Instruments used to analyse market expectations: risk-neutral density functions

Prices for derivative financial market products – especially option prices – are a good place to start when it comes to observing and analysing the expectations of financial market participants. Modern financial theory has developed a number of methods of extracting relevant information on market sentiment and making it readily comprehensible.[1] This article introduces a modern method which goes beyond a simple point estimation and presents the entire spread of market expectations of future exchange rate or interest rate developments, for instance. The applicability of this method is then illustrated using current examples. This includes, *inter alia*, the changes in US dollar/euro exchange rate expectations following the terrorist attacks in New York and Washington on 11 September (page 41). As a result, the article will demonstrate that this method can be used to obtain valuable additional information about market sentiment as derived from option prices.

Introduction

In practically all the financial theory methods used, it is assumed that price formation on the financial markets is the result of the com-

1 See also Deutsche Bundesbank, "The information content of derivatives for monetary policy", Monthly Report, November 1995, pages 17 to 32, and Deutsche Bundesbank, "Financial market prices as monetary policy indicators", Monthly Report, July 1998, pages 49 to 66.

plex interplay of rational market players ultimately concerned with evaluating financial assets and liabilities. For their current investment and portfolio decisions, market participants develop notions of how price-relevant factors are likely to change, resulting in the expectations being incorporated in the corresponding direct quotations. Let us take, for instance, the case of an investor who needs to decide between two investments each of which is in a different currency. He has to take account in his calculations not only of current domestic and foreign interest rates for the foreseeable investment period; he also needs to give due consideration to future exchange rates which might apply when an investment denominated in a foreign currency is exchanged back into his domestic currency. To hedge the foreign currency risk, various instruments are available on the forward and options markets at prices which represent, in essence, a market evaluation of that risk. Given certain conditions, information on the density function can be derived from the direct quotations for these instruments. This makes it possible not only to determine the mean of the expectations, but also to assess market players' expectations that a variable will exceed or fall below specific reference values.

Derivative financial instruments and indicators

Definition

Derivatives are normally defined as financial instruments the prices of which are derived from the value of a financial asset called the underlying asset (in this case the US dollar/euro exchange rate, see the box on page 33).

Common to derivative instruments is a contractual clause which specifies the conditions according to which the underlying transaction will be processed in the future. Derivatives are classified according to whether the future transaction is definitive or optional.

Forward transactions

The former category consists of forward transactions in which the two contracting parties commit themselves to buy or sell the underlying asset (in the example given here, the US dollar) at the agreed price when the contract reaches maturity. Depending on the actual spot rate development up to the agreed maturity date, this results in the risks of making a loss being balanced out against the chances of making a profit. The agreed forward rate can be interpreted under specific conditions as the value of the future spot rate expected by the market.

Options

The second category covers call options (put options), by which the options buyer acquires the right (but not the obligation) to purchase the underlying asset (here, the US dollar) from the option seller (to sell the underlying asset to the option seller) at a price fixed in the contract and at a specific date in the future. The market puts a value – the option premium – on the unilateral price change risk thus assumed by the option seller; the option buyer then pays this premium to the option seller. The amount of the premium is, *ceteris paribus*, determined by the difference between the guaranteed strike price and the expected change in the market price of the

Option prices – a glossary

Along with financial futures and forward rate agreements (FRAs), options are a member of the family of derivative instruments whose prices are "derived" from changes in the price of another financial instrument, called the "underlying" asset. Options are traded either on exchanges under standardised terms and conditions or over the counter (OTC) with terms and conditions tailored to customer needs. Derivatives can be used for hedging, arbitrage or speculative purposes, depending on market players' intentions and risk preferences. Option prices – also called option premiums – are calculated using complex mathematical formulae (option pricing models) which are derived from hedging strategies and the absence of arbitrage opportunities.

Call/put option: An option confers the right, but not the obligation, to buy (call option) or sell (put option) a fixed amount of an underlying asset at a price fixed in advance.

Strike price: The price fixed in the option at which the option holder may exercise his/her option to buy (sell) the underlying asset.

European/American option: In a European option, the holder of the option may exercise it only on the fixed expiry date. By contrast, the holder of an American option has the right to exercise it on any day over the entire duration up to the fixed expiry date.

In-the-money (ITM): A call (put) option is in-the-money if the spot price of the underlying asset is higher (lower) than the strike price. In this situation the holder of a call (put) option can buy (sell) the underlying asset from the option writer at the strike price and then resell (buy) it on the spot market at the current price, thereby making a profit.

Out-of-the-money (OTM): A call (put) option is out-of-the-money if the spot price of the underlying asset is higher (lower) than the strike price. For the option holder it is not worthwhile exercising the option.

At-the-money (ATM): A call (put) option is at-the-money if the strike price is exactly equal to the spot price of the underlying. The holder of the option stands to neither gain nor lose from exercising the option.

In the case of foreign exchange options where the value of the underlying asset is a certain amount of foreign currency, the ATM strike price often refers not to the spot price of the underlying asset at the time the deal is closed but to the current forward price of the currency in question ("at-the-money forward").

Combinations

Call and put options are standard option contracts which can be used to devise more complex combination strategies. Two widely used combined foreign exchange instruments are the risk reversal and the strangle.

Risk reversal: A combination of the parallel purchase of an out-of-the-money call option and the sale of an out-of-the-money put option. Both options expire on the same date and have strike prices which are equidistant in percentage terms from the forward rate at the time the agreement is concluded. The market price of the risk reversal can be used to determine whether market players' assessments of the appreciation and depreciation potential of the exchange rate are symmetrical.

