



BANK OF ENGLAND

# Working Paper No. 455

## Estimating probability distributions of future asset prices: empirical transformations from option-implied risk-neutral to real-world density functions

Rupert de Vincent-Humphreys and Joseph Noss

June 2012

Working papers describe research in progress by the author(s) and are published to elicit comments and to further debate. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee.



BANK OF ENGLAND

# Working Paper No. 455

## Estimating probability distributions of future asset prices: empirical transformations from option-implied risk-neutral to real-world density functions

Rupert de Vincent-Humphreys<sup>(1)</sup> and Joseph Noss<sup>(2)</sup>

### Abstract

The prices of derivatives contracts can be used to estimate ‘risk-neutral’ probability density functions that give an indication of the weight investors place on different future prices of their underlying assets, were they risk-neutral. In the likely case that investors are risk-averse, this leads to differences between the risk-neutral probability density and the actual distribution of prices. But if this difference displays a systematic pattern over time, it may be exploited to transform the risk-neutral density into a ‘real-world’ density that better reflect agents’ actual expectations. This work offers a methodology for performing this transformation. The resulting real-world densities may better represent market participants’ views of future prices, and so offer an enhanced means of quantifying the uncertainty around financial variables. Comparison with their risk-neutral equivalents may also reveal new and useful information as to how attitudes towards risk are affecting pricing.

**Key words:** Asset prices, derivatives, expectations, options, option-implied density, risk premia, probability density forecasting, probability measure.

**JEL classification:** G10, G12, G13.

---

(1) Bank of England. Email: [rupert.devincent-humphreys@bankofengland.co.uk](mailto:rupert.devincent-humphreys@bankofengland.co.uk)

(2) Bank of England. Email: [joseph.noss@bankofengland.co.uk](mailto:joseph.noss@bankofengland.co.uk)

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We are grateful to Jens Larsen, Stephen Taylor, Olaf Weeken, Peter Westaway, Robert Woods and Chris Yeates for useful comments. This paper was finalised on 2 May 2012.

The Bank of England’s working paper series is externally refereed.

Information on the Bank’s working paper series can be found at [www.bankofengland.co.uk/publications/Pages/workingpapers/default.aspx](http://www.bankofengland.co.uk/publications/Pages/workingpapers/default.aspx)

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH  
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email [mapublications@bankofengland.co.uk](mailto:mapublications@bankofengland.co.uk)

## Contents

Summary	3
1 Introduction	5
1.1 Assessing systematic differences between distributions	6
1.2 Removing systematic differences between distributions	9
2 Methodological overview	11
2.1 General framework	11
2.2 The beta distribution as a calibration function	12
2.3 Risk-neutral density estimation	14
2.4 The data	16
2.5 Parameter estimation	19
3 Results	19
3.1 The estimated calibration functions	20
3.2 The estimated real-world densities	21
3.3 A measure of the risk premium	24
4 Robustness checking	25
4.1 A comparison to a mean-shift calibration	26
4.2 Evolution of estimated beta parameters	28
5 Conclusion	32
Appendix 1: The beta PDF and the transformations it affords	34
References	37



## Summary

There is a strong tradition of central banks and other policymakers extracting information from the prices of financial securities. Derivatives contracts can provide information on the expected future path of their underlying asset's price that goes beyond its central expectation. They therefore offer an insight into the level of uncertainty surrounding future cash flows. The Bank of England regularly estimates probability density functions (PDFs) from options prices in order to obtain an indication of the weight investors place on different future prices.

However, such option-implied PDFs may not provide a true indication of the actual probabilities investors ascribe to certain outcomes. This is because such PDFs give an indication only of the probabilities investors would have in mind if they were 'risk-neutral', and did not consider the uncertainty around an asset's future pay-offs in assessing its value. In the likely case that investors are averse to this risk, this would lead to differences between the risk-neutral densities backed out of options prices, and the true 'real-world' probability densities considered by investors.

The resulting estimated 'real-world' PDFs offer a number of advantages over their risk-neutral counterparts. First, they afford an insight into market participants' actual views on future asset prices, and offer an improved quantification of the uncertainty around financial variables. Second, a comparison of the risk-neutral and estimated real-world PDFs reveals new information as to how investors' risk preferences are affecting derivatives prices. Finally, estimated real-world probability densities are directly comparable with other forecasts considered by policymakers that are *not* based on derivatives prices, for example those of GDP growth and inflation.

The approach examined here is empirical in that it compares the risk-neutral distribution generated directly from options prices to the *actual* distribution of prices as they are later observed. To the extent that the two show a systematic disparity over time, this may be exploited to adjust the risk-neutral densities over as yet unobserved future prices to estimate the agents' real-world expectations.

This work offers a robust means of transforming risk-neutral densities obtained from options contracts on the FTSE 100 and short sterling. The resulting real-world probability densities offer a superior average fit across the distribution of observed prices than their risk-neutral

counterparts. The resulting parameters appear stable over time, at least until the end of our data sample in June 2007. To the extent that this remained the case when the methodology was applied to prices since, it could form the basis of an operational method to better predict their future prices and enhance conjunctural analysis. It could also form the basis of more advanced work that aimed to condition this transformation on some other (macroeconomic) observable variable which may increase the method's predictive power.



## 1 Introduction

Uncertainty pervades the economy and monetary policy making must take account of that. Financial asset prices may reveal useful information about likely future states of the economy, of which central banks have long made use. In particular, the probability density functions (PDFs) derived from options prices are a key tool for quantifying the amount and directional bias of uncertainty around future asset prices, and hence to the economic outlook.<sup>1</sup>

The Bank of England regularly estimates PDFs from option prices for a wide range of assets. However, option-implied PDFs may be a biased indicator of the actual probabilistic view held by market participants. This is because such PDFs correspond to the probabilities that an investor would have if they were risk-neutral, but the agents that price the options might in fact be risk-averse. If that were the case, risk premia would lead to differences in the location and shape of risk-neutral densities (RNDs) and real-world densities (RWDs).

Estimates of RWDs could enhance our analysis in two respects. First, such PDFs would better represent market participants' actual views on future asset prices, and for quantifying uncertainty around key financial variables. Second, and perhaps more importantly, a comparison of the risk-neutral and real-world PDFs could reveal new and useful information about how attitudes towards risk are affecting pricing. For instance, a term structure of the equity risk premium could be derived, on a daily basis, from equity index PDFs for different horizons.

Estimated RWDs may also provide a measure of risk that is directly comparable to that contained in the Monetary Policy Committee's GDP growth and inflation forecasts, making it a more relevant tool for quantifying uncertainty around key financial variables. Information on the evolution of agents' risk preferences could add insight to our understanding of their likely behaviour.

This work offers a means of estimating RWDs that can be applied to options of different underlying assets, and of different maturities, on a daily basis. Previous Bank of England research has examined methodologies based on utility-function transformations of the RNDs

---

<sup>1</sup> Breeden and Litzenberger (1978) observed that the risk-neutral probability density of an asset's future price is proportional to the second derivative of the price of options written upon it with respect to their strike prices. This has now become a standard technique adopted by central banks for extracting information on the future course of asset prices from options contracts. See Clews *et al* (2000) for the implementation used by the Bank of England.

(Bliss and Panigirtzoglou (2004)) and estimating the RWD directly with a threshold-GARCH simulation of the underlying price process (Gai and Vause (2005)).

This study takes an empirical, density-forecasting approach to estimating the RWDs. Probability densities that reflect rational agents' real-world expectations should, on average, be unbiased predictors of the eventual outcomes. If risk-neutral PDFs are biased, the difference between the RND and the estimated RWD of prices as they are later observed provides an indication of the degree and nature of risk-aversion of the representative agent. To the extent that this difference displays a systematic pattern over time, it may be exploited to adjust the RNDs over as yet unobserved future prices to estimate the agents' real-world expectations. This paper follows Fackler and King (1990) in using the highly flexible, yet parsimonious, beta distribution function to deliver that calibration.

This section illustrates how such systematic differences between risk-neutral and real-world distributions can be identified, and reviews the existing literature on how these can be removed. A methodological framework is introduced in Section 2. Section 3 applies this to three-month options on two important UK financial instruments: FTSE 100 and short sterling. The results of these transformations yield time-series estimates of the equity risk premium and the term premium at that maturity. Section 4 examines the robustness of these results, and the stability of the estimated parameters over time; this is important if these parameters are to be relied upon to forecast RWDs of future prices. A final section concludes.

### *1.1 Assessing systematic differences between distributions*

Much of the recent literature on density forecasting is based on the probability integral transform (PIT), a method of testing the hypothesis that a sample of data are drawn from a particular candidate distribution (Rosenblatt (1952)).<sup>2</sup> The PIT of the realisation  $x_t$  is defined as:

$$y_t = \int_{-\infty}^{x_t} f_t(u) du \quad (1)$$

where  $f_t$  is the candidate probability density function.

If observations  $\{x_t\}$  are independent and the candidate PDF is identical to the true PDF, then the PIT,  $y_t$ , will be independently and uniformly distributed. Therefore, testing whether the observed

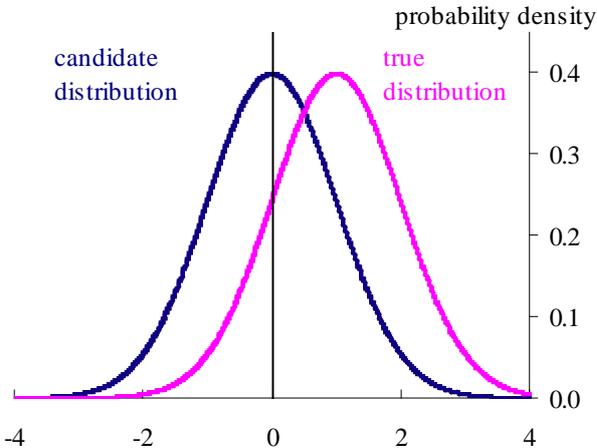
---

<sup>2</sup> This is also known as the inverse probability transform (IPT).

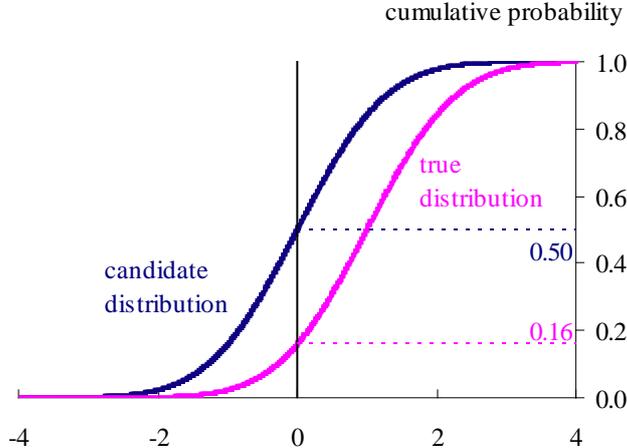
data was drawn from the candidate distribution is equivalent to testing whether the transformed data was drawn from the uniform distribution.

A quantile-quantile ( $q-q$ ) plot of the PIT provides an intuitive means of assessing visually any systematic difference between the candidate distribution and the true, unknown distribution from which the sample was drawn. Such a  $q-q$  plot shows the proportion of observations less than the  $n^{\text{th}}$  percentile of the candidate distribution.

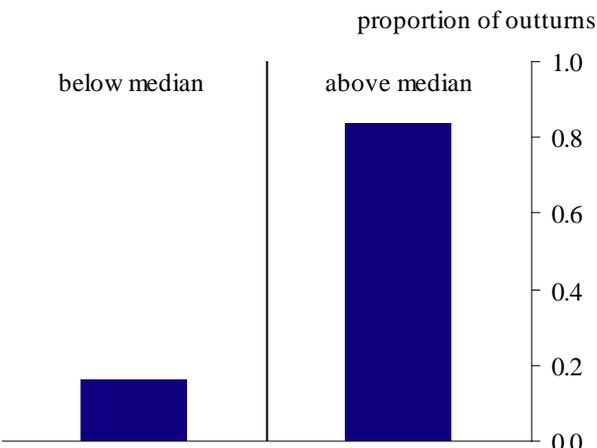
**Chart 1** If the candidate distribution understates the mean of the true distribution...



