Using option prices to measure financial market views about balances of risk to future asset prices

By Damien Lynch and Nikolaos Panigirtzoglou of the Bank's Monetary Instruments and Markets Division and George Kapetanios of the Bank's Conjunctural Assessment and Projections Division.

Probability density functions (pdfs), implied by prices of traded options, are often used by the Bank to examine financial market expectations about future levels of different asset prices. This article examines how information about one aspect of such expectations — views on balances of risk — for future asset prices may be inferred from the degree of asymmetry of an implied pdf. We first look at the general issue of choosing a statistic to summarise the degree of asymmetry of any pdf. The choice of units when measuring changes in the underlying asset price is then considered. Finally, we examine empirically the implications of using various asymmetry measures when relating the information from option-implied pdfs to market views about balances of risk to future asset prices.

1 Introduction

Financial assets are usually valued as the discounted sum of expected future cash flows from holding the asset. Viewed in this way, market prices may be thought to embody an aggregate 'market view' about expected future cash flows, discount rates and any other variables used in their valuation. These views, in turn, are likely to be related to the expected future economic environment and so asset prices may provide policymakers with a source of information about market expectations of future economic prospects.

Furthermore, derivatives traded on these assets allow market participants to take views on the future values of the assets themselves. Option markets are an important example of this. Previous work published by the Bank of England has illustrated the sort of forward-looking information that options embody.(1) One useful application for policymakers is that they can use option prices to infer a set of probabilities attached by financial markets to various future asset price levels. In the jargon this is referred to as an option-implied probability density function (pdf) for the price of the underlying asset in the future. The width of the pdf will reflect uncertainty about future asset prices. And the extent to which the pdf is asymmetric can potentially tell us about market views on the relative risks that future asset prices will be higher or lower, the so-called 'balance of risks'.(2)

In inferring this sort of information from implied pdfs it is important to bear in mind that the pdfs are extracted under the assumption that investors are risk-neutral, that is, investors do not require any compensation for bearing risk. However, investors are more likely to be risk-averse and so care about risk. As a result, the risk-neutral option-implied pdfs will reflect both investor preferences toward risk and market participants' 'true' pdf.(3)

Information about the shape of implied pdfs for different asset prices forms part of the information set regularly examined by the Bank in pursuing its two Core Purposes of monetary stability and financial stability. It is primarily of use in helping policymakers to understand market expectations about a range of future asset prices — and, by extension, perhaps the economy. For monetary stability, interest rate probability distributions implied by option prices are one way of assessing market views about risks around the path of expected future interest rates. Such views could reflect market uncertainty about the monetary policy reaction function or about the nature of exogenous risks facing future interest rates and the economy. Turning to financial

(1) See Bahra (1997) and Bliss and Panigirtzoglou (2002), Clews, Panigirtzoglou and Proudman (2000) and Moessner (2001) for examples of the use of option-implied pdfs by central banks.
(2) Another example of information about market perceptions of risks to asset prices that we can potentially extract from these distributions is the probability of extreme moves in asset prices. This is reflected in the concentration of probability in the tails of the distribution.
(3) We abstract from the effects of risk preferences for the purpose of this article. For more on this issue see Bliss and Panigirtzoglou (2004).

stability, information from option prices could be useful in monitoring and assessing potential risks to the financial system. For example, concentrations of probability in the tails of the probability distributions for future asset prices may indicate growing perceptions of a risk of unusual movements in asset prices.(1) More generally, indicators from implied distributions are widely used in the *Inflation Report*, *Financial Stability Review* and *Quarterly Bulletin* in interpreting and reconciling developments in a wide range of asset prices including interest rates, exchange rates, oil and equity prices.(2)

Some examples of pdfs for short-term UK interest rates (three-month Libor) are shown in Chart 1. The relative width of the pdfs suggests that market uncertainty declined between October 2002 and August 2003. In addition, relative to 2002, the pdf became more symmetrical in 2003. This suggests that market views had moved from attaching a greater risk to lower, relative to higher, future UK interest rates in October 2002, towards a more neutral view of the balance of risks by August 2003.

Chart 1
Option-implied pdfs for three-month Libor in six months' time

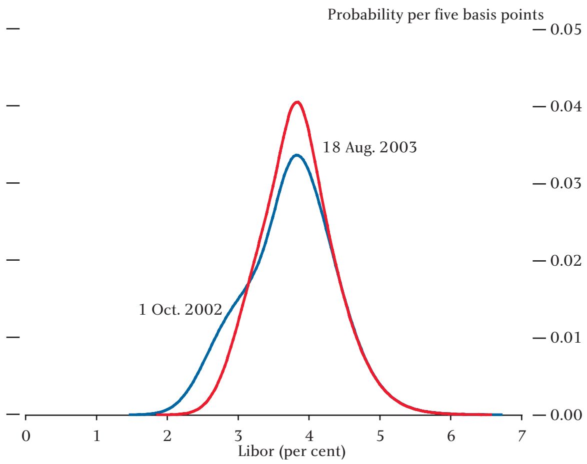

Comparing how probability distributions — or market views — have changed between two dates is easily achieved by such visual inspection. But we may wish to compare the shapes of distributions on many dates to say something about how these market views of the balance of risks have changed over an extended time period. In this case, visual analysis is less useful; instead we need to be able to measure the degree of symmetry or asymmetry of a distribution and, having done so, relate this measure to market views about the balance of risks. That is, we need a statistic — a number each day — that can summarise the shape of the distribution and that can be compared across different days. How might we go about constructing such a measure?

Statistical theory can provide us with some guidance. There are a number of well established measures for evaluating degrees of asymmetry of probability distributions. Examples include the skew coefficient or a comparison of aggregate probabilities above and below a particular point in the distribution. In deciding which measure to use, there are, however, a number of issues that we must first address. This article looks at what we need to consider in terms of choosing a measure of market views about balances of risk to future asset prices. We look at the general question of choosing a statistic to summarise the degree of asymmetry of any pdf, ie not just those implied by option prices.(3) Section 2 illustrates the potential pitfalls we face when making this choice. Our goal is then to identify measures of asymmetry that are consistent with the other information that we may take from a pdf — such as a view of the most likely outcome or the expected outcome.