Strangle: A combination of an out-of-the-money call option and an out-of-the-money put option, with both being held by the bearer. As in the risk reversal, both options expire on the same date and have strike prices which are equidistant in percentage terms from the forward rate at the time the agreement is concluded. The quotation of the strangle may serve as an indicator of extreme exchange rate fluctuations compared with the log-normal distribution.

Implied volatility: In standard option price models, the option premium for European options can be calculated as a function of contractually specified variables (duration and strike price), data obtainable directly from the market (interest rates and the spot price of the underlying asset) and the expected variance of the underlying asset, which is not directly observable.

Under a given set of parameters, the price of an option in currency units corresponds to exactly one volatility value, which means that one variable may be unambiguously derived from another. Implied volatility is that particular volatility which – using the standard calculation method as a basis – is compatible with the observed market price of the option. It measures the expected price dispersion of the underlying instrument during the option's duration. The distinct mutual convertibility of option premium and implied volatility using the standard option pricing model led, in the special case of OTC foreign exchange options, to the convention of negotiating implied volatilities directly instead of via option premiums. In OTC trading, therefore, quotations are given directly in units of implied volatility, or "vols", which then imply a specific option premium.

Deutsche Bundesbank

underlying asset over the time to maturity of the option in line with the forward rate. This difference is also called the intrinsic value of the option. Both call and put options can be classified according to whether the current constellation of the expected future spot and strike prices on the maturity date is associated with a profit or a (naturally, not realised) loss for the option holder. In the first case an option is referred to as being "in-the-money"; in the second it is said to be "out-of-the-money". If the expected spot price, or the spot price realised at maturity, is equal to the strike price, the option is said to be "at-the-money".

Individual market price indicators

The method of deriving a density function over the future exchange rate presented in this article is based on the prices of four over-the-counter (OTC) derivative financial market instruments; the information content of each instrument is dealt with separately. This involves quotations for the forward exchange rate, the at-the-money call option, the risk reversal and the strangle (for the definition of these terms and the content of the respective contracts, see the box on page 33). These four OTC market instruments are – unlike floor-traded contracts, which have standardised maturity dates – all newly concluded contracts with constant residual times (of, for example, one or three months) to their maturity or settlement dates. As a general rule, new contracts can be concluded on any trading day, meaning that the settlement date is deferred by one day for each additional trading day. The residual times to maturity in most OTC contracts are specified in whole months up to one year. It is possible for a longer residual time to the settlement date to be agreed but it is rather unusual. There is no market for contracts with a diminishing residual time to maturity and no direct quotations are therefore available for them. However, offsetting transactions can be used to close out the position during the residual maturity period.

Forward exchange rate

"Classic" foreign exchange forward contracts have long been used as a hedge against the risk of foreign exchange rate changes over the term of a foreign currency investment. In a perfect market, the principle of the absence of arbitrage opportunities of international financial transactions leads to a forward exchange rate which is fully determined by the current spot rate and the difference between interest rates in the domestic and external money markets. Moreover, if investors are risk-neutral, the forward rate determined in this manner reflects the future spot rate expected by market players. If this were not the case, foreign exchange dealers would resort to corresponding speculative transactions to try to absorb the difference between the forward rate and the expected spot rate until it is completely eradicated.

The foreign exchange forward rate may therefore be used appropriately as an indicator of the expected future exchange rate. However, viewed retrospectively, for most currency relations the correlation between the expected exchange rate changes measured in this way and the actual exchange rate changes observed later is rarely more than weak. The average actual fluctuation margin of exchange rate changes during the term of a contract is so large that the information content of market expectations alone – condensed by the forward rate to one point – is limited.

One explanation for this may be that time-varying risk premiums or the likelihood of structural shifts in the interrelation of currencies – which in a backward-looking analysis appear to be systematic – possibly play a role in determining the forward rate.

Implied volatilities

One indicator which goes beyond the point forecast of the forward rate and provides information about the relative dispersion of the future spot price of the underlying asset is the implied volatility derived from the option prices, of which it is an integral part, that are observable on the market. This is based on the calculation method used by market participants to compute a no-arbitrage price for options with a fixed exercise date (European options). To give a basic example, to calculate the option premium the price volatility of the underlying financial asset expected by market participants is included alongside negotiated contractual elements, such as time to maturity and the strike price, and variables directly evident on the market, such as zero-risk interest rates and the current spot price of the underlying asset. Conversely, it is thus possible to calculate the implied volatility – which cannot be ascertained directly but which is inherent in the option price model – by using the option premium and the known variables. As a general rule, although the calculation of the option premium is dependent on the method used, for OTC currency options there is a convention among dealers which consists in the implied volatilities being cited directly instead of the corresponding option prices or premiums, with the requested option premiums therefore being indicated only indirectly. This quotation practice is based on the Black-Scholes model adapted to calculate the foreign interest rate (although the validity of this model is not necessarily accepted).[2] The indicator of expectation thus observed in the market is a measure of the symmetrical percentage fluctuation margins of the future exchange rate expected by market participants. The implied volatilities therefore provide information about another important feature of the expectations prevailing in the market. In addition to the information specific to the forward rate, the implied volatility is a measure of the average future dispersion. However, no account is taken of other important structural features of the prevailing market sentiment, such as a possible asymmetry – relating to different probability assignments – between a presumed increase or decrease in the underlying financial market price (in the case of the exchange rate, the appreciation or depreciation risk) as well as the probability assessment of extreme exchange rate fluctuations.