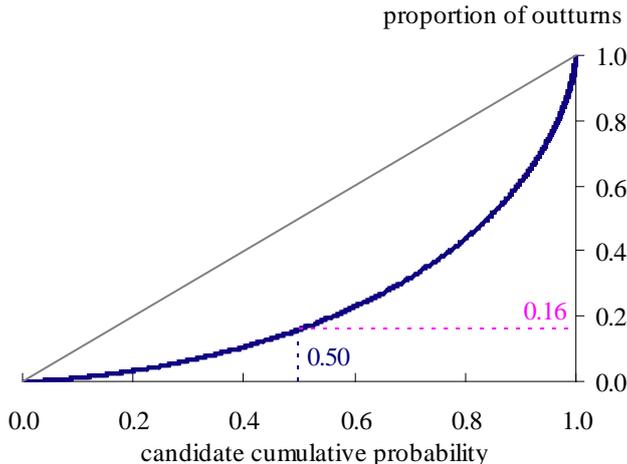
**Chart 2** ... ie the candidate distribution's PDF and CDF lie to the left...



**Chart 3** ... then there will be fewer outcomes below the median observed than expected ...



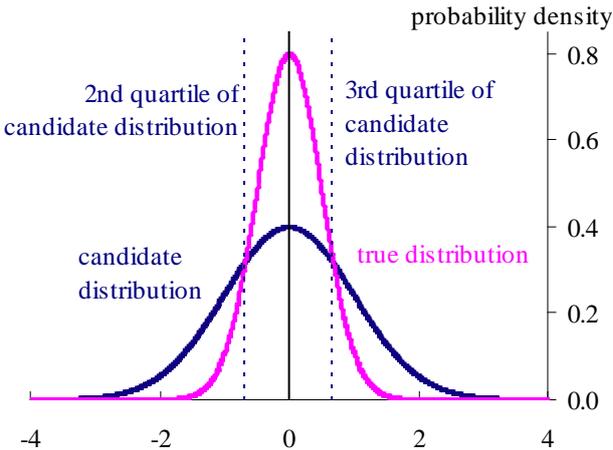
**Chart 4** ... the  $q-q$  plot of the PIT will bend around the straight line of gradient one.



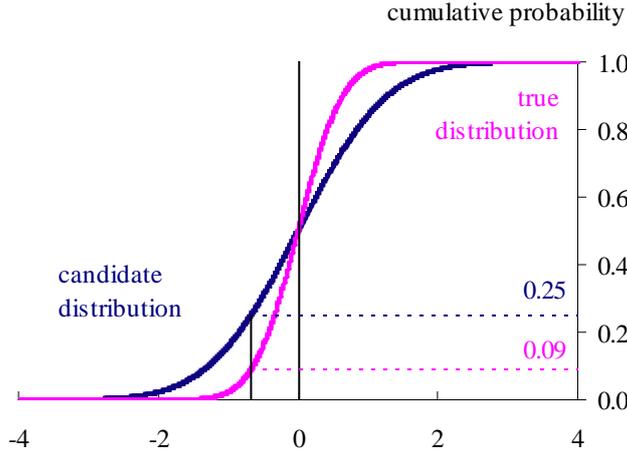
The key concept underlying a  $q-q$  plot is that, if the candidate distribution accurately describes the true distribution from which observations are drawn, then the proportion of observations less than the median of the candidate distribution should tend to one half, as the sample size approaches infinity. More generally, the proportion of out-turns less than the  $n^{\text{th}}$  percentile of the

candidate distribution should tend to  $n\%$ , as the same size increases. So if the candidate distribution and the true distribution coincide, then the  $q-q$  plot will be a straight line of gradient one. Alternatively, if the candidate distribution understates the mean of the true distribution, then the proportion of observations less than the  $n^{\text{th}}$  percentile will tend to a value less than  $n\%$ , so the  $q-q$  plot will sag below the gradient-one line. This is illustrated in Charts 1-4.

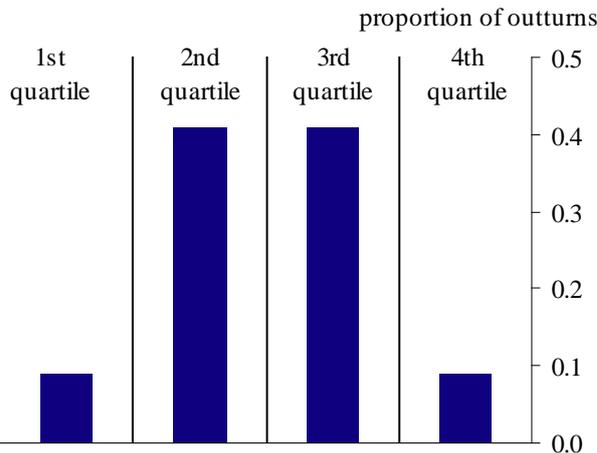
**Chart 5** If the candidate distribution overstates the variance of the true distribution...



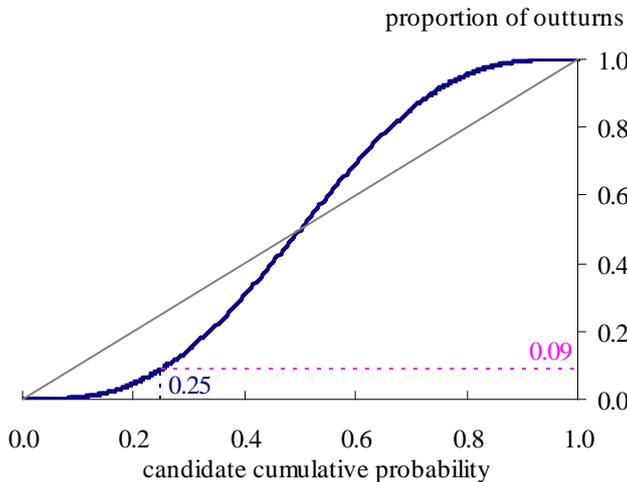
**Chart 6** ... ie the candidate distribution's PDF and CDF are wider...



**Chart 7** ... then there will be fewer outcomes in the tails observed than expected ...



**Chart 8** ... so the  $q-q$  plot of the PIT will bend around the straight line of gradient one.



If the candidate distribution overstates the variance of the true distribution, then there will be fewer extreme outcomes than the candidate distribution suggests. This means that the proportion of observations less than the  $n^{\text{th}}$  percentile will initially tend to a value less than  $n\%$ ; after the median it will tend to a value greater than  $n\%$ . So the cumulative PIT form an 'S'-shape, bending around the gradient-one line. This is illustrated in Charts 5-8.

Fackler and King (1990), which forms the foundation for this paper, use  $q-q$  plots to illustrate that option-implied RNDs provide a poor probabilistic description of agricultural commodity price out-turns. Liu *et al* (2007) also employ  $q-q$  plots to compare the PIT to a uniform distribution, which as previously explained is equivalent to checking whether the empirical distribution function of the PIT lies on the gradient-one line.

## 1.2 *Removing systematic differences between distributions*

There are two main strands to the literature on adjusting option-implied RNDs to produce an estimate of the RWD.

The first of these seeks to estimate the utility function agents use when assessing uncertain outcomes. The estimated parameters of this utility function determine investors' preferences towards risk. Bliss and Panigirtzoglou (2004) establish that option-implied RNDs do not provide good forecasts of future realisations of either the S&P 500 or the FTSE 100. They then examine RWDs that had been adjusted under the assumption of investors' preferences are represented by both power utility and exponential utility functions. The parameters of those utility functions were selected such that forecasting performance was optimised. These utility adjustments were successful in that they failed to reject the hypothesis that the adjusted RWDs were an unbiased predictor of future realisations.

Alonso *et al* (2006) apply the analysis of Bliss and Panigirtzoglou (2004) to the Spanish IBEX-35 index, and extend the analysis to include habit-based preferences and time-varying risk aversion. Anagnou-Basioudis *et al* (2005) study RNDs for the S&P 500 and the sterling-dollar exchange rate estimated using four different methodologies: generalised beta, normal inverse Gaussian, a combination of two lognormals and a smoothing-spline. They find that option-implied RNDs are both biased and inefficient predictors of realised distributions, and that the bias cannot be removed by a simple mean adjustment. Such a mean shift would correspond to a constant risk premium, and is the risk-neutral to real-world transformation of the Black-Scholes option-pricing paradigm. However, once the RNDs are adjusted using a power utility function they are unable to reject the assumption of no bias and efficiency, which corroborates the findings of Bliss and Panigirtzoglou (2004). In contrast, Weinberg (2001) fails to reject the hypothesis that a simple mean adjustment adequately removes the bias in the distributions formed from S&P 500 options.

The second strand of the literature also seeks to transform the observed RNDs into the RWD but does not require the function performing the calibration function to be a utility function purporting to represent investors' risk preferences. This approach is more empirical, in that it seeks the function that best fits the RNDs obtained from a given set of observed data, without insisting that this function fit an underlying model of agents' behaviour. That this approach is less ambitious in its aims seems sensible, given that we do not observe the states of the world that determine a given asset's price, only the price itself to one of many states of the world has given rise. A complete model of investors' risk preferences, which would be capable of pricing any asset, would require these individual states of the world to be observed, as it is levels of consumption in these different states of the world over which investors have preferences. It is therefore more realistic to attempt to find a means only of calibrating a RND to its real-world equivalent for a given asset class, rather than an all encompassing model of investor preferences.

Fackler and King (1990) find that the CDF of a beta distribution can be used to calibrate the RND to observed outcomes, removing the systematic difference between the two. They find the CDF of the beta distribution offers a parsimonious yet flexible calibration function. Although it depends on only two parameters, it nests many simple forms of transformation such as a mean shift, mean-preserving changes in variance, and changes involving mean, variance and skewness. A more detailed exposition of the various transformations that the beta CDF can deliver is presented in an appendix. Shackleton *et al* (2010) also use more sophisticated methods based on the model of price dynamics in Heston (1993) to obtain densities at multiple horizons. This includes the use of nonparametric calibration functions.

While the approach of Fackler and King (1990) does not aim to restrict itself to estimating transformations of the RND to the RWD that represent investors' preferences, Liu *et al* (2007) derive conditions under which the parameters of the beta CDF could represent the preferences of a risk-averse representative agent. They estimate the parameters of the calibration function and a power-utility function, based on a series of non-overlapping option contracts with one-month maturity. So the calibrations that result from the empirical approach may be compatible with those arising from a coherent model of investor preferences, even if the strength of the approach lies in how it is flexible enough to consider a broad class of calibration function that impose no such restriction.

## 2 Methodological overview

### 2.1 General framework

The calibration approach estimates real-world probabilities as some function of the risk-neutral probabilities. That function is called the calibration function. Equation (2) expresses this general relationship in terms of cumulative real-world and risk-neutral cumulative probability density functions  $F^P$  and  $F^Q$  (cumulative counterparts to the probability density function  $f$  in (1)):

$$F^P(x) = C(F^Q(x)). \quad (2)$$

Equation (2) shows that the choice of calibration function,  $C()$ , must be restricted to functions which are themselves cumulative probability functions, ie have a range  $[0,1]$  and are non-decreasing. It can be recast in terms of probability densities:

$$f^P(x) = C'(F^Q(x)) f^Q(x) \quad (3)$$

The multiplicative form of equation (3) is consistent with that of the fundamental asset pricing equation, which states that an asset's price at time  $t$ ,  $p_t$ , is equal to the expected product of its pay-off at time  $t+1$  across different states of the world,  $x_{t+1}$ , and stochastic discount factor  $m_{t+1}$ , a measure of investors' willingness to delay consumption:

$$p_t = E_t^P [m_{t+1} x_{t+1}]. \quad (4)$$

Using the definition of covariance, and noting that the expectation of the stochastic discount factor,  $m_{t+1}$  is the reciprocal of the risk-free rate, this can be immediately rewritten as,

$$p_t = \frac{1}{R^f} E_t^P [x_{t+1}] + \text{cov}(m_{t+1}, x_{t+1}) \quad (5)$$

so that the price of an asset is its expected pay-off, discounted at the risk-free rate, plus a risk premium coming from the covariance of the pay-off with the stochastic discount factor.