We begin by recognising the need to specify a reference point, around which to look at asymmetry. To this end, we employ a so-called 'loss function' which we may combine with information from a pdf to guide us in our choice. Section 3 introduces the concept of a loss function and looks at some alternative functions and what they imply for our choice of reference point. Section 4 continues with this framework in obtaining measures of pdf asymmetry that are consistent with common reference points in a pdf.

Having set out some measures of asymmetry, we look at how the choice of units when measuring changes in the underlying asset price can affect the shape of the option-implied pdfs. For example, should we look at a

(1) The usefulness of the information from option prices is subject to the liquidity of the option contracts. In particular, a lack of liquidity in contracts that are far away from the money — that is, in the tails of the implied distributions — means that reliably estimating the tails can be difficult. Liquidity also tends to diminish the longer the time to expiry of the contract. In addition, there may be liquidity differences between call and put options so that the upper and lower tails of the pdf may differ due to liquidity premia. As a result, implied pdf asymmetry can reflect factors other than market views on balances of risk. The reliability of the pdf estimates is also subject to the smoothing of prices by exchanges in calculating settlement prices and the discrete nature of both option price tick sizes and exercise prices.
(2) For examples see *Inflation Report*, August 2004, Section 6.2; *Financial Stability Review*, December 2003, pages 12–13; *Quarterly Bulletin*, Summer 2004, page 116.
(3) However, throughout our discussion we illustrate our thinking using examples with pdfs implied by option prices.

probability distribution for changes in the level of the asset price or changes in the logarithm of the asset price? This could have implications for how we relate a measure of asymmetry for pdfs to market views on the balance of risk to future asset prices. Section 5 motivates this point further and sets out our analysis.

Finally, we consider some empirical findings on the relationships between the various measures of asymmetry suggested as well as with other frequently used measures.

2 Pdfs and asymmetry — setting up the problem

To help to understand the problem of choosing an appropriate measure of asymmetry, consider the example of someone, say a policymaker, presented with a set of probabilities for different values of a random variable occurring. In principle, this random variable could be anything — from the number of sunny days in the United Kingdom during summer to the level of the sterling/euro exchange rate in six months' time. Continuing with our asset price focus, let us take the random variable to be the level of the FTSE 100 equity index in six months' time.

As for our interest rate example above, the policymaker can plot the probability associated with each potential FTSE 100 level to get a visual idea of the distribution of probability — we do this in Chart 2. Now the policymaker, seeking to summarise the information in the distribution, would like to choose a 'point' estimate of the future level of the FTSE 100, six months hence. Such point estimates are usually chosen from the 'centre' of the distribution and are often referred to as measures of central tendency. Common examples include the outcome with the highest probability — the mode, or the 'expected' outcome — and the mean, calculated as the sum of all outcomes, weighted by their probability.

Let us suppose the policymaker chooses the mean outcome as a point estimate. Relying on a sole point estimate may not be advisable and so the policymaker will also want to know the spread of outcomes around this point and whether the risks around this point are stacked in one direction more than another (ie the 'balance of risks' mentioned earlier). How might the policymaker measure the balance of risks around the point estimate? One way might be to measure the difference between the probability, in aggregate, attached to outcomes above and below the point estimate. Alternatively, we could look to statistical theory and use the well established method of calculating the degree of skew of the probability distribution.(1) We use both and compare the results.

Chart 2 shows that the distribution has a longer lower than upper tail. Visually we would say it has negative asymmetry or that it is negatively skewed. Calculating the skew coefficient confirms this: it is around -1. From this we might infer that the risks around our reference point — the mean — are tilted towards lower, rather than higher, outcomes. However, measuring the probability attached to the outcome being above and below the mean, we arrive at a different conclusion.

Chart 2
Option-implied pdf for FTSE 100 in six months' time on 22 March 2004

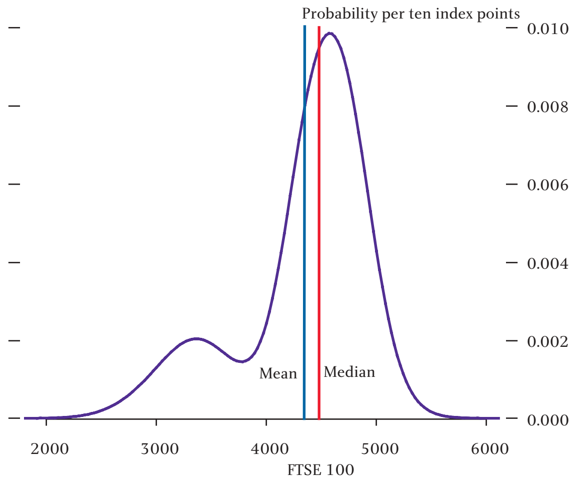

To see this, Chart 2 also plots the median of the distribution: that point at which there is equal probability, in aggregate, attached to outcomes above and below it. The mean outcome lies below the median, so the probability attached to outcomes above the mean is greater than that attached to outcomes below the mean. We find that there is about 20 percentage points more probability attached to FTSE 100 outcomes above the mean than to those below the mean. So from this we might infer that the balance of risks around the mean FTSE 100 outcome is actually positive — in contrast to what the skew measure had indicated. How can we understand this difference?

(1) The skew is calculated by summing up each of the distances between potential outcomes and the mean, raised to the power of three, and multiplied by the probability attached to each outcome. This sum is then divided by the third power of the standard deviation of the pdf to adjust for any effects due to changes in the width of the pdf.

