is a longer or fatter tail on the left side of the distribution, and vice versa.

Implied volatility: is that value of the volatility of the underlying instrument which, when input in an option pricing model will return a theoretical value equal to the current market price of the option.

Bibliography

Black, F. and Scholes, M. 1973. “The pricing of options and corporate liabilities.” *Journal of Political Economy* 81: 637-654.

Bliss, R. and Panigirtzoglou, N. 2000. “Testing the stability of implied PDFs.” *Journal of Finance* 59: 407-446.

Breeden, D. T. and Litzenberger, R. H. 1978. “Prices of state-contingent claims implicit in option prices.” *Journal of Business* 51(4): 621–625.

Campa, J., Chang, K. and Reider, R. 1998. “Implied Exchange Rate Distributions: Evidence from OTC Option Markets” *Journal of International Money and Finance* 17: 117-160.

Clews, R., Panigirtzoglou, N. and Proudman, J. 2000. “Recent developments in extracting information from option markets” *Bank of England Quarterly Bulletin* 40(1): 50–60.

Cox, J. and Ross, S. 1976. “The valuation of options for alternative stochastic processes” *Journal of Financial Economics* 3: 145–166.

Elder, R., Kapetanios, G., Taylor, T. and Yates, T. 2005. “Assessing the MPC’s fan charts” *Bank of England Quarterly Bulletin* 45(3): 326-348.

Gutiérrez, J. and Vincent-Humphreys, R. (2010). ‘A quantitative mirror on the Euribor market using implied probability density functions’, Working Paper, European Central Bank.

Wright, J. H. 2016. “Options-Implied Probability Density Functions for Real Interest Rates”, *International Journal of Central Banking* 12(3):129-149.