Alternatively it may be rearranged thus:

$$\begin{aligned} p_{t+1} &= \int_{\mathcal{S}} m_{t+1}(s) x_{t+1}(s) d\pi_{t+1}^P(s) = E_t^P [m_{t+1}(s)] \int_{\mathcal{S}} \frac{m_{t+1}(s) x_{t+1}(s) d\pi_{t+1}^P(s)}{E_t^P [m_{t+1}(s)]} \\ &= E_t^P [m_{t+1}(s)] \int_{\mathcal{S}} x_{t+1}(s) d\pi_{t+1}^Q(s) \quad \text{where} \quad \pi_{t+1}^Q(s) = \frac{m_{t+1}(s)}{\int_{\mathcal{S}} m_{t+1}(s) d\pi_{t+1}^P(s)} \\ &= \frac{1}{R^f} E_t^Q [m_{t+1}(s)] \end{aligned} \quad (6)$$

The function  $\pi_s^Q(s)$  defines a new probability measure: the *risk-neutral* measure. It is termed ‘risk-neutral’ because under this measure the market-clearing price of assets is such that an investor is indifferent between the asset and a certain payment equal to the asset’s expected pay-off, discounted with the risk-free interest rate. This defines a risk-neutral agent. Equation (7) links these actual (real-world) probabilities and their risk-neutral probabilities.

$$\pi_{t+1}^P(s) = \frac{E_t^P[m_{t+1}(s)]}{m_{t+1}(s)} \pi_{t+1}^Q(s) \quad (7)$$

This ‘state-dependent’ stochastic discount factor,  $m_{t+1,s}$ , allows for the pricing of any claim that delivers pay-off  $x_{t+1,s}$  at time,  $t+1$ , in each possible state of the world  $s$ .

In practice, we observe only the price of assets, rather than states of the world that give rise to them. This underlines an important limitation on the information we can hope to recover from those prices. Multiple states of the world may lead to identical asset prices. And two agents with similar yet different preferences might buy derivative contracts at the same prices, giving rise to identical calibration functions, if their preferences happen to be identical over the subset of states of the world we observe. It is therefore only possible to recover a calibration function, and corresponding estimated RWD, that is specific to a given asset, not a general calibration function that can be used on any asset.

Put in terms of equation (4), we are able to recover a stochastic discount factor,  $m_{t+1,x}$  that is the discount factor expected across all states of the world that lead to that observed price; that is

$$p_t = E_t^P[m_{t+1,x} x_{t+1}] \text{ where } m_{t+1,x} = E[m_{t+1,s} | x_{t+1,s} = p_{t+1}].$$

This stochastic discount factor therefore contains less information than the general version in (4); in particular it does not represent the risk preferences of any particular observable agent or investor in the underlying asset.

## 2.2 The beta distribution as a calibration function

We follow Fackler and King (1990) and impose that the calibration function is a beta distribution. This has a number of distinct advantages. Since the beta distribution nests the uniform distribution, it allows for the risk-neutral and real-world measures to be identical without imposing any transformation. This would correspond to the testable restriction that  $j = k = 1$ , in the case that the RNDs genuinely describe the real-world. In addition, while being parsimonious – there are only two parameters to be estimated – the beta functional form nests

within it a range of simple transformations, such as shifts in the mean of the distribution, mean-preserving changes in its variance, and adjustments to its skew. Of course, with only two degrees of freedom, the transformation is unable to affect all three of these moments independently; see Appendix 1 for an assessment of how changes in these three moments are related.

The standard beta PDF, given in equation (8), is defined over the finite interval [0,1]; the parameters  $j$  and  $k$  are the shape parameters.

$$f(z|j,k) = \frac{(z)^{j-1} (1-z)^{k-1}}{B(j,k)} \quad j, k > 0 \quad (8)$$

Because any CDF is, by definition, defined over the bounded interval [0, 1], the use of the standard beta PDF has become popular in the literature.<sup>3</sup> In the case when  $z = F_{t,\tau}^Q(x)$ , the option-implied risk-neutral cumulative probability, at time  $t$  for options expiring at time  $\tau$ , the standard beta PDF and CDF are given by equations (9) and (10).

$$f^B(F_{t,\tau}^Q(x)|j,k) = \frac{(F_{t,\tau}^Q(x))^{j-1} (1-F_{t,\tau}^Q(x))^{k-1}}{B(j,k)} \quad j, k > 0 \quad (9)$$

$$F(F_{t,\tau}^Q(x)|j,k) = \frac{B(F_{t,\tau}^Q(x)|j,k)}{B(j,k)} \quad (10)$$

The term  $B(j, k)$  is the beta function, given in equation (11); the incomplete beta function is given in equation (12)

$$\begin{aligned} B(j,k) &= \int_0^1 u^{j-1} (1-u)^{k-1} du & (11) \\ &= \frac{\Gamma(j)\Gamma(k)}{\Gamma(j+k)} \end{aligned}$$

$$B(F_{t,\tau}^Q(x)|j,k) = \int_0^{F_{t,\tau}^Q(x)} u^{j-1} (1-u)^{k-1} du \quad (12)$$

Combining the form of the calibration function given in equation (2) with equation (9) the real-world PDF becomes:

$$f_{t,\tau}^P(x) = f_{t,\tau}^B(F_{t,\tau}^Q(x)|j,k) f_{t,\tau}^Q(x) \quad (13)$$

where  $f_{t,\tau}^B(x|j,k)$  is the PDF of the beta distribution.

---

<sup>3</sup> See Fackler and King (1990) and Liu *et al* (2007).

Estimating the RWDs from the set of historical RNDs and price out-turns therefore amounts to estimating the parameters  $j$  and  $k$  of the beta distribution.

### 2.3 Risk-neutral density estimation

In this paper, the RNDs are estimated using a refinement of the non-parametric method set out in Bliss and Panigirtzoglou (2004).<sup>4</sup> That method exploits the result of Breeden and Litzenberger (1978), which relates the RND to the curvature of the call price function, that is

$$f_{t,\tau}^Q(K) = e^{-r_f \tau} \frac{\partial^2 C(S_t, K, T, t)}{\partial K^2}. \quad (14)$$

Where:  $S$  is the price of the asset,  $K$  is the strike price of the option,  $\tau = T - t$  is the option's time to maturity and  $r_f$  is the 'risk-free' rate.<sup>5</sup>

However option prices, and therefore the call price function, are observed only for a discrete set of strike prices. So in order to use this result, those discrete points must first be used to generate a call price function that is a twice-differentiable function of the strike price.

The procedure for estimating RNDs consists of five stages. First, the observed option prices are filtered to remove those prices that violate the no-arbitrage restrictions of monotonicity and convexity (Chart 9).<sup>6</sup> The data are then translated from (*strike, price*) into (*delta, implied volatility*) using the Black-Scholes formulae, as in Malz (1997). This does not assume that the Black-Scholes option-pricing paradigm holds true. Rather, this transformation should be viewed as purely for numerical convenience, which facilitates the third stage: fitting a natural smoothing spline through the transformed data (Chart 10). The parameter which controls the trade-off between smoothness and goodness-of-fit was set to 0.98. The spline is evaluated at 1,000 delta values and transformed back into (*strike, call price*) space before a numerical algorithm is used to compute the second derivative (Charts 11 and 12).

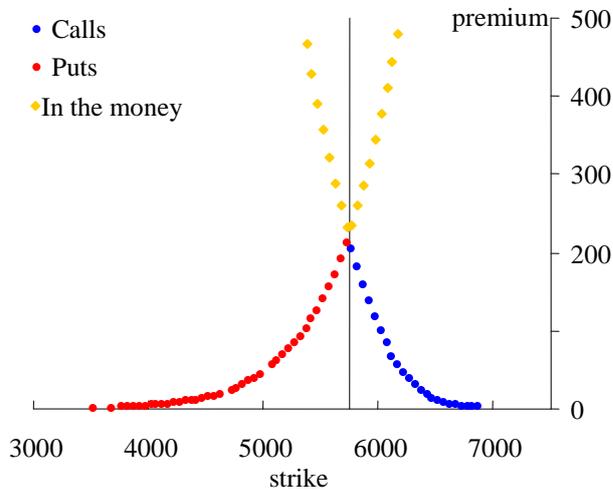
---

<sup>4</sup> This has become the standard technique for estimating RNDs adopted by other central banks (see de Vincent-Humphreys and Puigvert (2012)). For further details see Clews *et al* (2000).

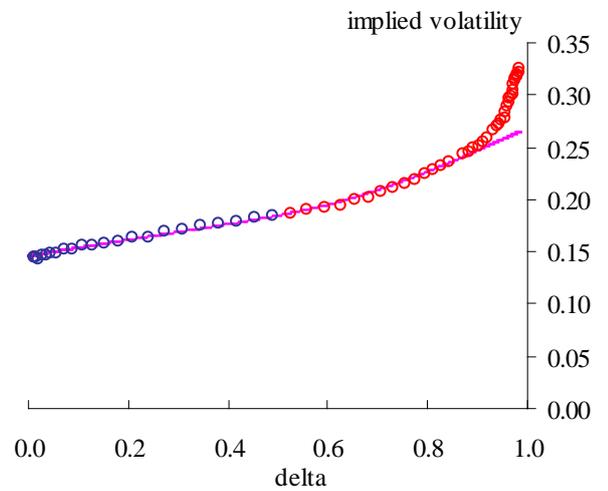
<sup>5</sup> We use three-month Libor as the risk-free rate for FTSE PDFs, for short-sterling PDFs we assume  $r_f = 0$ .

<sup>6</sup> Violation of those no-arbitrage restrictions in settlement prices does not imply a riskless profit opportunity, once bid-ask spreads are taken into account. It would, however, lead to negative probability densities.

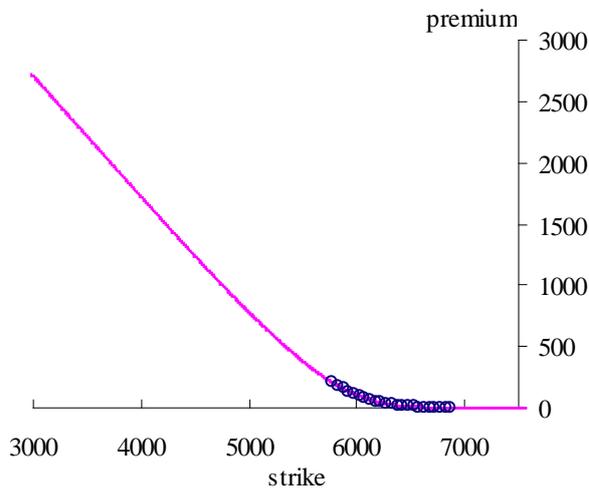
**Chart 9: FTSE option prices on 15 June 2001, expiry Sep 2001. No prices violate monotonicity or convexity.**



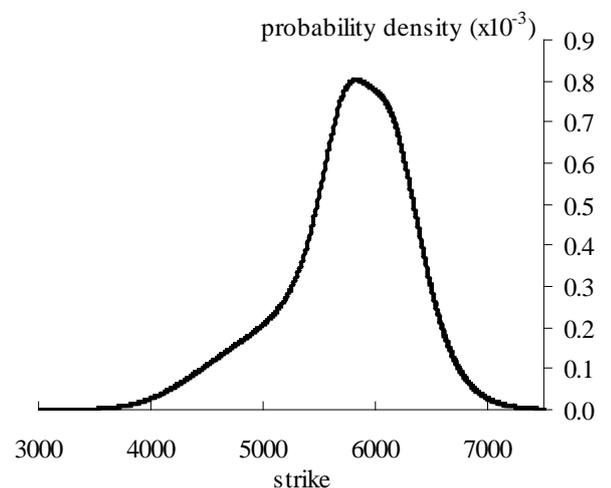
**Chart 10: Volatility smile for the option prices in Chart 9. A natural smoothing spline has been fitted through the 63 observed prices.**



**Chart 11: Estimated call price function: the transforming of the volatility smile of Chart 10 back into (strike, price) space**



**Chart 12: Estimated FTSE RND, the second derivative of the call price function in Chart 11**



## 2.4 *The data*

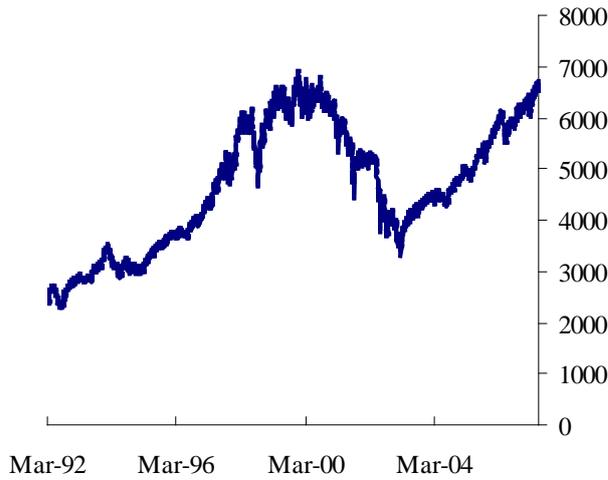
### 2.4.1 *FTSE 100 Index*

Options on the FTSE 100 equity index are traded on the London International Financial Futures and Options Exchange (LIFFE) with expiry dates of the third Friday of March, June, September and December of each year, plus additional months such that the nearest four calendar months are always available for trading. RNDs were estimated for each expiry date, using data on options expiring at the following expiry date. At this three-month horizon, the FTSE data set comprises 60 non-overlapping observations spanning the period March 1992 – June 2007. We leave as further work an extension up to the present day; for further details see Section 5.