Risk reversal

The probabilities of a specific percentage exchange rate appreciation or depreciation may well be assessed differently. The risk reversal is a financial derivative quotation which can

2 The basic model for calculating option prices was provided by Black, F. and M. Scholes (1973), "The Pricing of Options and Corporate Liabilities", Journal of Political Economy, Vol. 81 (2), pages 637 to 654, and was modified for use on the exchange markets by Garman, M. B. and S. W. Kohlhagen (1983), "Foreign Currency and Option Values", Journal of International Money and Finance, Vol. 2 (6), pages 231 to 237.

be used by market analysts to gauge market sentiment in this respect. This is a strategy combining standard option contracts, i.e. the parallel purchase of an out-of-the-money call option and sale of an out-of-the-money put option. If market participants consider it equally likely that the exchange rate could move by a specific percentage in either direction, the risks incurred at both positions cancel each other out, leaving the risk reversal price at zero. By contrast, if the players on the foreign exchange market estimate the potential loss incurred by a put option as a result of the exchange rate moving below the strike price as being higher than the potential profit incurred by a call option as a result of the exchange rate moving above the strike price, the risk reversal has a negative value. In fact, it is clear that, over time, prices quoted for the risk reversal are frequently different from zero, which means that expectations of exchange rate changes moving to a specific extent in one direction or another are asymmetrical.

Strangle

As explained above, the implied volatility in an at-the-money call or put option is a standardised parameter which, in the context of log-normal distributed price changes, adequately describes the average deviation of the exchange rate from its mean movement. However, if one tries to describe the total relative dispersion of the price changes in financial market data by this variable only, the implied volatility of an at-the-money call option, for instance, is inadequate because in reality extreme exchange rate volatilities can be observed more frequently than might be suggested by the log-normal distribution. This can also be seen from the fact that the implied volatility – contrary to the assumptions of the standard option price model – is a variable which changes when the underlying conditions are otherwise constant and the strike price varies, therefore preserving only a local volatility measure dependent on the strike price rather than a general measure. As a rule, the implied volatility increases the more remote the strike price, as established in the option, is from the forward rate ("volatility smile"). The inference is that market participants expect an exchange rate fluctuation margin above that which is compatible with the implied volatility of an at-the-money call option. The direct quotations of the combination of simultaneously acquired out-of-the-money call and put options, known as the strangle, are a measure of these major exchange rate fluctuations expected by market participants. The holder of this instrument only receives a pay-off at maturity if the exchange rate is above the strike price of the call option or below that of the put option. Because the options considered in the context of the strangle are out-of-the-money at the time of purchase (i.e. in the case of call options, the forward rate is below the strike price), the presupposition is that exchange rates will have changed markedly by the exercise date. The willingness to pay for the strangle therefore increases with the risk perceived by market players of an exceptional exchange rate development up to maturity, with the result that the uncertainty assessment based on the observed implied volatility may be appropriately supplemented by the information which can be derived from the price quotations for this instrument.

Implied risk-neutral density functions

Concept and approach

The bits and pieces of information contained in the price quotations of the aforementioned derivative foreign exchange instruments can be gathered together by calculating an implied density; it is not necessary to resort to the use of structural information such as a certain random process for the price development of the underlying asset or a certain option pricing model.[3] This is possible since, assuming that investors are risk-neutral, the value of an option corresponds to the current value of the expected payment to the holder of the option discounted at the risk-free interest rate. Thus, an option premium implicitly reflects the probability assumed by market players that the spot rate of the underlying asset will be higher or lower than the fixed strike price at the expiry date. If at a certain point in time there are several quotations for options with varying strike prices for a certain underlying asset, each of the implied individual probabilities may be used to approximate an implied density; the precision of these estimates increases proportionately to the number of options with varying exercise prices. In an ideal case, a continuum of different option premiums can be extracted from the market and used to derive the exact implied density function.

In the real world of OTC foreign exchange trading, however, there are only a few actively traded call and put options at varying strike prices. The risk reversal and the strangle are not separate instruments in their own right but rather combinations of standard option contracts. To approximate as closely as possible the ideal case of an option pricing quotation which is continuous in its dependency on different strike prices, the few actually observed market prices are supplemented by plausible interim values which take account of what is known as the volatility smile. It is observed that the implied volatility of options that are either in-the-money or out-of-the-money is generally higher than the implied volatility of at-the-money options.

Computation method

These extended data are then used to extract the implied risk-neutral density of the expected changes in the value of the underlying asset. The method used here is set out in detail in the Annex to this article. The density function calculated in that manner reflects

– implied volatility in the width,

– the price of risk reversal in the skewness, and

– the quotation of the strangle in the fatness of the tails.

The chart on page 38 is a stylised depiction of how various market price constellations of at-the-money call options, risk reversals and strangles captured simultaneously can be converted into various densities at a given point in time; they can only be interpreted as density functions if it is assumed that market participants are either indifferent to potential

3 The calculation method is taken from: Malz, A. M. (1997), "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices", The Journal of Derivatives, Vol. 5 (2), pages 18 to 36.