3 Pdfs, reference points and loss functions — a framework for thinking about asymmetry

In our example above, the first step in measuring the degree of asymmetry of a pdf lay in our choice of a point estimate or a reference point in the distribution. Having suggested some points of central tendency as candidates, we arbitrarily picked the mean. We then sought to apply some well known measure of asymmetry to the distribution and to infer something about the balance of risks around this reference point. But this choice is not as innocuous as it may seem. Below we consider the choice of a reference point in a more structured framework — that is, in the context of loss functions — and show that this choice can depend, either implicitly or explicitly, on the preferences of the person making the estimate.(1) Our choice of asymmetry measure should also be consistent with these preferences. Thus we should not be surprised that simply picking an *ad-hoc* reference point and applying an *ad-hoc* measure of asymmetry around it can provide conflicting indications about the balance of risk in the distribution. So our challenge is to use this preference dimension or loss function framework to derive measures of asymmetry associated with common measures of central tendency that we might use as point estimates.

In principle, any point in a probability distribution may be used to provide a reference point. Can we set ourselves some criterion against which to judge which point is best? An obvious reference point is our best point estimate of the future value of the variable, given our preferences. So at a simple level, one criterion is to say that an estimate is 'best' when it exactly matches the subsequent outturn and that it suffers a loss of 'quality' when it differs from it. Pursuing this line, we can quantify this loss of quality by using what is termed a loss function — a mathematical function that sets out the loss or penalty incurred in picking an estimate that is not the best that we could choose.(2) In economics, loss functions are complementary to utility functions — with the former we measure a cost/dissatisfaction associated with a particular event whereas the latter measure the benefit/satisfaction associated with an event. As a result, an individual's preferences as revealed in their utility function will also be revealed in the corresponding loss function.(3)

How we measure this loss might depend on the purpose for which an individual selects a point estimate. So, for example, one individual may express a preference that puts all emphasis on getting the forecast exactly right such that incorrect forecasts are equally bad, be they incorrect by a very small margin or a very large margin.(4) Another may be more willing to accept small errors so that forecasts that are close, but not equal, to the actual outturn are valued more (or penalised less) than those that are very far away from it.

Some visual examples may help to cement the idea; Chart 3 shows three loss functions that are often used in economics, engineering and other sciences: the quadratic, indicator and absolute loss functions. The horizontal axis shows the set of point estimates for the random variable that we may choose from. Also marked on the horizontal axis is the location of the actual outturn. The vertical axis measures the loss that one would incur for each point on the horizontal axis as the point forecast, given the actual outturn.(5) For all three

Chart 3
Alternative loss functions

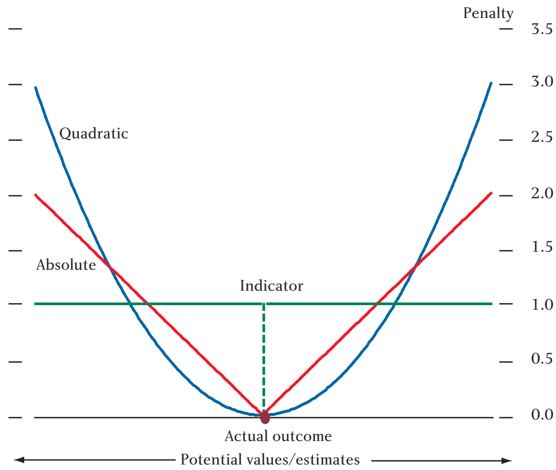

(1) In our example, we have chosen to refer to this person as a policymaker. However, when thinking about pdfs implied by option prices we are examining the 'market's' probability distribution or forecasts so should we not be concerned with the preferences of the market? For now, we continue with our example and address this point at the end of Section 4.

(2) This is analogous to a problem in engineering where products from a production process need to be monitored/assessed to see if their quality matches the desired specifications of the product. Loss functions are often employed as tools to deal with the problem. See Joseph (2004) for more details.

(3) In a previous *Quarterly Bulletin* article (see Vickers (1998)), possible loss functions for the Monetary Policy Committee were examined in terms of theoretically describing potential preferences of the Committee in pursuing its inflation objectives as specified in the Bank of England Act 1998. For more examples of applications involving loss functions in economics, see Svensson (2004).

(4) Consider someone placing a conventional bet on a horse race: the nature of the bet will mean that picking the winner and getting the forecast right is crucial and so the punter will place no value on forecasts that are incorrect by a small (eg second place) or a large (eg second-last place) margin.

(5) To facilitate a comparison of the three loss functions, the losses calculated under the quadratic and absolute loss functions are normalised by dividing by the average loss under the respective function.

functions, a zero loss is incurred when the estimate that is picked is the same as the outturn. Picking points away from this best estimate incurs positive loss and the different mathematical functions are designed to show some alternative 'loss schemes'. The three alternative schemes imply three different 'attitudes' towards alternative point estimates.

- Beginning with the most basic loss function — the indicator loss function — all points that are different from the best estimate are deemed to incur the same penalty. In other words, all value is placed on picking the best estimate and all other potential forecasts are viewed as being of equally poor quality.

- In contrast, the quadratic and absolute loss functions penalise different points according to how far they are from the best point estimate (see the appendix for mathematical definitions). Those points that are 'close' to the best estimate incur a smaller penalty under the quadratic loss function than under the absolute loss function. But moving further away from the best estimate, the quadratic loss function gradually begins to penalise mistakes more than the absolute loss function.

The criterion set out above was based on choosing the best estimate as the one that is equal to the actual outturn. Of course, this is of little practical use to us — the reason that we are picking an estimate is because we do not know what the outturn will be. So, remaining within the loss framework, we instead need to think about the loss we would be expected to incur were we to pick a given point as our best estimate.(1) To quantify expected losses we need to use the information we have about the probabilities attached to different outcomes — the probability distribution. So we can identify the best estimate to choose as the point that we expect to result in the smallest loss, given the probability that we attach to each outcome being realised. In this sense the 'best' estimate now depends both on the probability distribution of possible outcomes, as seen by the person selecting the estimate, and on individual preferences.