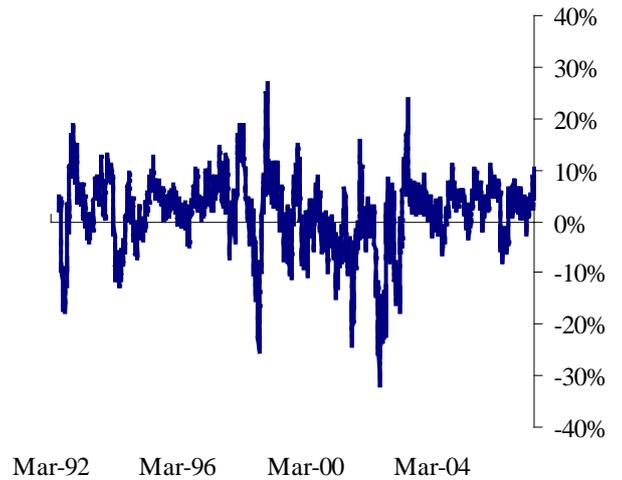
Chart 13 shows the evolution of the FTSE 100 over the sample period; Chart 14 shows the corresponding three-month log returns. As an example, Chart 15 shows the RND estimated from option settlement prices on 15 June 2001, for out-turns on 21 September 2001; the value of the underlying index at the time of the option's expiry is also indicated. Over that time, the FTSE 100 fell from 5723 to 4434. Chart 16 recasts the data in terms of cumulative probability, according to the option-implied distribution estimated from options three months previously. This time series passes a test for serial independence, with auto correlations that do not deviate from zero, at a 95% significance level (Chart 17).

Chart 18 plots the corresponding  $q$ - $q$  plot, which shows that over this sample the option-implied PDFs overstated the probability of out-turns below the median, and understated the probability of out-turns above. This, like the stylised example shown in Charts 1-4, is consistent with the candidate distribution understating the mean of the true distribution, and with investors requiring a positive premium to compensate them for bearing the risk of holding equities.

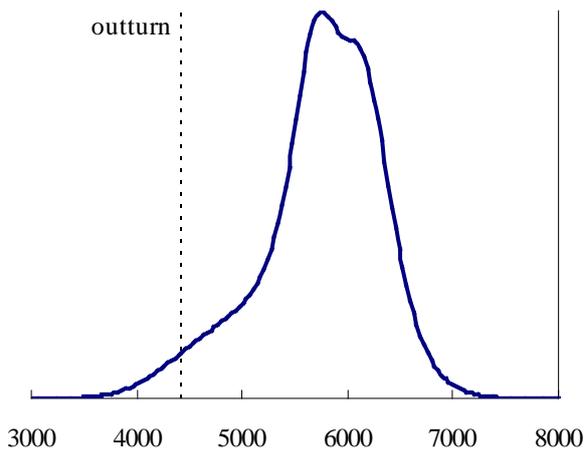
**Chart 13: FTSE 100 price history**



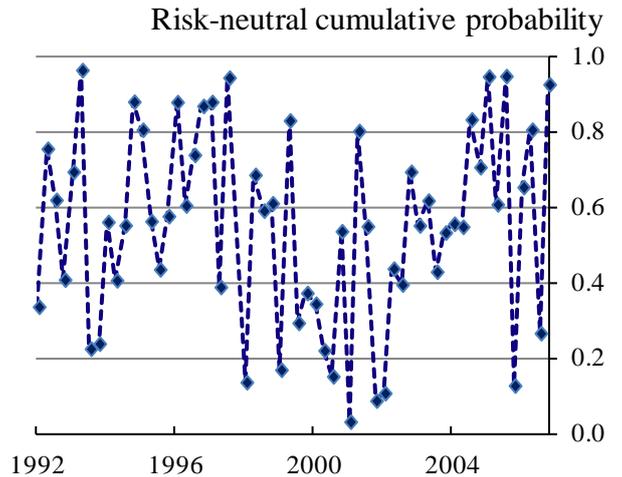
**Chart 14: FTSE three-month log-returns**



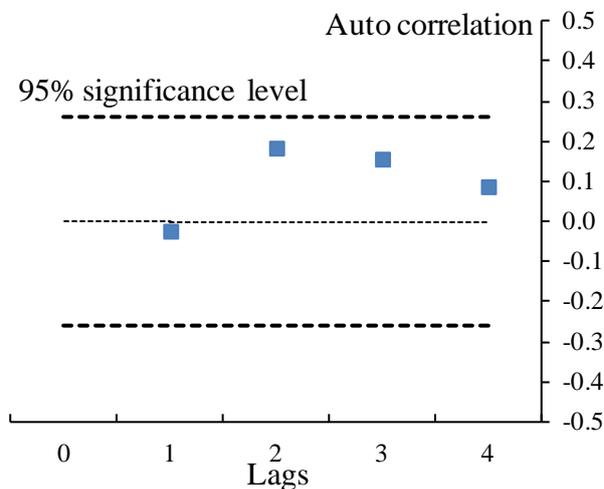
**Chart 15: FTSE RND estimated using 15 June 01 option prices for 21 Sep 2001 out-turns**



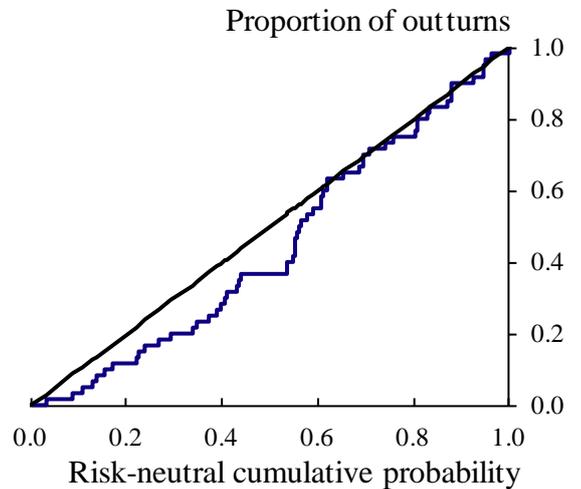
**Chart 16: Cumulative risk-neutral probabilities of FTSE out-turns, according to the RND estimated three months previously**



**Chart 17: Auto correlation of cumulative risk-neutral probabilities of FTSE out-turns**



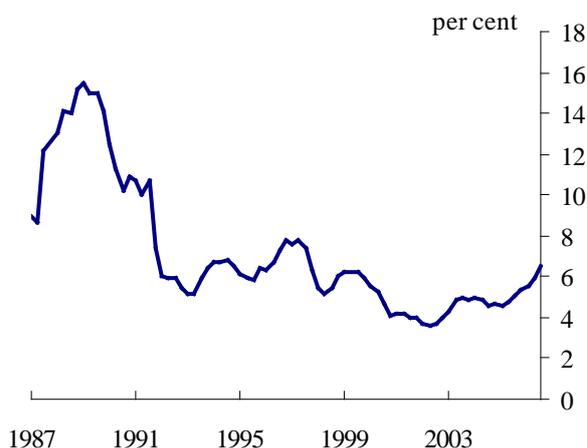
**Chart 18: *q-q* plot of three-month FTSE out-turns Mar 1992 – Jun 2007**



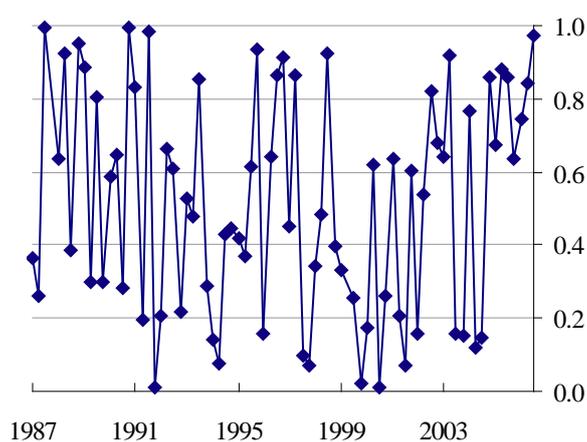
## 2.4.2 Short sterling

Options on short sterling futures contracts are also traded on LIFFE, with the same quarterly expiry schedule as FTSE options. The short sterling data set comprises of 80 observations spanning the period 1987–2007. Out-turns of the three-month Libor over this period are shown in Chart 19, and in terms of their risk-neutral cumulative probability, in Chart 20. The short sterling  $q$ - $q$  plot, presented in Chart 21, lies closer to the gradient-one line, suggesting that for this asset and maturity, option-implied RNDs may be a better representation of the true distribution prior to transformation, than those of the FTSE 100.

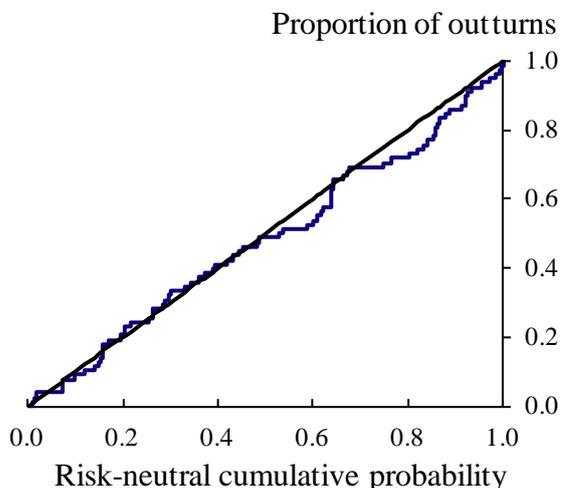
**Chart 19: Three-month short sterling spread history**



**Chart 20: Cumulative probabilities of Libor out-turns, according to the RND estimated three months previously**



**Chart 21:  $q$ - $q$  plot of three-month Libor out-turns Mar 1992 – Jun 2007**



## 2.5 Parameter estimation

Recasting (13) in terms of cumulative density functions, the real-world CDF,  $F_{t,\tau}^P(x)$  becomes

$$F_{t,\tau}^P(x) = F_{t,\tau}^B(F_{t,\tau}^Q(x) | j, k) \quad (15)$$

where  $F^B$  is the CDF of the beta distribution. If this candidate distribution function accurately describes the distribution from which the data,  $x_i$ , are drawn, it follows that:

$$F_{t,\tau}^P(x_i) = F_{t,\tau}^B(F_{t,\tau}^Q(x_i) | j, k) \sim U(0,1) \quad (16)$$

so that  $F_{t,\tau}^Q(x_i)$  is beta-distributed with parameters  $j$  and  $k$ . Our problem therefore reduces to estimating the parameters  $j$  and  $k$  that give rise to a beta distribution that ‘best fits’ the cumulative risk-neutral probabilities.

These parameters are estimated *via* the maximum likelihood method applied to the available sample of data. This selects the parameter values that produce the beta distribution that is ‘most likely’ to have resulted in the observed data.<sup>7</sup> The parameters that lead to the highest mean value of the beta PDF when it is taken over the observed data; that is:

$$(j, k) = \arg \max \log L(j, k; Y) = \sum_{i=1}^N \log L(j, k; Y_i) = \sum_{i=1}^N \log \frac{Y_i^{j-1} (1 - Y_i)^{k-1}}{B(j, k)}, \quad (17)$$

where the beta function,  $B(j, k)$ , is given by  $B(j, k) = \int_0^1 u^{j-1} (1 - u)^{k-1} du$ . This optimisation is

achieved by a numerical procedure that tries alternative parameter values, stopping when a maximum is found. The CDF of the beta distribution that results is the beta distribution function that, on average, best transforms the CDF of the risk-neutral distribution into that observed empirically.