Risk-neutral density functions

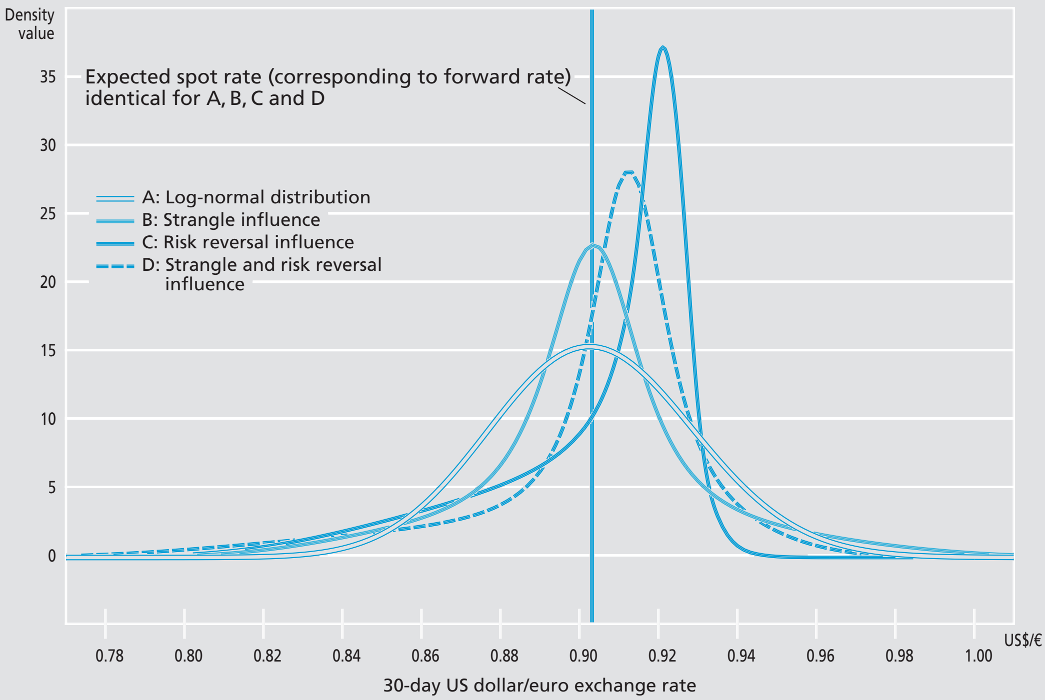

Deutsche Bundesbank

price risks (risk-neutral) or that risk premiums are unimportant.

All densities presented in the chart have the same mean value and are based on the same implied volatility. The mean value implied by the density functions corresponds to the forward rate at the time of computation. An increase in the implied volatility of the at-the-money option would be reflected in the probability distribution having a wider range, so the distance between the density margins would widen.

Log-normal distribution ...

Risk-neutral density function A corresponds to the reference case of a log-normal distribution and represents a situation in which percentage appreciation and depreciation potentials are symmetrically assessed by the market and – measured in terms of the log-normal distribution – no extreme exchange rate movements are expected.

... and the influence of the strangle and the risk reversal

The consequences of a rise in the price of the strangle are shown by density function B. The pattern of the density is sharper above the mean value than the log-normal distribution, whereas its tails are fatter, making extreme exchange rate movements in both directions more likely. The impact of a change in the price of the risk reversal on the calculation of the risk-neutral density can be recognised using the implied risk-neutral density denoted by C. In the case of a negative risk reversal, the implied density "leans" to the right, putting its peak to the right of the average expected spot rate and making it more likely that the US dollar will appreciate by a given percentage than that it will depreciate by the same percentage. Density D represents the combined influences of the risk reversal in C and the strangle in B on the derivation of the implied probabilities.

Advantages of this method of representation ...

The advantage of representing these simultaneously captured market quotations as an implied density function is that complex information can be presented in a comprehensible manner. Not only the mean market expectation but also the potentially asymmetrical or – in terms of the log-normal distribution – extreme range of the expectations around the expected value are filtered out of the data. Those using the implied risk-neutral density thus no longer need to resort to interpreting a point expectation but can see the overall picture of market opinion at a glance.

... and limitations

However, caution is warranted when interpreting risk-neutral density functions and their significance. It must not be forgotten that the relative unavailability of options with various strike prices and the consequent need to resort to an interpolation procedure considerably reduce the informative value of the method. When choosing the option prices to be used for the computation, it should be ensured that the quotations used are actually sufficiently liquid and therefore representative. In addition, it would be desirable to base the computation on the widest possible range of observations. Particularly the derivative prices of options that are far out-of-the-money or far in-the-money – called "wing options" – could improve the information content of the calculated density functions. But this is often where the trouble lies, since it is precisely those wing options that are likely to be best suited to more accurately gauging market expectations of potential extreme swings which are either not available at all or contain high liquidity premiums which distort market expectations of future price developments. When using the method explained here to calculate density functions, therefore, often the only available avenue is to use relatively closely at-the-money option premiums which are generally comparatively liquid, making it possible to disregard the distortions they cause. However, caution is warranted when making interpretations that relate to probabilities outside the range of the density function spread out by the strike prices of the strangles and risk reversals used.

Problems caused by "technical" market tensions ...