We show in the appendix that the 'best' point estimate differs across the three loss functions as follows:

- For an indicator loss function, the best point estimate is the mode of the distribution (the most likely outcome). Intuitively, this makes sense; remember that our exercise is to choose the estimate that will minimise our expected loss, given the probability attached to each point actually occurring. And our loss function is such that we have an all-or-nothing character. It then follows that the logical thing to do must be to pick that point that is most likely to occur — the mode — as our best forecast given our preferences.

- The mean of the distribution (or the average of all possible outcomes) is the best point estimate under a set of preferences given by a quadratic loss function.

- Finally, the best point estimate with an absolute loss function is the median (the point in the distribution such that there is equal probability of the outcome being higher or lower than it).

4 Characterising and measuring pdf asymmetry

The previous section looked at how we could use loss functions as a tool in helping us to choose our best point estimate from our probability distribution. How can we use the concept of a loss function to arrive at measures of pdf asymmetry that are consistent with the best point estimates that our loss framework provides us with? We start by measuring asymmetry in terms of the difference between the expected losses attached to outcomes above and below the point estimate.(2) By doing so we are assessing the balance of risks around a reference point (ie making a relative assessment of the upside and downside risks).

So for a given distribution and loss function, we first need to compute our best estimate and then we may use the relative expected losses around this reference point to compute the associated asymmetry measure. Taking the three loss functions mentioned above, we can derive the asymmetry measures that are consistent with each of them; these are shown in Table A1.1 in Appendix 1. We briefly set out their key properties here:

- Taking the indicator loss function first, we mentioned earlier that the mode was the

(1) It is important to acknowledge that focusing on 'expected loss' is in itself a preference-based choice. For example, one could choose to minimise the modal loss or the median loss.
(2) We standardise the difference in expected losses by dividing by the total expected loss.

best point estimate and hence that is our reference point. Under this loss function, the difference between expected losses above and below the mode is shown to be the difference between the aggregate probabilities of the future outcome being above, and below, the mode. Put simply, our asymmetry measure is just the difference in the probability masses above and below the mode.

- Turning to the absolute loss function, the loss appears slightly less intuitive, with the relative expected losses now the standardised sum of probabilities weighted by the distance, in absolute terms, of each potential outcome from the reference point (the median). This has a simple form driven by the difference between the mean and the median.

- Finally, asymmetry under the quadratic loss function is measured in terms of squared distances of future outcomes from the mean outcome weighted by the corresponding probabilities. As a result, those outcomes that are further away from the mean will have a proportionately greater influence in determining the magnitude/sign of our measure. Just how much influence they have will be determined by their probabilities. In this sense, this asymmetry measure is closely related to the statistical measure of skew.

For all measures, a positive (negative) number indicates a greater expected loss attached to outcomes above (below) the central projection than to those below (above). In terms of the option-implied pdfs, the positive (negative) number would indicate that the market views the balance of risks to point to a relative upward (downward) risk to asset prices. In the case of a unimodel symmetric probability distribution, the mean, mode and median all coincide and so the best point estimate under each of the three loss functions is the same point in the probability distribution. In this case the asymmetry of the distribution will be zero, reflecting the fact that the expected losses above and below the single reference point are equal, regardless of the loss function with which they are measured. So, for example, the bell-shaped 'normal' probability curve — a frequently used symmetric distribution — has zero asymmetry under all three loss functions discussed above. This feature of the normal distribution means that it is a useful benchmark when assessing the degree of asymmetry of probability distributions.

To illustrate how the loss function combines with a pdf to produce an expected loss function we examine an option-implied pdf for the FTSE 100 equity index and the three loss functions in Chart 4. The lower part of Chart 4 illustrates the expected losses at each index level for the three loss functions, with the upper part showing the FTSE 100 pdf. That is, for each level of the FTSE 100, we evaluate the expected loss were that level chosen as our best estimate. Plotting the resulting expected losses against associated FTSE 100 levels provides the expected loss functions in the lower part of Chart 4. The expected loss for the indicator loss function is minimised on the mode; that for the quadratic on the mean; and that for the absolute on the median. The asymmetry of expected loss for projections above and below the best estimates is also evident, reflecting the negative asymmetry of the probability distribution. This negative asymmetry arises because market participants are paying more for insurance against a large fall in the FTSE 100 than they are for protection against a corresponding large rise in the index.

Chart 4
FTSE 100 pdf, expected loss functions and optimal central projections

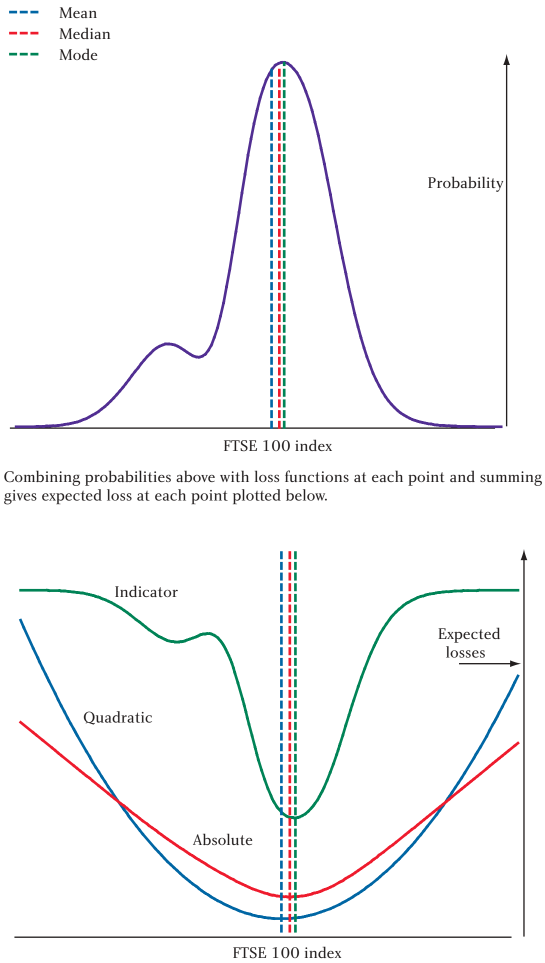

We should now be able to understand why we obtained seemingly inconsistent measures of asymmetry for the FTSE 100 pdf in the example in Section 2. Our loss-based framework implies that loss functions, central projections and asymmetry are linked. It was the arbitrary mixing of the reference point and asymmetry measures that created the inconsistency. Recall that in Section 2, we compared two measures of asymmetry — the skew and difference in upper and lower cumulative probabilities, using the mean as our point of reference. But our analysis has shown that the mean is associated with the quadratic loss function and so neither of these measures may be fully consistent with it. Instead we should be using the measure of asymmetry appropriate for the quadratic loss function.(1) Using appropriate measures of asymmetry with different central projections should provide consistent measurement of the degree of pdf asymmetry. We demonstrate this empirically in Section 6.