## 3 Results

Data are sampled at a frequency of three months, equal to the maturity of the options. The period between each observed option price and the observed price of the underlying asset therefore do not overlap, so that observations are independent of each other. Results this way are better suited to hypothesis testing (see Section 4.1), as they are free of serial correlation that would be induced by overlapping observations, as would be the case if prices of options with a year to maturity were sampled at, eg, a quarterly frequency.

<sup>7</sup> For further background, and a sample application of this technique, see Liu *et al* (2007).

Using these non-overlapping observations does however reduce the number of available data points. This reduces the efficiency of the parameter estimates obtained *via* maximum likelihood estimation,<sup>8</sup> and decreases the power of the statistical tests we apply to assess how well our calibrated distribution fits that of the data. The 20 years of data available for the two markets considered here, options on the FTSE 100 and short sterling, mean that we are left with over 60 non-overlapping observations. However were this approach to be applied to some assets for which the options market is relatively nascent, for example, foreign exchange rates, this reduced sample size could be more problematic. If the methodology described here were being used operationally, overlapping observations could be used in order to maximise the available data.

### *3.1 The estimated calibration functions*

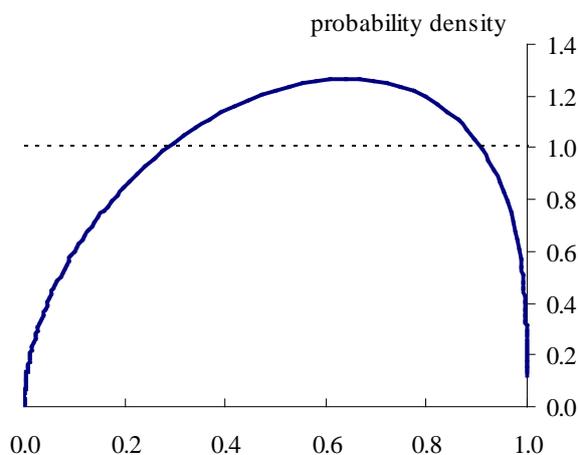
The estimated calibration function is the CDF of a beta distribution, estimated by finding the parameters of beta distribution that transforms the CDF of the option-implied distribution into the CDF of the distribution that maximises the likelihood of the realised returns. It is these transformed distributions that we term the estimated real-world densities (RWDs).

For options on the FTSE 100 with a three-month maturity, the estimated beta parameters were  $j = 1.56$  and  $k = 1.31$ . A plot of the corresponding density function (Chart 22) illustrates exactly how this calibration function affects the option-implied RNDs. Chart 22 shows that the beta PDF is asymmetrical, and takes values greater than one between 0.3 and 0.9. That means that the actual (ie real-world) probability of price out-turns with an option-implied cumulative probability between 0.3 and 0.9 is understated by the option-implied RND. Consequently, the calibration process augments such probabilities, increasing the mean of the candidate distribution (as in Charts 1-4). Because all probability densities must integrate to one, the calibration process reduces the option-implied probability of price out-turns outside that range. Chart 23 adds the estimated beta CDF to the FTSE 100  $q$ - $q$  plot. Another way of viewing the calibration function estimation is finding the line of best fit – out of the class of beta functions – through the empirical  $q$ - $q$  plot.

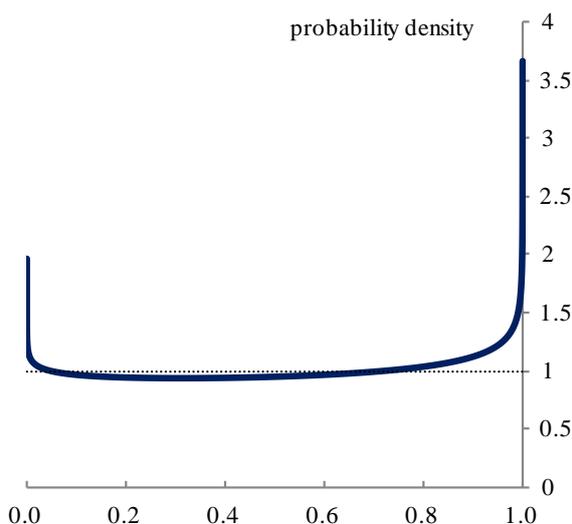
---

<sup>8</sup> It is a well-known property of maximum likelihood estimators that they converge in probability to their true value as the number of observations increases.

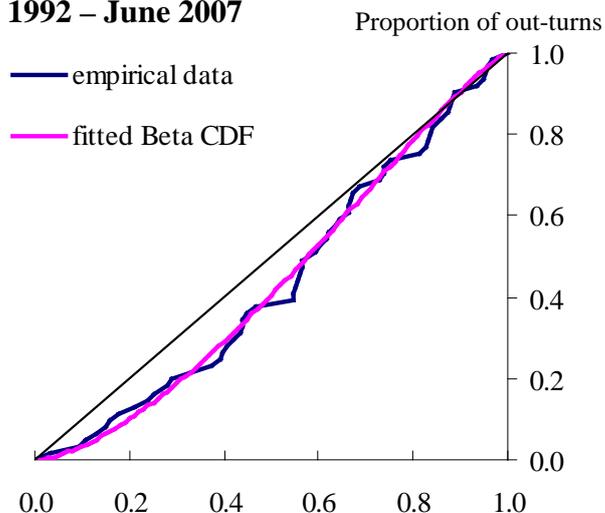
**Chart 22: Estimated calibration function for three-month FTSE (a beta PDF,  $j=1.56$  and  $k=1.31$ )**



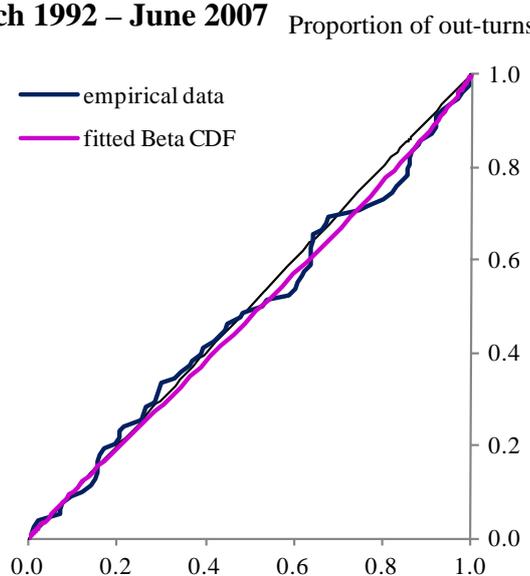
**Chart 24 – Estimated calibration function for three-month short sterling (a beta PDF,  $j=0.05$  and  $k=0.08$ )**



**Chart 23: Estimated beta CDF overlaid over  $q-q$  plot of three-month FTSE out-turns Mar 1992 – June 2007**



**Chart 25 – Estimated beta CDF overlaid over  $q-q$  plot of three-month Libor out-turns March 1992 – June 2007**



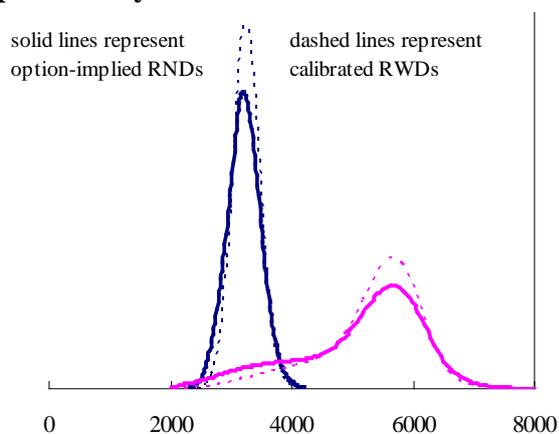
For three-month short sterling, the estimated beta parameters were  $j = 0.05$  and  $k = 0.08$ . The calibration function is plotted in Chart 24. This contrasts to that obtained for the FTSE 100 in that it is close to one over the majority of its domain, broadly symmetrical, and has spikes in its tails. This suggests that the risk-neutral density does not differ significantly from the estimated real-world density around the median of the distribution; but that it underestimates the actual probability of outcomes in its tails. The effect of the transformation of the short sterling RND is therefore concentrated at the tails of the distribution, whereas that for the FTSE 100 is focused more at its central moments, shifting its mean and, to some extent, its variance. The corresponding beta CDF is overlaid on the  $q-q$  plot in Chart 25.

### 3.2 The estimated real-world densities

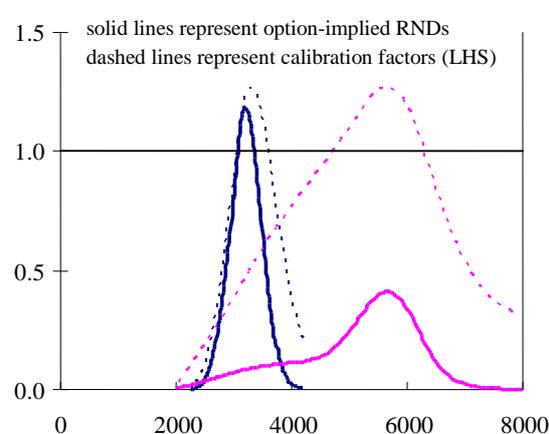
The RWDs are generated through the point-wise multiplication of the calibration function and the option-implied RND (equation 13). The RNDs and RWDs are compared in solid and dashed lines in Chart 26.

While the unconditional estimation of  $(j,k)$  means that the beta PDF is constant across the entirety of the sample data on which it is calibrated, the *function* performing the transformation of each RND will vary within this. This is because when used as a calibration function, the argument of the beta PDF is the cumulative risk-neutral probability, and this will vary with the shape of the option-implied PDFs. This is illustrated in Chart 27, which plots the two radically different RNDs in Chart 26 (solid lines), alongside their associated calibration functions (dashed lines). A given set of parameters for an asset class can therefore give rise to different calibration functions, dependent on the shape of the RND.

**Chart 26: Risk-neutral and real-world probability densities**

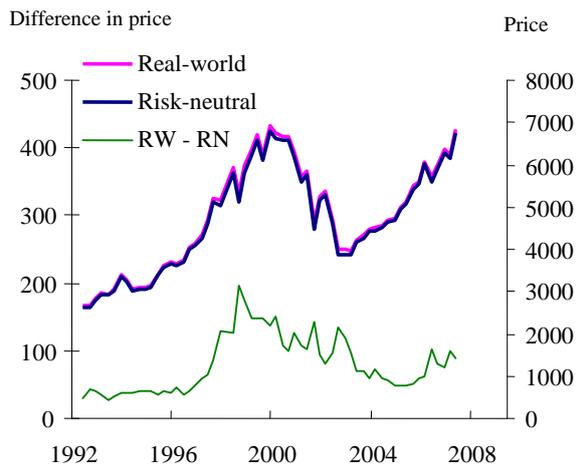


**Chart 27: Two risk-neutral densities and their calibration factors**

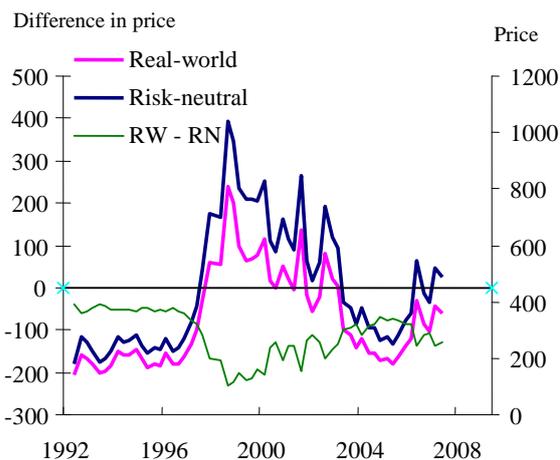


Once the set of RWDs have been estimated, the moments of those distributions can be computed numerically. Charts 28-31 compare the means and variances of the calibrated RWDs and the option-implied RNDs. Note that in Chart 28, the means of the RN and RW distributions for short sterling are close to indistinguishable over much of the data, as the transformation leaves the mean close to unchanged. Their sharp widening on 16 September 1992 (Black Wednesday) likely reflects the speed with which risk premia were changing, and so risk-neutral prices derived from the prices of options at their inception, differ more markedly from the actual out-turn.