The calculation and interpretation of risk-neutral density functions could also lead to problems if, in certain market situations, there are sudden surges in demand which are reflected in sharp fluctuations in the premiums for standard option contracts. As a case in point, there is a strong reciprocity between standard derivatives and specialised or "exotic" options which cover price risk only up to certain upper or lower limits ("knock-out" or "knock-in" options) and cease to function as insurance once the spot rate reaches the barriers fixed in the contracts. If a large number of these thresholds are close together at a given point in time because, for instance, they match certain technical-analysis-based "resistance lines", then, if the borderline is unexpectedly exceeded, the dealers, now openly exposed to price risks, may well demand such a large quantity of hedging instruments that the prices in question will rise sharply, and with them the implied volatilities of these instruments. In such a situation both the holders of the knock-out options, especially the option sellers in question, who, by buying standard option contracts in such a situation, are trying to cap the potential loss of the hedging positions they entered into in connection with the sale of the knock-out options. The meaningfulness of the derivative financial market prices and the density function extracted from them is limited in times of such "technical" market disruptions.[4]

... and risk premiums

When interpreting the risk-neutral densities as pure measures of probability, however, the same caveats apply as those that need to be observed when interpreting the forward rate as an expected value of the future spot rate. As the name suggests, these are density functions derived under the assumption that investors are risk-neutral. This ignores the fact that investors will generally demand a premium for incurring risk and tend more towards risk-aversion rather than risk-neutrality. Only under the assumption of market participants' risk-neutrality, however, are the "pure" probabilities of future exchange rate developments reflected in the calculated densities. Otherwise, they additionally contain a component influenced by the individual risk preferences, which cannot be separately captured and isolated. In practice, the quantitative importance of risk premiums tends to be small and quite possibly affects the mean but not the form of the density, so the distortions that could occur in the risk-neutral densities when calculating according to the method described here are not very large. That is especially the case if one focuses less on the specific density function and more on its change over time.

Possible use

Taking the aforementioned limitations into account, however, the implied risk-neutral density functions and their changes can undoubtedly be used to derive important information on the pattern of the market players' expectations and risk assessments. Four examples have been selected to illustrate this phenomenon. They concern the influence on market players' exchange rate expectations of

- the European Central Bank's intervention in the foreign exchange markets at the end of September 2000;
- the surprisingly unfavourable US economic data announced at the end of October 2000;
- the publication of the unexpectedly sharp decline in the Ifo business climate index in March 2001, and
- the impact of the terrorist attacks in New York and Washington on 11 September 2001.

Case in point: intervention on 22 September 2000

The fall of the euro against the US dollar in autumn 2000, which was exaggerated with

4 A market situation of this kind is described in: Malz, A. M. (1995), "Currency Option Markets and Exchange Rates. A Case Study of the U.S. Dollar in March 1995", Current Issues in Economics and Finance, Federal Reserve Bank of New York, Vol. 1 (4).

Risk-neutral density function: potential application

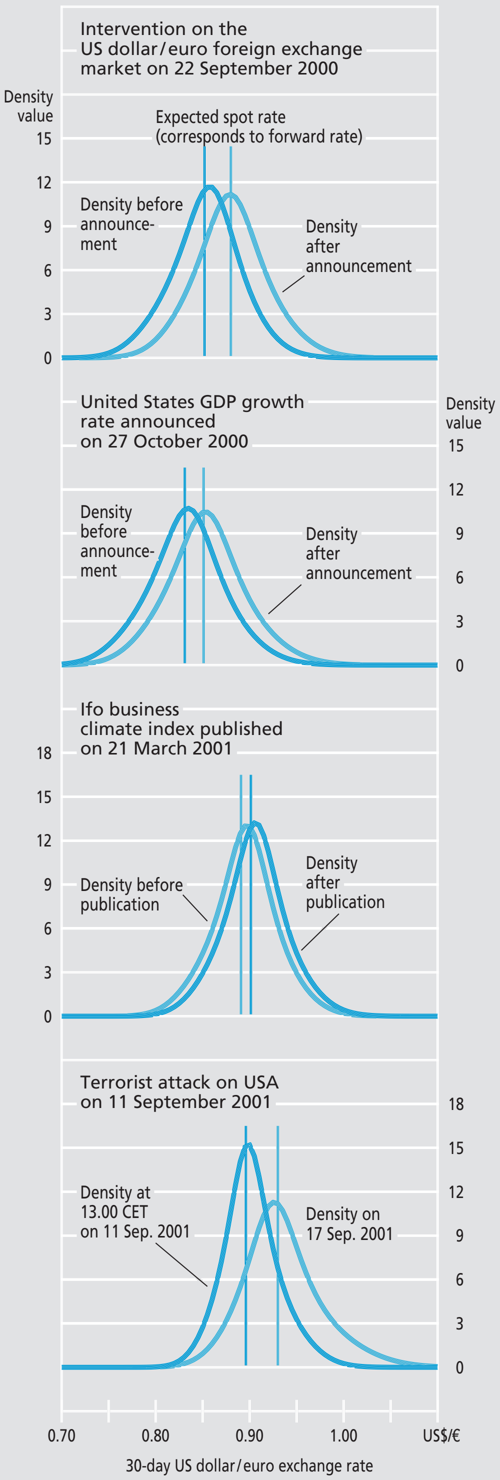

Deutsche Bundesbank

respect to the fundamentals, and the attendant risks to price stability conspired to lead the European Central Bank, in concert with the Federal Reserve and the Bank of Japan, to intervene in defence of the euro on 22 September 2000. The adjacent chart illustrates the risk-neutral density functions with an expectations horizon of one month calculated from the relevant option prices one day before and one day after those interventions.[5] A comparison of the two risk-neutral densities before and after intervention shows that the probability of extreme events (measured in terms of the value of the strangle) was not affected by the intervention operations, whereas the market players' assessment of the general dispersion of the future exchange rate (in the sense of the implied volatility of the at-the-money call option) rose slightly; the rightward-leaning density prior to intervention on the foreign exchange market (captured by the negative price of the risk reversal) straightened up and had a mean value that had shifted distinctly to the right. Thus, after intervention – given the virtual absence of an increase in uncertainty – the market players expected the exchange rate of the euro to be somewhat higher one month later than in the pre-central bank intervention scenario and did not rate the risk of a pronounced devaluation of the euro any higher than that of the euro appreciating by the same amount. However, premature conclusions about the sustainability of the impact of intervention should not be drawn from this event. Some