Though we set out our example using a 'policymaker', we have already noted that it may be applied to anyone wishing to summarise the information in a probability distribution for any variable. But our emphasis from the start lay with the information that we can get from option prices on market views about future asset prices. And the set of probabilities that we extract from option prices are market probabilities rather than those of the policymaker.(2) This might beg the question: whose preferences should we be concerned about in choosing a loss function? Strictly speaking, if our aim is to summarise market views it should be those of the market. The views we are considering are aggregate market views — the result of many interactions of individual heterogeneous agents and we have no idea what might be a reasonable loss function. But we argue that we need not be so concerned with this point. What is important is that, in considering market views, we use asymmetry measures that are consistent with alternative central projections. The loss functions may be viewed simply as tools that allow us to identify these measures for commonly used central projections.

Until now our discussion has focused on 'how' we can measure pdf asymmetry. But before we can interpret this measure of asymmetry in terms of market views about where the risks lie, on balance, to asset prices in the future, we need to think about 'what' it is that we want to measure the asymmetry of. Our examples above used pdfs of the level of a random variable — the FTSE 100 — but is there a case to be made for looking at the pdf of future logarithmic changes in the level of financial variables instead? And does this affect our interpretation of asymmetry as an indicator of market views on the balance of risks to asset prices?

5 Asset price levels, logarithmic changes and option-implied pdf asymmetry

The shape of a pdf will depend on the units with which we choose to measure the variable; whether, for example, we look at levels of asset prices (or, equivalently, simple proportional changes in price levels) or logarithmic changes in asset prices.(3) But why might we choose to look at units such as logarithmic changes instead of asset price levels themselves?

When evaluating the performance of different investment assets — such as equities, bonds and futures contracts — logarithmic growth rates are often preferred to simple (proportional) changes in asset prices for a number of reasons:

- Asset prices cannot be negative, which means that the distribution of possible asset price levels should naturally be asymmetric. Looking at the logarithm of the underlying asset price may allow us to get around this because the logarithm of positive numbers does not have a lower bound at zero.

- In addition, for assets like exchange rates, logarithmic changes are not dependent on the way prices are quoted. That is, a given appreciation of sterling against the euro implies the same depreciation of the euro against sterling when changes are calculated in logarithmic terms. That is not true when calculated using levels.

- A further advantage of logarithms, when considering probability distributions, lies in the equivalence of pdfs in log levels and pdfs in log changes. That is, as the price today is known, the logarithmic change over some future horizon is

(1) However, we show in a later section that, empirically, the measure of asymmetry based on the quadratic loss function is strongly associated with the skew measure.
(2) As noted earlier, the risk-neutral nature of the option-implied pdfs means that the implied probabilities will reflect both market views on probabilities and compensation for risk. The latter factor means that the probabilities themselves are also likely to reflect the preferences of the person selecting the estimate.
(3) By changes in asset prices we mean the change in the asset price at some horizon relative to today's futures price for that horizon.

simply the logarithm of the price level in the future minus a constant.

- Finally, logarithmic changes (or growth rates) and their sum provide a better description of the actual change over a period than the sum of simple proportional changes.(1)

So how will pdfs in terms of logarithmic changes for asset prices differ from those in terms of price levels? Let us consider the special case where asset prices are assumed to be 'lognormally' distributed. What do we mean by this? A random variable, say an asset price, is by definition said to have a lognormal distribution if the logarithm of the asset price is normally distributed. And a lognormal distribution for the level of asset prices necessarily means that simple proportional changes of the asset price level will also follow a lognormal distribution. In contrast, the logarithmic price level or logarithmic price level changes would have a normal distribution.

Chart 5 illustrates the difference for two theoretical random variables – x_1 and x_2. It shows the frequency distributions for 30,000 random observations for x_1, drawn from a lognormal probability distribution. Taking the logarithm of the lognormal variable x_1, we obtain normally distributed random observations for x_2. We can see that x_2 takes both positive and negative values while x_1 observations are only positive.

Chart 5
Examples of lognormal and normal probability distributions

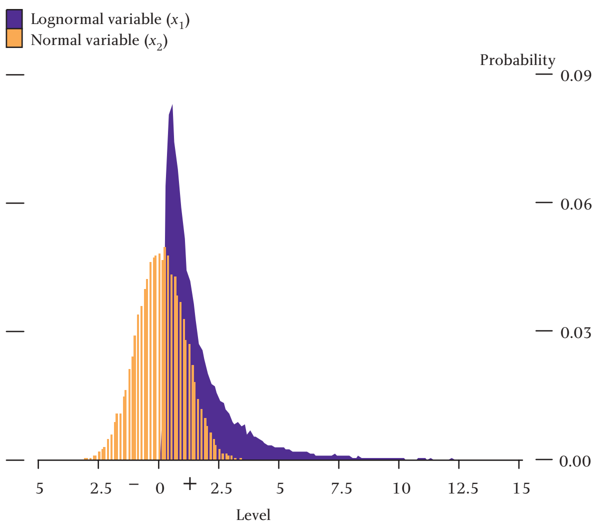

Furthermore, x_2 is symmetrically distributed, in contrast to the asymmetric distribution for x_1.