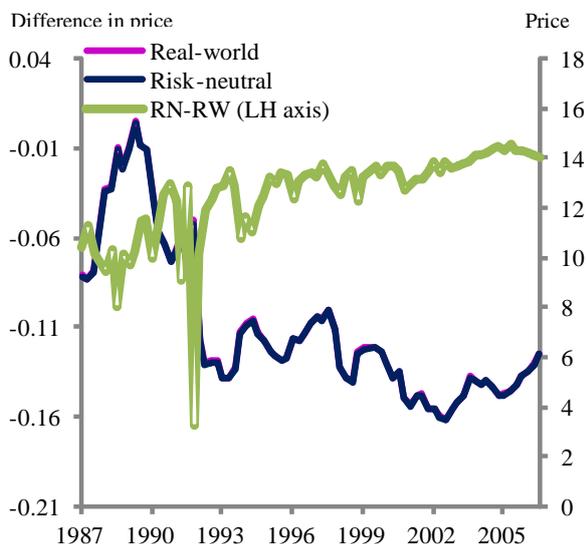
**Chart 28: Time-series mean of RNDs and RWDs for three-month FTSE**



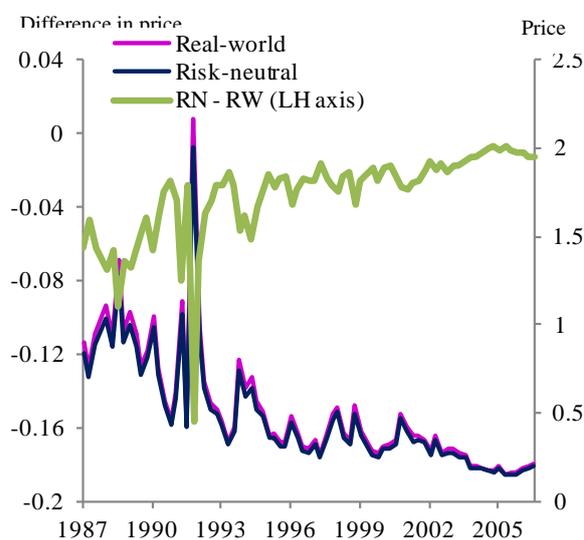
**Chart 29: Time-series variance of RNDs and RWDs for three-month FTSE**



**Chart 30: Time-series mean of RNDs and RWDs for three-month short sterling**



**Chart 31: Time-series variance of RNDs and RWDs for three-month short sterling**



Since this spike in 1992 however, the risk premia resulting from the short sterling contracts appears more stable than that for the FTSE, and this stability appears to have grown over time up to 2007. This difference is perhaps surprising, given that you might expect some drivers of risk premia to be common to both markets. This greater stability of risk premia derived from short sterling may, however, reflect the greater uncertainty around equity prices than those of short-horizon fixed-income instruments.

### 3.3 A measure of the risk premium

The means of the RNDs and RWDs can be combined to provide a measure of the risk premium embedded in the price of a futures contract. The current price of a security can be written as the discounted expectation of a future spot price, under different probability measures, provided that each expectation is discounted at the appropriate rate. The risk-neutral measure must be discounted using the risk-free rate, and the real-world measure discounted at a premium over the risk-free rate:

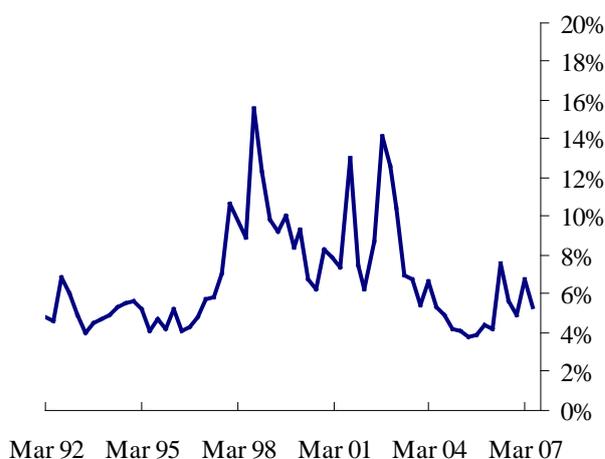
$$S_0 = e^{-rT} E^Q[S_T] = e^{-(r+\mu)T} E^P[S_T]. \quad (18)$$

That risk premium can therefore be recovered once the expectations (means) of the two measures have been estimated:

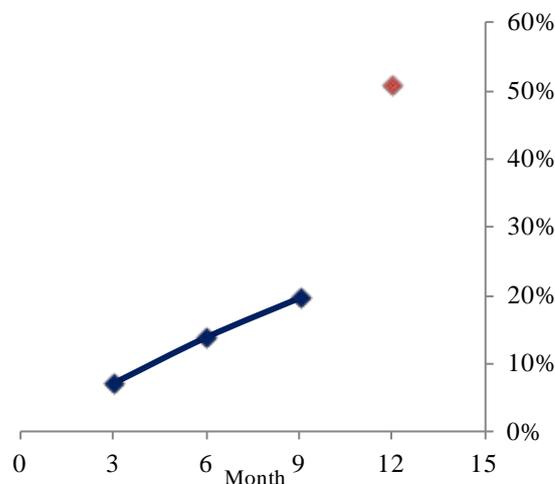
$$\mu = \frac{1}{T} \ln \left( \frac{E^P[S_T]}{E^Q[S_T]} \right). \quad (19)$$

Chart 32 plots this measure of the risk premium, estimated from three-month FTSE PDFs. This measure pertains to the remaining lifetime of the option, so because options of different maturities trade simultaneously, this method can be used to produce a *term structure* of the equity risk premium Chart 33. This slopes upwards, reflecting the higher degree of uncertainty over prices further into the future, and hence the higher risk premium demanded by investors. The risk premium derived from the twelve-month contract is absurdly high; however we are cautious of reading too much into this result as it is derived from only three contracts that are available at this long-maturity and so may be spurious.

**Chart 32: Three-month equity risk premium estimate for the FTSE 100**

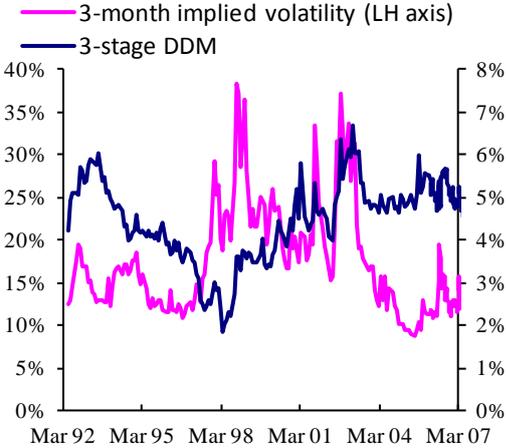


**Chart 33: Term structure of equity risk premia for the FTSE 100 (19 Mar 2004)**

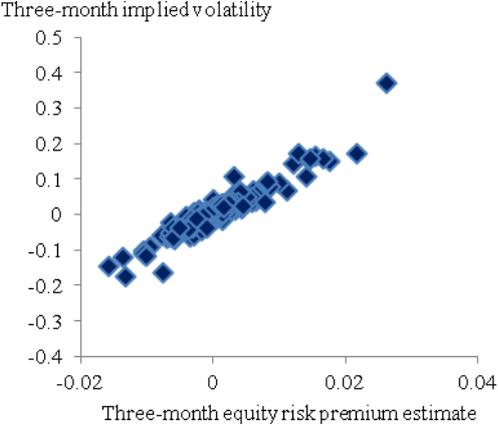


Two other indicators of an equity risk premium are the estimates available from a dividend discount model (DDM)<sup>9</sup> and FTSE at-the-money implied volatility. Chart 34 compares these two indicators to the measure derived from PDFs. The similarity of estimated risk premium with at-the-money implied volatility is striking, as is the correlation of their first differences in Chart 35. However this is perhaps unsurprising: at-the-money implied volatility is commonly taken as a simple measure of the level of risk around the current price of options’ underlying asset and so captures risk aversion. The risk premia from the DDM is rather different. Most strikingly, it falls in 1998/99 when both measures of risk premia rise sharply, perhaps in response to the events of LTCM’s failure and the Asian crisis. However, this may be due to the shortcomings of the DDM, particularly its use of analysts’ earnings forecasts that are not always updated swiftly in response to news. The comparison with the DDM also does not compare like with like: the PDF-derived measure is a short-term risk premium over a horizon of one year. In contrast, since the DDM includes forecasts of future dividends over a longer range of future time periods, the risk premium from the DDM is over a far longer horizon.

**Chart 34: Other measures of the equity risk premium**



**Chart 35: First differences of three-month risk premium estimates versus implied volatility**



**4 Robustness checking**

Two checks on the robustness of the above results are presented here. Firstly, that results of a similar quality – in terms of their success at producing an estimated distribution that matches the observed distribution – cannot be produced by using a simpler methodology. One such

<sup>9</sup> The DDM finds the discount rate at which expected future dividends must be discounted in order for the resulting equity price to equal that observed in the market. The difference between this discount rate and the risk-free rate is the resulting equity risk premium. For further details see Inkinen *et al* (2010).

methodology might simply be to ‘shift’ the RND by the average amount by which its mean differs from that of the observed distribution of prices.

Second, we assess how parameter estimates evolve over time when estimated over an expanding window of data. That the estimated parameters are stable over time is important if this approach is to be used operationally. Our discussion here has involved finding the parameters that best adjust RNDs to the RWDs corresponding to a sample of observed prices. In practice, if this approach were to be used to forecast the RWD of future as yet unobserved prices, then it would only be likely to improve that forecast if the *ex post* parameters, estimated once those out-turns are known, were not significantly different to the *ex ante* parameters previously estimated from the sample without those out-turns. In other words, the ability of this method to improve our density forecasts depends on the estimated parameters being stable over time.

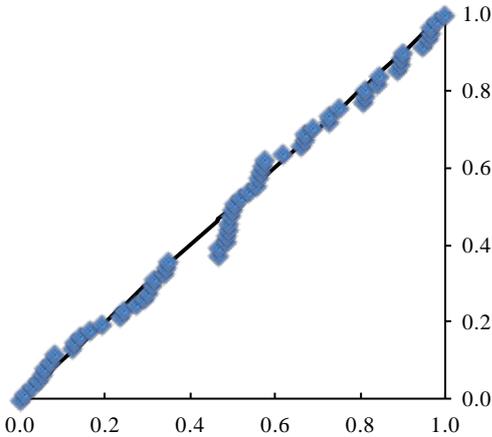
#### 4.1 *A comparison to a mean-shift calibration*

A simpler method for transforming risk-neutral option-implied distributions into real-world distributions is to simply ‘shift’ the RND by the average difference between the means of the RND and the RWD. This ‘mean-shift’ transformation is embodied in the Black-Scholes option-pricing framework that assumes that the underlying price is log-normally distributed (see Girsonav (1960)). While displacing each RND by the average excess return would, by construction, improve its fit to the existing data to which it is calibrated, it may not improve its fit to out-of-sample data generated in future.

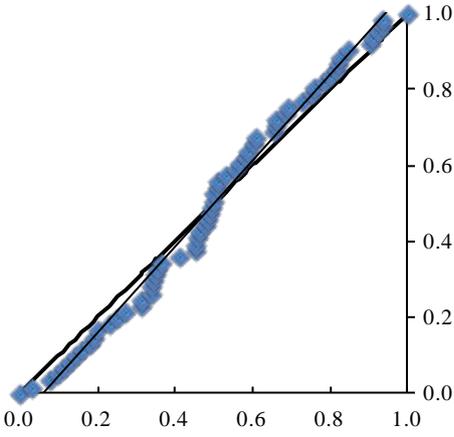
The mean shift is nested within the transformation afforded by the beta calibration function, and can be effected by imposing the restriction that  $j = 0$  in the estimation procedure described in Section 2.5.

The  $q$ - $q$  plots in Charts 36 and 37 compare the relative success of the two methods at matching the observed distribution. Unsurprisingly, the central percentiles of the distribution are closer matched under the mean shift, and lie closer to the diagonal. However, away from the mean the fit is far poorer, with a systematic bias in the tails of the distribution with points lying above or below the diagonals of the  $q$ - $q$  plot. This is unsurprising because it is only the first moment which has been altered. In contrast, the beta-shifted distribution gives a much closer average fit of the percentiles across the entire distribution. Although compensating for a constant risk premium improves the performance of the mean as a point forecast it does not improve the performance of the distributional forecast.