5 The risk-neutral densities are calculated on the basis of bid-ask middle prices of the relevant derivative financial-market prices which Citibank requests from its London dealers daily at 12 noon GMT.

days later, the aforementioned changes in the calculated density had already receded.

Case in point: announcement of US business data ...

The expectations indicator, besides studying how the central bank's own instruments work, also affords the possibility of studying how other market-relevant factors and events influence market players' assessments. In October 2000, as signs of a gradual economic slump in the United states were mounting, provisional data on US economic growth in the third quarter of 2000 were published. At 2.7 %, growth lagged far behind the 3.5 % expected by market analysts. The chart on page 41 shows the change in the implied risk-neutral densities under the cloud of this new information. On the day after the announcement, the mean expectation of the US dollar/euro exchange rate one month later was somewhat higher than that prior to publication. At the same time the market assessment of price risks on both sides – after having been biased more towards depreciation than towards appreciation by the same amount – tended to be symmetrical. The measures of the mean fluctuation margins and of extreme price fluctuations, however, remained unaffected.

... and publication of the Ifo survey in March 2001

On this side of the Atlantic, it was possible to observe a similar yet opposite reaction to the announcement of unexpectedly bad business figures. In March 2001 the Ifo Institute announced that the German Business Climate Index had dropped by a surprisingly large margin (3.1 index points, although analysts had predicted that the drop would only be 0.5 %). This distinct worsening of the assessment of the business situation in Germany also left its mark on market expectations of the outlook for exchange rates. In concrete terms, the mean expectation for the US dollar rate was pushed slightly downward, and the risk-neutral density tended to continue to lean more closely towards devaluation. Moreover, market players seemed to be generally less certain, since the mean range of future exchange rate fluctuations expanded slightly. Only the fatness at the tails, and thus the assessment of extraordinary exchange rate movements, remained unchanged following the announcement of the index value.

Case in point: terrorist attack in the United States

The 11 September terrorist attack in the United States had a severe impact on market players' US dollar/euro exchange rate expectations. In the period prior to the attack the US dollar regained some of the ground it had lost owing to the gloomier outlook for economic growth in the United States. This was the environment shaken by the news of the attacks in New York City and Washington. The chart on page 41 contains the implied risk-neutral density function computed from option prices immediately preceding the attacks and from market prices one week afterwards (since US financial markets had been closed for a time immediately after the attack). Unlike in the pre-terrorist attack assessment, after 11 September market players expected the spot rate of the euro to be higher one month later. The general assessment of risk shot upwards at the same time. In purely arithmetic terms the value of the implied volatility was 30 % higher than before the attack. In the chart this is evident in the "more ducked" and wider span of the density. The probability of extreme exchange rate fluctuations increased as well. That is evinced by the density having fatter tails. Moreover, the market, which even before the attack had tended to expect the euro to appreciate distinctly rather than to depreciate by the same amount, gave a higher assessment of the potential for asymmetrical exchange rate movements. The value of the risk reversal, which determines the extent of the symmetry break, doubled as a result of the events. In the chart, this is shown by the density leaning more strongly to the left.

At the same time, though, this example makes it clear that caution is warranted when interpreting this indicator. As mentioned above, risk premiums, which are factored out when calculating the densities, may have played a considerable role in the days following 11 September. Moreover, even one week after the attack, the liquidity situation in the OTC market for foreign-exchange options in London had not returned to normal – the rather high bid-ask spreads for otherwise quite accessible instruments were a visible sign of this – meaning that the market assessments could not be estimated as "accurately" as is usually possible.

Summary and evaluation

Financial market prices give a central bank access to information about how the market players assess the future at any given time. This knowledge is important to monetary policy practitioners. However, it is neither useful nor possible to apply indicators mechanistically. Uncertainty surrounding the capture and conversion of the observed market prices into expectation indicators argue against such an approach; financial market data are at times buffeted by special institutional factors or can be distorted by market tension. It will therefore be necessary to continue to reassess the meaningfulness of the presented risk-neutral density function and to apply it to other financial markets. Experience with it will permit economists to make a more broadly based assessment of the instrument described in the foregoing and its ability to describe market expectations.

Annex

Calculating implied risk-neutral density functions

The method used here comprises three steps.