What significance does this have for pdf asymmetry? Suppose we are considering an option-implied pdf for asset price levels or changes. In terms of its shape, the lognormal distribution for levels or simple price changes would have a natural positive asymmetry under each of our asymmetry measures. But the pdf for logarithmic price levels or changes would be symmetric. In this sense, by looking at the asymmetry of the pdf for logarithmic changes, what we are really considering is the excess skew in asset prices — that is, how asymmetric the pdf for asset prices is relative to some 'natural' benchmark (which we take to be lognormal). A further illustration of this point in the context of option pricing and implied volatilities is provided in the box on page 450.

How realistic is this assumption? Much of the empirical finance literature has shown that probability distributions for historical logarithmic changes in asset prices, especially for equity indices, exhibit non-normal features. This is especially so for short-horizon changes such as those at the daily frequency. But we focus on pdfs for much longer horizon changes — those over three, six, nine and perhaps twelve months. For such horizons the evidence in the literature is less clear: empirically it is difficult to estimate reliably probability distributions for changes over these horizons due to insufficient numbers of independent past observations. However, at a theoretical level, the Central Limit Theorem is sometimes cited to reason that logarithmic changes at these horizons may be better approximated by a normal distribution than short-horizon changes.(2)

To illustrate the effect of using an asymmetry measure from the implied logarithmic changes pdf (or equivalently, the logarithmic level pdf) and the implied price level pdf to assess market views on the balance of risks to asset prices, Chart 6 shows time series of an asymmetry measure from each of the six months ahead implied pdfs for oil prices.

It is clear that the two series are highly correlated. The level difference between the two means that asymmetry (skew) in level space would imply a positive balance of risks most of the time. But this is not the case for

(1) Consider an asset price which changes from 100 to 150 in period 1 and back to 100 in period 2. The sum of proportional (arithmetic) changes is 0.50 – 0.33 = 0.17 while the sum of logarithmic (geometric) changes is 0.41 – 0.41 = 0. The sum of arithmetic changes is positive despite the price of the asset at the end of period 2 being the same as in period 1.
(2) See Campbell, Lo and MacKinlay (1997, page 19) for more details.

Option-implied pdfs, the Black-Scholes model and implied volatility smiles

The assumption that asset price levels (or simple proportional changes in asset prices) are lognormally distributed is frequently used in the pricing of option contracts. For example, the Black-Scholes (1973) model, a benchmark model for option pricing, is consistent with this assumption. One of the reasons it is a useful benchmark for option pricing is because the logarithmic growth rate of asset prices in the Black-Scholes model is normally distributed. In practice, however, the implied pdfs that we observe often deviate from that implied by the Black-Scholes model. Nonetheless, if we look at pdfs based on logarithmic changes then the model may still be useful as a benchmark for assessing asymmetry.

To see why this may be so, let us consider the volatilities derived from the prices of option contracts (often referred to as option premia). The Black-Scholes formula can be used to infer, from the option premium and other characteristics of the contract, the 'implied volatility' of the price of the underlying asset.(1) If the Black-Scholes model is correct then this implied volatility should provide a measure of the expected volatility of the underlying asset over the remaining life of the option contract. Plotting the Black-Scholes implied volatilities across different exercise prices is called the 'implied volatility smile'. The information that goes into a pdf is essentially the same as that on which the relevant volatility smile is based.(2) In fact it is the slope of the implied volatility smile that determines the shape, and hence the degree of asymmetry, of the implied pdf. A flat volatility smile is consistent with the assumptions of the Black-Scholes framework, and so is often used as a convenient benchmark for assessing deviations from the Black-Scholes implied distribution for logarithmic changes (ie the normal distribution).(3)

In practice, a flat volatility smile is rarely observed. Charts A and B provide an example of an (interpolated) volatility smile implied by FTSE 100 option contracts with December 2003 expiry (as of 3 November 2003), together with the corresponding implied probability density function for logarithmic changes. For comparison, the volatility smile and implied pdf under a Black-Scholes framework are also shown.(4) The observed smile for the FTSE 100 is downward sloping and thus deviates from the Black-Scholes flat volatility smile. That is, the implied volatility smile suggests that investors are paying higher premia for contracts with low FTSE 100 strike prices than suggested by the Black-Scholes model.(5) This is then reflected in the implied pdf with lower outcomes (ie more negative logarithmic changes) having relatively more probability than implied by the Black-Scholes normal pdf. Consequently the associated FTSE 100 implied pdf is not normal, in contrast to the Black-Scholes benchmark.