**Chart 36: *q-q* plots for RND, under beta transformation**



**Chart 37: *q-q* plot under mean shift**



These observations are confirmed in the summary statistics compared in Table 1. While the root mean-squared forecast error for the mean-shifted sample was reduced from 379.04 to 375.59, the joint log-likelihood was reduced from -424.55 to -426.54, reflecting how the mean shift produces a transformation that is less close in likelihood terms. The log-likelihood ratio marginally fails to reject the null hypothesis that two distributions coincide at the 10% significance level.

**Table 1: The results of the beta versus the mean-shift calibration**  
(*P-values are given in brackets*)

	<b>RND</b>	<b>Beta calibration</b>	<b>Mean-shift</b>
<b>Log likelihood</b>	-427.87	- 424.55	- 426.54
<b>Root mean square error</b>	375.59	379.04	375.59
<b>Log-likelihood test statistic</b>		3.98 (0.1367)	
<b>Kolmogorov-Smirnov test statistic</b>	0.15689 (0.3990)	0.0865 (0.9678)	0.0851 (0.9722)
<b>Cramer-von-Mises test statistic</b>	0.3090 (0.7180)	0.0445 (0.9963)	0.1304 (0.9736)

In assessing the relative fit of the beta calibration versus the mean-shift transformation, we first consider the well-known Kolmogorov-Smirnov (KS) test statistic: the maximum deviation between the percentiles of the distribution. This is the maximum deviation from the diagonal over the points of the *q-q* plot. As, for this data, the maximum deviation between the two quantiles occurs close to the mean of the distribution (visible around the centre of Chart 35), the mean shift actually produces a lower KS statistic than the beta calibration.

As we are seeking to optimise the fit across the entirety of the distribution, however, not just their maximum deviation, we prefer the Cramer-von-Mises (CvM) test statistic.<sup>10</sup> This is the

<sup>10</sup> See Anderson (1962) for the implementation used here.

average squared deviation between the quantiles of the two distributions. This is perhaps a more salient test statistic, as it considers the average distance from the diagonal across all the points of the  $q$ - $q$  plot, and is reduced, from 0.1304 to 0.0445, by using the beta calibration compared to the mean shift. This test is therefore able to formalise the improvement we observe visually in the  $q$ - $q$  plot in Chart 35 compared to Chart 36.

Whilst the CvM test statistic is greatly reduced by the use of the beta calibration instead of the mean shift, neither it nor the KS test fails to reject the hypothesis at a 10% significance level that the distribution obtained from the mean-shift transformation, and that of the data, are identical, with both test statistics giving  $p$ -values in excess of 0.97. Neither test has sufficient power, when used with a sample of only 60 observations, to distinguish between the two distributions. This is a consequence of the using non-overlapping observations of options prices and realised prices. Using overlapping observations would vastly increase the size of the sample, but would introduce auto-correlation into the data that would hinder our ability to carry out meaningful statistical tests.

The estimation of the calibration function, and therefore, the RWDs, is unconditional. It is based on the premise that, over time, investors do not make systematic errors in their own RWD forecasts. It is informed only by past out-turns, and how they are distributed across the option-implied RNDs. The next section assesses how sensitive the calibration function estimated using the data on the FTSE 100 and short sterling options, is to the particular sample of past out-turns used.

#### *4.2 Evolution of estimated beta parameters*

The main assumption underlying this methodology is that there is a systematic difference between the RND and RWD, which is both stable over time, and can be adequately described by a beta function. If this is true, then as the sample of price out-turns increases, the estimate of that systematic difference should become more precise. In other words, the estimated beta parameters should converge to some constant values.

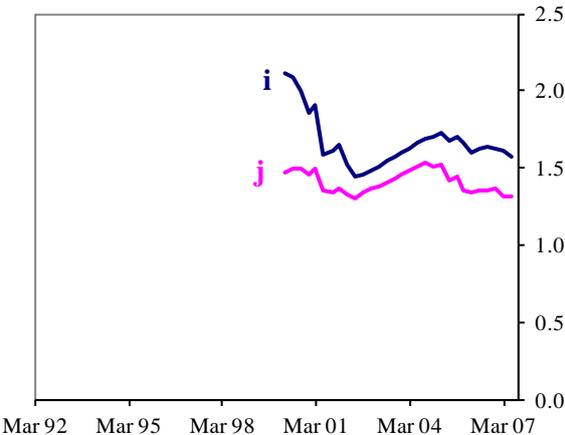
That has important implications, one practical, one economical. First, if the approach were to be used operationally, it would require the adjustment of a RND whose actual out-turn had not yet been observed. This adjustment would therefore have to be made using the parameters obtained from past calibrations. Were the parameters to show a large degree of variation over time, this would frustrate forecasters’ ability to determine this choice of parameters reliably.

Second, if the results of such a risk-neutral to real-world transformation process are to be used to inform our economic understanding of the past, the arrival of new data about the very recent past should not materially alter our understanding of the more distant past. For instance, it would be undesirable if estimates of the implied risk premium were to be substantially revised a number of years later, if the estimated calibration parameters for that larger sample turned out to be materially different to the calibration parameters previously estimated.

4.2.1 Evolution of estimated beta parameters: FTSE 100

The evolution of the estimated beta parameters ( $j, k$ ), as the data estimation window expands, are presented in Chart 38. For each date, this chart plots the ( $j, k$ ) values estimated using all quarterly data from March 1992 up until that date. So the first plotted points for March 2000 are estimated using the 31 observations until that date, whereas the final data points are estimated using the full sample of 61 observations.

**Chart 38: The evolution of the estimated beta parameters for FTSE 100 over time**



Although the estimated beta parameters may now be starting to show signs of stabilisation (Chart 38), expanding the estimation window had previously led to noticeable fluctuations. That variation can be better understood by considering the transformation of the beta parameters

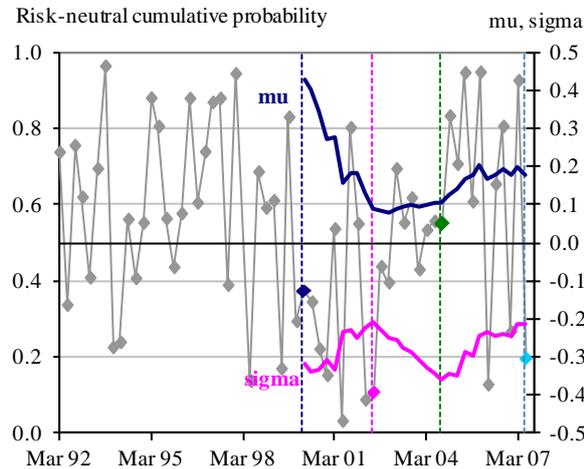
$(\sigma(j,k), \mu(j,k))$ , defined below. These can be thought of as representing an adjustment to the location and shape of the distribution.

*location adjustment:*  $\mu = j - k$

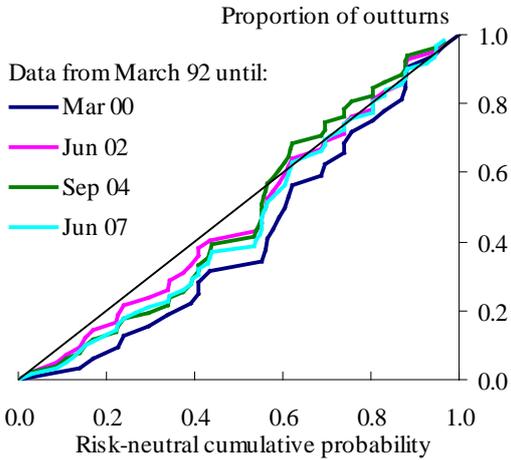
*shape adjustment:*  $\sigma = 1 - k$

These measures are shown in Chart 38, alongside the cumulative probability of each out-turn, according to the RWD estimated three months previously. The cumulative probabilities are summarised in the four PIT  $q-q$  plots in Chart 40, constructed using the out-turns from March 1992 to each of the four vertical dashed lines in Chart 39. Between March 1992 and March 2000 the majority of out-turns were drawn from the upper half of the RNDs, so the RNDs appeared to understate the mean and overstate the variance of the RWDs. Note how the March 2000 line in Chart 40 sags far below the gradient-one line, and also the kink around a risk-neutral cumulative probability of 0.5. The optimal calibration needed to map this sample of out-turns onto the RNDs therefore needs to increase the means of the RNDs (large positive  $\mu$ ) and decrease their variances (sufficiently negative  $\sigma$ ).

**Chart 39: The estimated beta parameters and cumulative probabilities of the FTSE out-turns according to the RND estimated three months previously**



**Chart 40:  $q-q$  plots of three-month FTSE out-turns expanding periods**

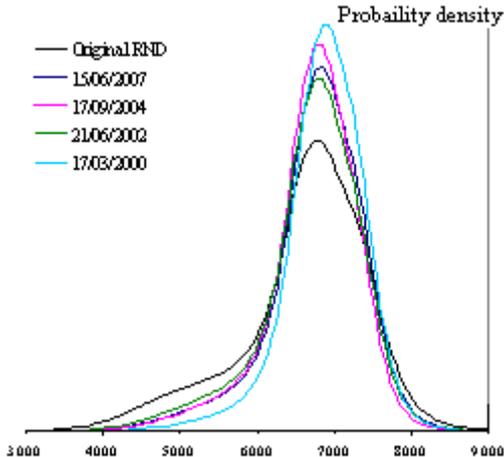


However, adding the March 2000–June 2002 out-turns changes the pattern of past out-turns notably. In contrast to the earlier period, the majority of these additional out-turns were less than the median of their corresponding option-implied RND. This means that, on average, the dispersion of out-turns across the quantiles of the option-implied RNDs becomes a lot more balanced. That is reflected in the PIT  $q-q$  plot: the pink line in Chart 40 lies much closer to the

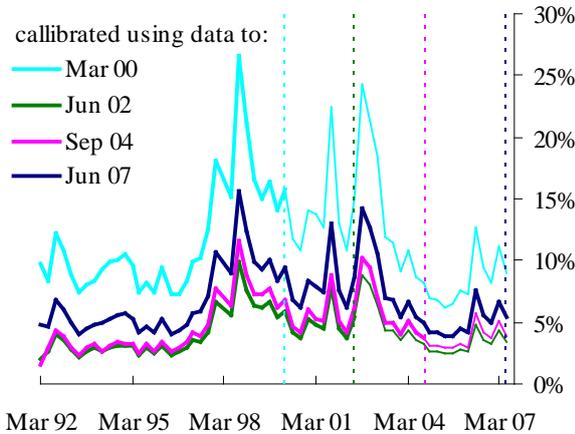
gradient-one line, nor is the kink as stark. The magnitude of both the mean and variance correction therefore decreases. Similarly, the addition of June 2002–September 2004 data, which are clustered around the median, tend to increase the extent to which RNDs, on average, overstate the true variance. Therefore, the estimated calibration must deliver a stronger variance reduction.

What matters more, however, is the sensitivity of the calibrated RWD to this variation in the estimated beta parameters. Chart 41 illustrates how the prevailing estimate of the 17 March 2000 three-month FTSE real-world PDF changes as the estimate of the beta calibration parameters is updated. Changes in the mean of the estimated RWD translate into changes in the estimate of the option-implied risk premium. Chart 42 therefore shows how the historical option-implied risk premium series changes for the different estimates of the beta parameters. Although the change in the estimated beta parameters from Mar 2000 to Jun 2002 corresponded to a marked revision of the historical risk premia series, more recent updates of those parameters had less of an impact.