First, the prices (recorded simultaneously and expressed as implied volatilities) of OTC foreign exchange options with various strike prices expressed as deltas are interpolated (I). Second, they are converted into an option price function which is continuous in the strike price (II). Third, the implied risk-neutral density is derived by twice numerically differentiating this function with respect to the strike price (III).[6] Before the individual steps in the

6 The calculation method – based on quotations for European options – is taken from: Malz, A. M. (1997), "Estimating the Probability Distribution of the Future Exchange Rate from Option Prices", The Journal of Derivatives, Vol. 5 (2), pages 18 to 36.

calculation process are discussed, a few explanations on the price quotations of foreign exchange options will be given as a basis from which to start.[7]

OTC quotations

In OTC trading, prices of foreign exchange options are expressed as implied volatilities, which are then converted by dealers using the Garman-Kohlhagen formula (the Black-Scholes formula adapted to calculate foreign currency interest) into the price in currency units, called an option premium, to settle their transactions. The use of the Garman-Kohlhagen formula is merely a market convention and does not imply that market participants necessarily accept the validity of the underlying model. This convention enhances market transparency and simplifies settlement since – in contrast to quotation using currency units – the exchange rates on the spot market, which change minute by minute, do not necessarily cause the price expressed in volatility units to react. Ideally, the option prices quoted in this fashion merely reflect dealers' changing subjective volatility assessments. The Garman-Kohlhagen option pricing formula can be written as follows:

$$c(S_t, \tau, X, \sigma, r, r^*) = e^{-r^*\tau} S_t \Phi(d_1) - e^{-r\tau} X \Phi(d_2)$$

where (1)

$$d_1 = \frac{\ln(S_t / X) + \left(r - r^* + \frac{\sigma^2}{2}\right)\tau}{\sigma\sqrt{\tau}}$$

$$d_2 = d_1 - \sigma\sqrt{\tau}$$

where S_t is the spot rate at the time the option contract is concluded, τ the time to maturity, X the strike price of the option, r the domestic interest rate and r^* the foreign currency interest rate. $\Phi(.)$ is the cumulative standard normal distribution. The parameter σ symbolises the expected price volatility of the underlying asset of the option during time to maturity τ, and c(.) is the option premium to be agreed given this expected volatility. As the contract-specific elements S_t, τ, X, r and r^* are known when the contract is concluded, the option pricing formula, given the option premium c(.), implies precisely one value for σ which can solve equation (1). Therefore the variable σ is also called implied volatility, although the term may in some cases be misleading in connection with OTC foreign exchange options, since in that case dealers negotiate σ directly and thus conversely imply an option premium c(.).

According to conventions among dealers, the strike price is likewise not expressed in currency units but in deltas, which measure the sensitivity of option prices to changes in the price of the underlying asset, which in this case is the exchange rate. Since, *ceteris paribus*, the option's delta decreases continuously as strike prices rise, the strike price and the option's delta are unambiguously connected, which means one can be transformed into the other using the Garman-Kohlhagen formula. The advantage of expressing the strike price as delta is that the "distance" between the strike price and the current forward rate is normalised. This means that contracts with the same delta always show forward rate/strike price constellations with, in terms of the forward rate, the same percentage distance between the strike price and the forward rate. Specifically, in out-of-the-money call options (in-the-money call options) the contracting parties mostly select strike prices such that they corres-

7 Risk-neutral density functions can also be calculated using US foreign exchange options. See the discussion paper by Craig, B. and J. Keller, "The Empirical Performance of Option Based Densities of Foreign Exchange", Economic Research Centre of the Deutsche Bundesbank, to be published at the end of the year.

pond to a delta of 0.25 (0.75) or, less frequently, to 0.10 (0.90).

Analytically speaking, the delta of a call option results from the first derivative of the option price formula with respect to the spot rate of the underlying asset.

$$\delta = e^{-r^{*}\tau}\Phi(d_1),$$

where $\Phi(d1)$ corresponds to the value of the cumulative standard normal deviation evaluated at d_1.

OTC quotations for combinations

This quotation convention exists not only for standard option contracts (call and put options) but also for various combinations such as the risk reversal and the strangle.

A risk reversal is a combination of the purchase of a call option and the sale of a put option, the strike price deltas of which (here: δ = 0.25 for the call option, δ = – 0.25 for the put option) are equidistant from the forward rate. Its price therefore corresponds to the difference, expressed in implied volatility units or "vols", between the two instruments used to construct it (see the glossary on page 33). The formula for the quotation of a risk reversal is expressed as:

$$RR_t^{\delta = 0.25} = \sigma_t^{\delta = 0.25} - \sigma_t^{\delta = 0.75},$$

where $\sigma_t^{\delta = 0.25}$ ($\sigma_t^{\delta = 0.75}$) represents the implied volatility price of a call option (put option) with a delta of 0.25 (– 0.25).[8]

A strangle is the parallel purchase (or sale) of a call and put option with – in absolute terms – the same delta (see the glossary on page 33). The price of this combination is expressed as the deviation of the average of the call and put options contained therein from the price of an at-the-money call option.

The formula is:

$$ST_t^{\delta = 0.25} = 0.5\,(\sigma_t^{\delta = 0.25} + \sigma_t^{\delta = 0.75}) - \sigma_t^{\text{"at-the-money"}},$$

where $\sigma_t^{\text{"at-the-money"}}$ is the implied volatility price of an at-the-money call option.

Rearranging terms will enable us to take the market prices of the risk reversal and the strangle and to deduce from them the implied volatilities of the call options underlying those combinations, the delta of which is 0.25 and 0.75, respectively:

$$\sigma_t^{\delta = 0.25} = ST_t^{\delta = 0.25} + \sigma_t^{\text{"at-the-money"}} + 0.5\,RR_t^{\delta = 0.25}$$

$$\sigma_t^{\delta = 0.75} = ST_t^{\delta = 0.25} + \sigma_t^{\text{"at-the-money"}} - 0.5\,RR_t^{\delta = 0.25}$$

Besides the quotation of the at-the-money call option ($\sigma_t^{\text{"at-the-money"}}$) with a delta of around 0.5, there are now two additional option prices expressed in implied volatilities with strike prices expressed in delta.