Chart A
FTSE 100 volatility smile and corresponding flat volatility smile

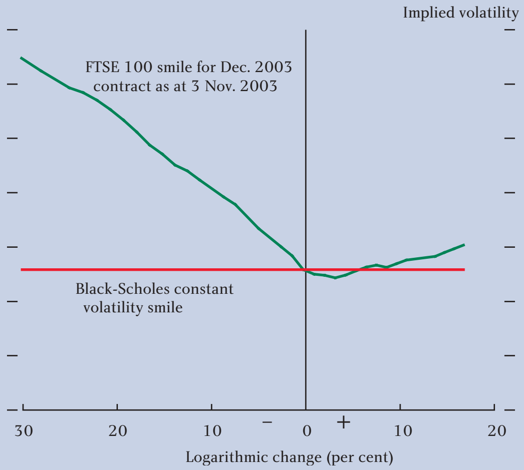

Chart B
FTSE 100 implied pdf and corresponding normal pdf

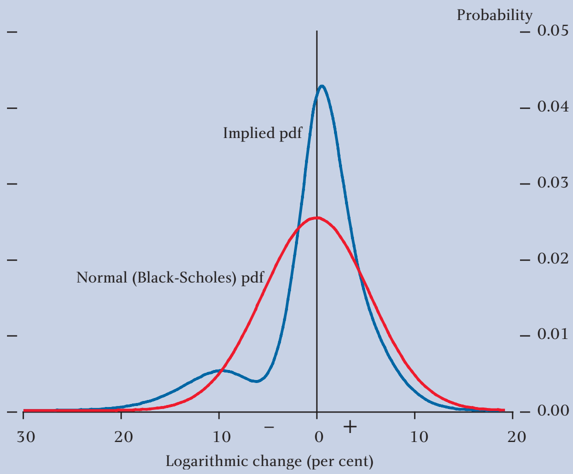

(1) The 'implied volatility' is the annualised standard deviation of logarithmic changes in the price of the underlying asset over the remaining life of the option contract.
(2) All of the techniques for extracting pdfs from option premia have as their input a set of option prices corresponding to different strike prices or equivalently a set of implied volatilities with corresponding strike prices. More information on the technique used to extract the implied pdfs in this article can be found in Clews, Panigirtzoglou and Proudman (2000).
(3) More specifically, an underlying asset stochastic process with constant volatility is consistent with the Black-Scholes framework.
(4) The normal pdf is fitted with the same mean and variance as that of the FTSE 100 implied pdf. Logarithmic change is with respect to the current futures price.
(5) There is a one-to-one positive relationship between option premia and implied volatility so one may think of implied volatility as a transformed premium.

Chart 6
Six months ahead option-implied oil price asymmetries(a)

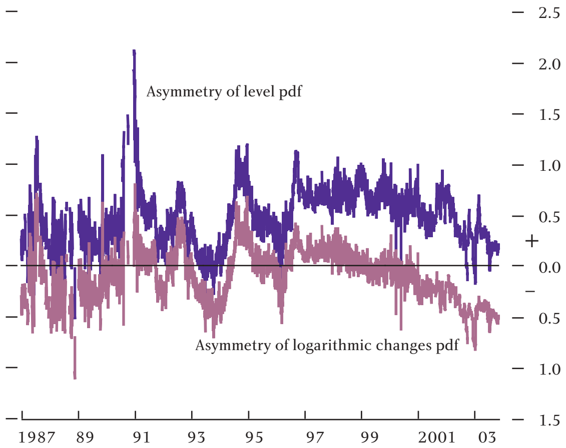

(a) These statistics are obtained from pdfs derived from option prices on the West Texas Intermediate light, sweet crude oil contract traded on the New York Mercantile Exchange.

logarithmic changes. Asymmetry on average is very close to zero. Here we take zero to represent the benchmark level of skew for both pdfs. However, if the level pdf is naturally asymmetric then zero would not be the appropriate benchmark level of skew. On the other hand, if the distribution of longer-horizon logarithmic changes is closer to a normal distribution then we may use zero as the benchmark for the corresponding pdf asymmetry series. So focusing on the asymmetry of the logarithmic changes pdf provides a more straightforward read on asymmetry and market views about the balance of risks.

6 Empirical comparisons of alternative asymmetry measures

Thus far we have set out and considered some theoretical aspects on appropriately measuring asymmetry of option-implied pdfs and on relating this measurement to market views. We now turn to some empirical analysis to support our reasoning, using option-implied pdfs that we estimate on a daily basis. In addition, we compare the measures of asymmetry recommended above with other often-used measures of asymmetry.

Appendix 2 shows the empirical associations between our measures of quadratic, absolute and indicator asymmetry from implied pdfs for logarithmic changes, over a large sample of option-implied probability distributions. Table A2.1 shows correlations between the measures, while Table A2.2 shows the percentage of days for which the measures have the same sign. These suggest that the absolute, indicator and quadratic asymmetry measures are all very highly correlated. In addition, the percentage of days when the quadratic and absolute asymmetry measures have the same sign is high (89%). However, the percentage of days for which the indicator asymmetry measure has the same sign as either the quadratic or absolute measure is somewhat lower — around 78%–80%.

It may be surprising that the three asymmetry measures do not have the same sign for an even higher proportion of our sample. One possible reason for this is measurement error. This is especially the case with mode-based statistics, as the mode is difficult to estimate accurately relative to the other points of central tendency. In addition, most of the observations where the three asymmetry series have different signs are where the pdf is nearly symmetric.

Overall, these findings suggest that the measures are fairly consistent in measuring the asymmetry of the implied pdf, but may provide different signs for asymmetry at times when measurement error is high and/or asymmetry is close to zero.

Finally, we compare our asymmetry measures with a number of other commonly used measures:

- The skew (third central moment standardised by the standard deviation) of the logarithmic changes implied pdf. This is the preferred measure of asymmetry in analyses of balances of risk for most asset prices at the Bank and is often reported in Bank publications.

- The difference between the mean and the mode and between the mean and the median of the logarithmic changes implied pdfs standardised by the corresponding standard pdf deviation.

- The risk reversal — or the difference between the costs of insurance against increases in the underlying asset price (beyond a certain level) and insurance against decreases.(1) The risk reversal is regularly traded and quoted by investment banks in the over-the-counter foreign exchange options

(1) So a positive (negative) risk reversal suggests market participants are paying more (less) for insurance against increases in the underlying asset price than against decreases and thus suggests that the balance of risks for the underlying asset price is positive. Formally, it is the difference between equally out-of-the-money (25-delta) call and put-implied volatilities and gives an idea of the slope of the implied volatility smile (see the box on page 450). We standardise it by dividing by the at-the-money implied volatility. As a benchmark, a lognormal distribution (which is positively skewed) has a risk reversal of zero.

market for example. As a result, it is the preferred measure of balances of risk for exchange rates in analysis at the Bank.

Each of these measures is very highly correlated with the three loss-based asymmetry measures. There are strong relationships between the sign of the quadratic loss asymmetry measure (which corresponds to the mean central projection) and both the skew and risk reversal; and between the mean minus the median and the absolute loss asymmetry (which corresponds to the median central projection).

This suggests that both the skew and risk reversal are reliable measures if one chooses the mean of the implied pdf as a central projection. In addition, the difference between the mean and the median is useful when the median is chosen as a central projection. In contrast, the lower same-sign percentage statistic between the mean minus the mode and the indicator loss asymmetry measure suggests that the mean minus the mode is a less reliable indicator of asymmetry when using the mode as a central projection.