Calculating the implied risk-neutral density function

I. Interpolating the market prices

These three prices, expressed in implied volatilities, form the basis for an interpolation where it is as-

8 Because of the put-call parity, the delta of the put option, – 0.25, corresponds to a delta of the call option of around 0.75. The negative signs of the delta of the put option are left out in floor parlance.

sumed that the non-observed volatility prices of options with differing levels of delta form a sickle-shaped pattern around the at-the-money call option, so they can be approximated by a parabola. This non-linear arrangement of the implied volatility prices is also called the volatility smile. In algebraic terms, this phenomenon can be expressed as a second-order polynomial:

$$\sigma_t^\delta = \alpha_0 + \alpha_1\delta + \alpha_2\delta^2,$$

where σ_t^δ denotes the implied volatility price of a call option with a delta δ. Since this equation only contains three variables α_0, α_1, α_2, it can be unambiguously solved at any point in time using the three simultaneously observed OTC quotations of the at-the-money call option, the risk reversal and the strangle.

II. Transformation into an option price function where the strike price is continuous

Using the Garman-Kohlhagen formula in (1), every δ/σ quotation can be numerically assigned to a pair consisting of the strike price X and the option premium c(.) with the same informative value. In the following, it will be shown that numerically differentiating the option premiums twice with respect to the strike price will lead to the desired implied probability distribution.

III. Implied risk-neutral density function

The option premium c(t, X, T) of a European call option with the strike price X and time to maturity $\tau = T\text{-}t$ at time t is the result, if market players are risk-neutral, of discounting the expected value of the option's pay-offs at the date of expiry T by the interest rate r of a risk-free investment. Since a call option is only exercised if price movements are favourable, i.e. only if the exchange rate S_T is higher than the strike price X, the value of the option is between zero and the difference between the exchange rate S_T and the strike price X on the date of expiry. The formula for this is:

$$\begin{aligned} c(t, X, T) &= e^{-r\tau} E[\max(S_T - X), 0] \\ &= e^{-r\tau} \int_X^\infty (S_T - X)\, \pi(S_T)\, d(S_T), \qquad (2) \end{aligned}$$

where $\pi(S_T)$ is the probability density function over the set of potential realisations of S_T assumed by dealers at the time of the transaction. This means that the observable option prices c(.) imply information on the density functions assumed by dealers. The information content increases in line with a rise in the number of independent option premiums for options with varying strike prices which may be taken from the market. In an ideal case, a continuous function is available. In that case, it is possible to extract the implied probability density function by numerically differentiating (2) twice with respect to the strike price X. Taking the first derivative of the option price function yields the following relationship:

$$\begin{aligned} \frac{\partial c(t, X, T)}{\partial X} &= -e^{-r\tau} \int_X^\infty \pi(S_T)\, d(S_T) \\ &= -e^{-r\tau} [1 - \Pi(X)] \\ &= -e^{-r\tau} P(S_T > X), \end{aligned}$$

where $\Pi(S_T)$ is the cumulative distribution function over the set of possible realisations of S_T and $P(S_T > X)$ the probability that the exchange rate will exceed the strike price X on the day of maturity. If the first derivative is evaluated at two different points X_1 and X_2, the difference between $P(X_1)$ and $P(X_2)$ (where $P(X_1) > P(X_2)$) may be used to calculate the probability implied by market players that S_T will be between X_1 and X_2. For infinitesimally small variations in the strike price this probability may be extracted by once again differentiating the option price function twice with respect to the strike price

$$\frac{\partial^2 c(t, X, T)}{\partial X^2} = e^{-r\tau}\,\pi(X). \qquad (3)$$

The second derivative of the option price function therefore corresponds to the (discounted) value of the probability density function over the set of realisations of S_T, evaluated at strike price X.

As mentioned above, the practical implementation of this approach to calculating the implied risk-neutral probability density functions causes a problem: there exists neither a uniform analytical option price formula nor a large number of market prices for options with different strike prices. Through the interpolation in I and the step-by-step mathematical transformation in II, a continuum of option premiums c(t, X, T) at varying strike prices X is established. By forming the difference quotient for, in principle, infinitely small step sizes, the difference quotient in (3) is approximated. The numerically derived cumulative distribution function $\hat{\Pi}(X)$ is calculated as follows:

$$\hat{\Pi}(X) = 1 + e^{r\tau}\left[\frac{c(t, X, T) - c(t, X - \Delta X, T)}{\Delta X}\right]$$

where $\frac{c(t, X, T) - c(t, X - \Delta X, T)}{\Delta X}$ is the first-order difference quotient of the numerically existing option price function and ΔX the step size of the discrete differences (which can be made infinitely small). The first-order difference quotient of the cumulative distribution function $\hat{\pi}(X)$ then yields the desired implied risk-neutral probability density function $\hat{\Pi}(X)$:

$$\hat{\pi}(X) = \frac{\hat{\Pi}(X) - \hat{\Pi}(X - \Delta X)}{\Delta X}$$

When interpreting the implied probabilities gleaned in this manner, however, it must be borne in mind that information on the mass of probabilities between and beyond the areas defined by the strike prices of traded options is chiefly dependent on the method of interpolation. That is particularly true of the two ends of the probability density function.