7 Conclusions

The above analysis explores many of the issues involved in measuring and interpreting probability distribution asymmetry. It is worth emphasising that much of the analysis arises out of a need to summarise how the information in a probability distribution evolves over time. To do this, it is necessary to use a framework based on loss functions in order to ensure that our measures of asymmetry are consistent with other information that we can take from a probability distribution. But if we were not interested in summarising the information in the probability distribution, an analysis of the probability distribution would not need to involve consideration of loss functions.

The article has focused on the analysis of two summary measures of a probability distribution, the central reference point/point estimate and the asymmetry. These are both shown to depend on an assumed loss function and, as a result, matching the two is important. Taking three commonly used point estimates, we derive the corresponding measures of asymmetry as the difference between expected losses above and below a decision-maker's best point estimate. For the symmetric normal distribution, the asymmetry measures are all zero under the three loss functions we consider. Given this, the normal distribution is a commonly used benchmark when examining asymmetry.

Turning to the units in which we measure changes in asset prices, we show that this choice will affect the shape — and thus the degree of asymmetry — of a probability distribution. For example, under the popular metric of logarithmic changes in asset prices, the normal distribution benchmark coincides with that implied by the Black-Scholes option-pricing model.

Taking these general considerations into account, the article finally turns to the specific task of relating the information in probability distributions implied by option prices to market views about the asymmetry or balance of risks to future asset prices. Empirically, we found the loss-based asymmetry measures to be fairly consistent in measuring the asymmetry of option-implied pdfs, but they may provide different signs for asymmetry at times when measurement error is high or when asymmetry is close to zero. Other well known measures of asymmetry were found to be reliable indicators. That is, the risk reversal and implied pdf skew are useful when the mean is used as a point estimate, while the difference between the mean and the median is useful when the median is used as a point estimate.

Appendix 1: Loss functions and associated central projections and asymmetry measures

Table A1.1

Loss function	Central projection(1) $(\hat{x})$	Asymmetry
Indicator $L(x,\hat{x}) = (1 - I(x-\hat{x}))$	Mode: $x_{mode} = \max_{x} f(x)$	$Prob\ (x \geq x_{mode}) - Prob\ (x \leq x_{mode})$
Absolute $L(x,\hat{x}) = \lvert x-\hat{x} \rvert$	Median: x_{50} where $\int_{-\infty}^{x_{50}} f(x)dx = 0.5$	$\dfrac{Mean - Median}{\int_{-\infty}^{\infty} \lvert x - x_{50} \rvert f(x)dx}$
Quadratic $L(x,\hat{x}) = (x-\hat{x})^2$	Mean: $\bar{x} = \int_{-\infty}^{\infty} xf(x)dx$	$\dfrac{\int_{\bar{x}}^{\infty} (x-\bar{x})^2 f(x)dx - \int_{-\infty}^{\bar{x}} (x-\bar{x})^2 f(x)dx}{Variance}$

(1) Where *f(x)* refers to the probability distribution of *x*.

Appendix 2: Relations between alternative measures of asymmetry for option-implied pdfs for logarithmic changes in oil price(1)(2)

Table A2.1: Correlations

	Absolute asymmetry	Quadratic asymmetry	Indicator asymmetry	Risk reversal	Skew	Mean-mode	Mean-median
Absolute asymmetry	1.000						
Quadratic asymmetry	0.953	1.000					
Indicator asymmetry	0.813	0.848	1.000				
Risk reversal	0.952	0.999	0.849	1.000			
Skew	0.951	0.999	0.845	0.999	1.000		
Mean-mode	0.883	0.923	0.986	0.923	0.921	1.000	
Mean-median	1.000	0.952	0.814	0.951	0.951	0.885	1.000

Table A2.2: Percentage of observations with same sign

	Absolute asymmetry	Quadratic asymmetry	Indicator asymmetry	Risk reversal	Skew	Mean-mode	Mean-median
Absolute asymmetry	100.0						
Quadratic asymmetry	88.8	100.0					
Indicator asymmetry	79.6	78.3	100.0				
Risk reversal	88.3	97.5	76.5	100.0			
Skew	88.8	99.4	77.9	98.0	100.0		
Mean-mode	85.6	84.3	94.0	82.5	83.9	100.0	
Mean-median	100.0	88.8	79.6	88.3	88.8	85.6	100.0

(1) These statistics are obtained from pdfs derived from option prices on the West Texas Intermediate light, sweet crude oil contract from 1987–2000, traded on the New York Mercantile Exchange.
(2) See Section 6 for definition of risk reversal, mean-mode and mean-median asymmetry measures.

References

Bahra, B (1997), 'Implied risk-neutral probability density functions from option prices: theory and application', *Bank of England Working Paper no. 66*.

Black, F and Scholes, M (1973), 'The pricing of options and corporate liabilities', *Journal of Political Economy*, Vol. 81, pages 637–54.

Bliss, R B and Panigirtzoglou, N (2002), 'Testing the stability of implied probability density functions', *Journal of Banking and Finance*, Vol. 26, pages 381–422.

Bliss, R B and Panigirtzoglou, N (2004), 'Option-implied risk aversion estimates', *Journal of Finance*, Vol. 59, pages 407–46.

Campbell, J Y, Lo, A W and MacKinlay, A C (1997), *The econometrics of financial markets*, Princeton University Press.

Clews, R, Panigirtzoglou, N and Proudman, J (2000), 'Recent developments in extracting information from options markets', *Bank of England Quarterly Bulletin*, February, pages 50–60.

Joseph, V R (2004), 'Quality loss functions for nonnegative variables and their applications', *Journal of Quality Technology*, Vol. 36, pages 129–38.

Moessner, R (2001), 'Over the counter interest rate options', *Research Papers in Finance*, 1/2001, Centre for Central Banking Studies, Bank of England.

Svensson, L E (2004), 'Optimal policy with low-probability extreme events', *CEPR Discussion Papers*, No. 4218.

Vickers, J (1998), 'Inflation targeting in practice: the UK experience', *Bank of England Quarterly Bulletin*, November, pages 368–75.