**Chart 41: RWD estimates for the FTSE 100 over different estimation windows**



**Chart 42: Variation in estimated ERP series, depending on beta parameter estimation window**

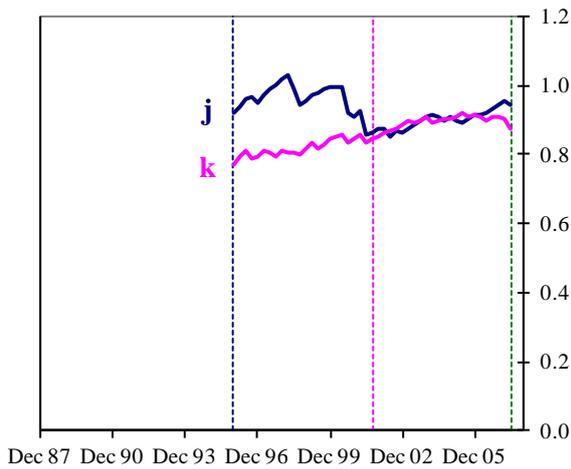


4.2.2 Evolution of estimated beta parameters: short sterling

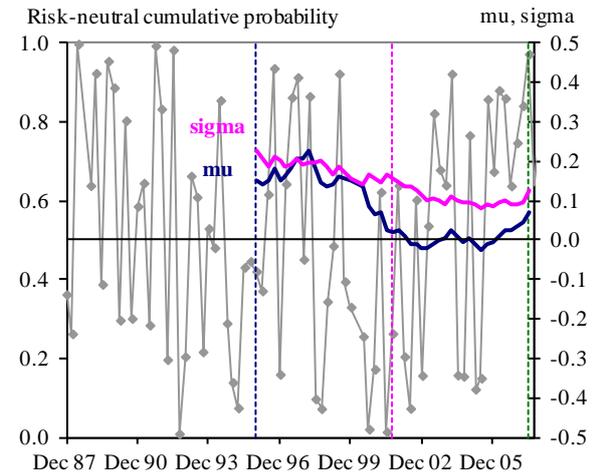
The evolution of the three-month short sterling beta parameters are shown in Chart 43. The corresponding location and shape adjustment is shown in Chart 44. These appear more stable than those for the FTSE 100, suggesting that the nature of the transformation required to obtain the RWD from the RND changes less over time. This may suggest that short-term interest rate markets are less susceptible to changes in investor preferences and risk premia, resulting from market developments, than equity markets.

As might be expected, this does not generate material revisions to the estimate of a past RWD (Chart 45).

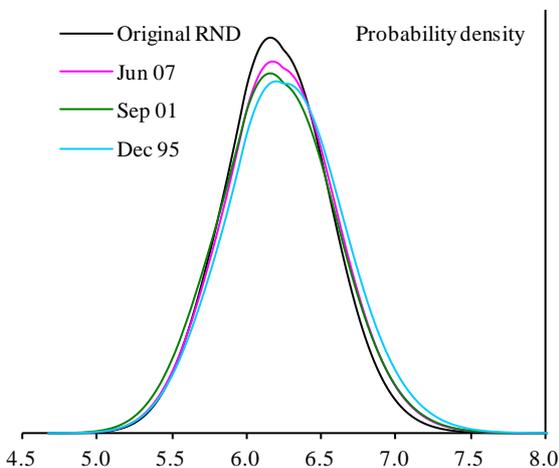
**Chart 43: The evolution of the estimated beta parameters for short sterling over time**



**Chart 44: The estimated beta parameters and cumulative probabilities of the FTSE 100 out-turns according to the RND estimated three months previously**



**Chart 45: RWD estimates for short sterling over different estimation windows**



## 5 Conclusion

This study offers a means of transforming the risk-neutral probability densities of future asset prices obtained from derivatives contracts into estimates of the ‘real-world’ probability densities that might better reflect the probabilities considered by market participants. The approach taken

is ‘empirical’ and uses a beta CDF to transform the risk-neutral density into that which best fits observed prices at the contracts’ expiry.

The methodology is applied to options contracts on the FTSE 100 and short sterling, but is highly flexible and could be applied to derivatives contracts on any underlying asset. The resulting real-world probability densities offer a superior average fit across the distribution of observed prices than their risk-neutral counterparts. Besides offering an improved insight into the perceived uncertainty surrounding future prices, a comparison of risk-neutral and real-world probabilities reveals new information as to how attitudes towards risk affect asset prices, and allows a measure of investors’ risk aversion to be backed out from the difference in the means of the two distributions. The resulting calibrations also appear fairly robust when applied to an expanding data set of prices that is updated in light of new observations, up until June 2007.

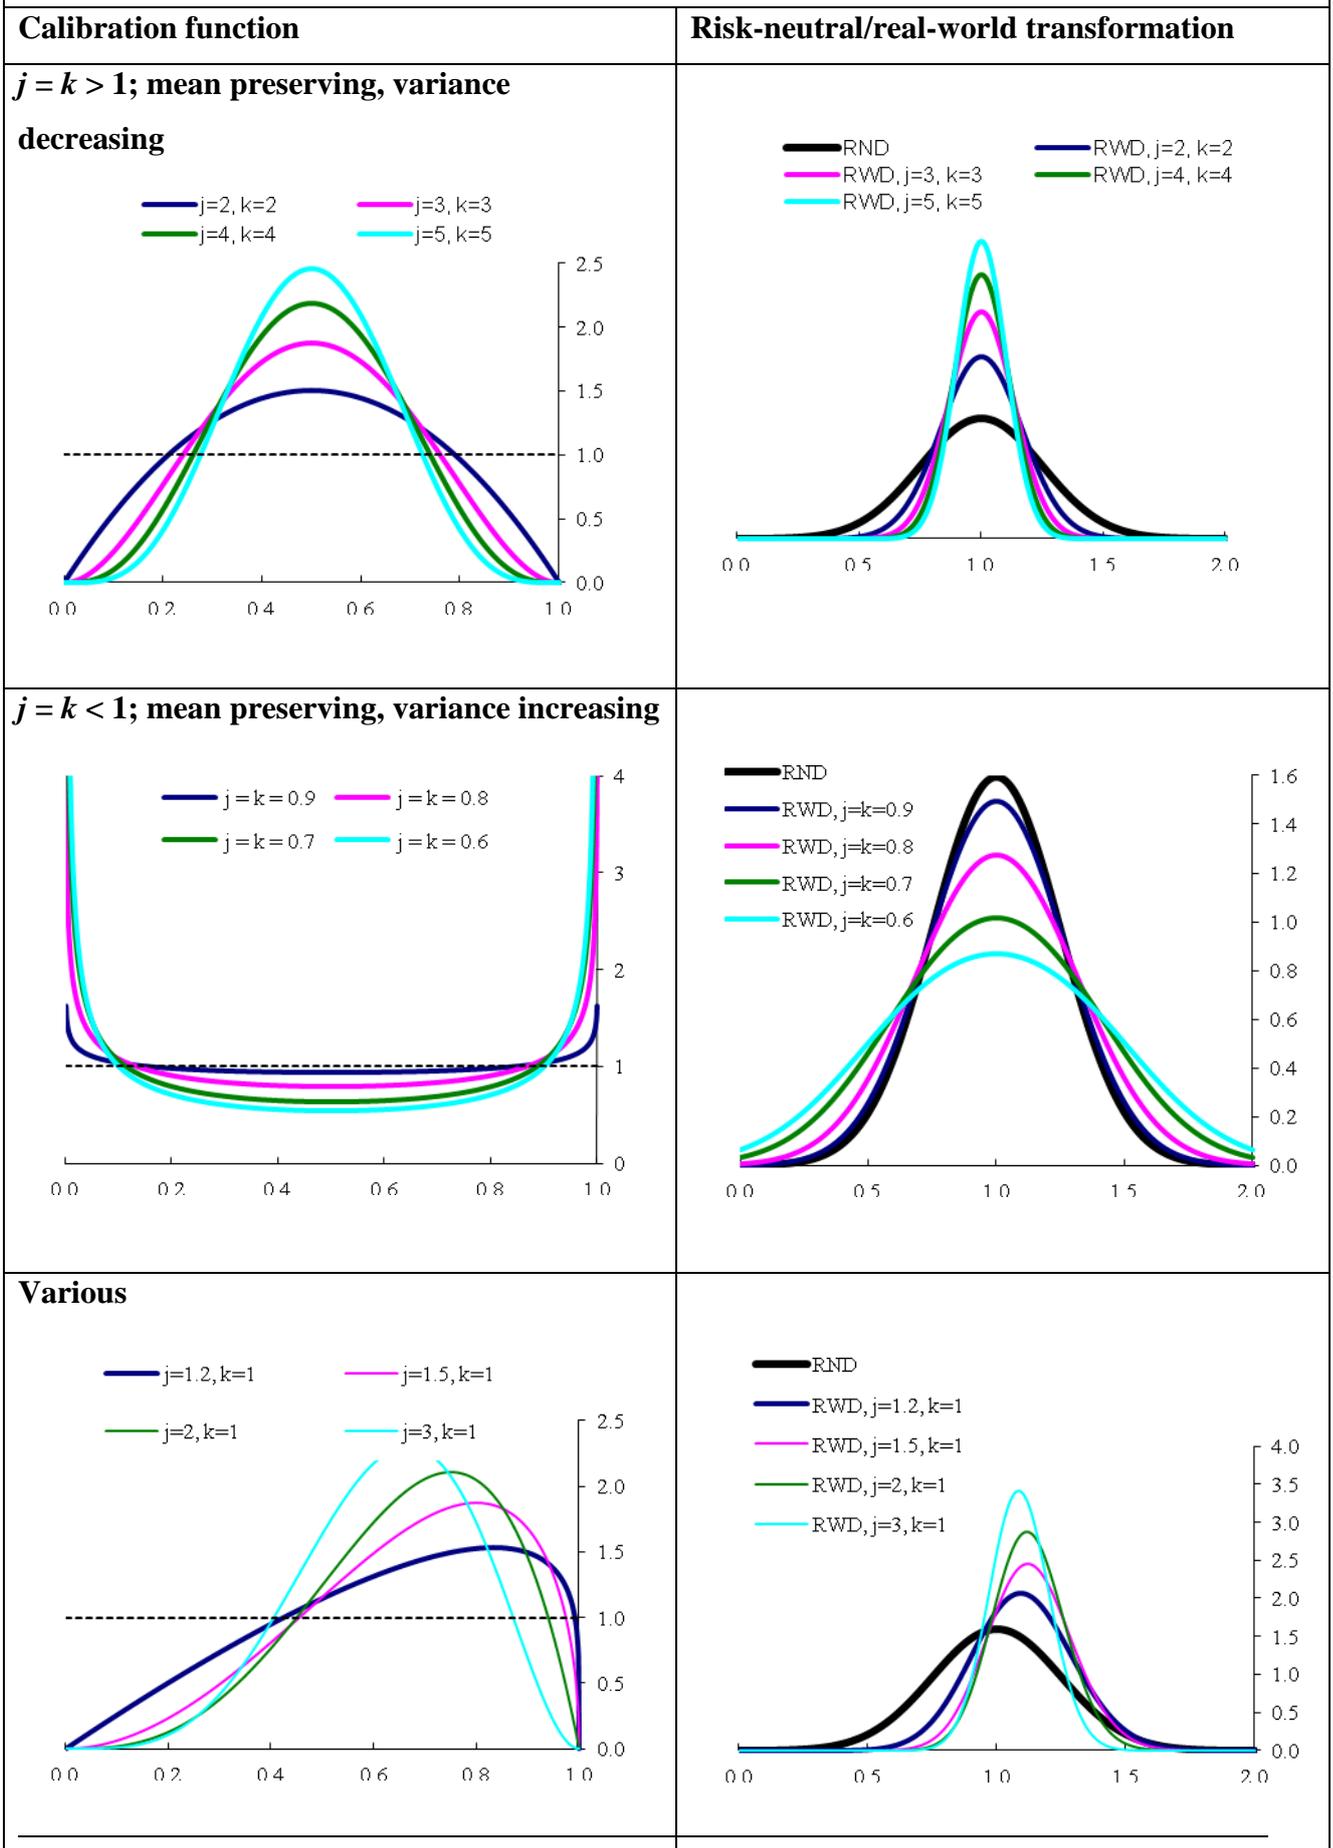
We leave as further work the extension of our results beyond June 2007 up to the present day. Such an extension may need to consider how best to ‘operationalise’ the method so that the day-to-day calculation of real-world PDFs can be used for policy purposes. Given the volatility of financial markets in the intervening period, and the possible volatility, therefore, of the estimated parameters that control the calibration of the risk-neutral to the real-world density, this may warrant some consideration of whether to ‘condition’ the parameter estimates on some other observable variable – for example that relating to the real economy. Doing so may increase the method’s predictive power.

## **Appendix 1: The beta PDF and the transformations it affords**

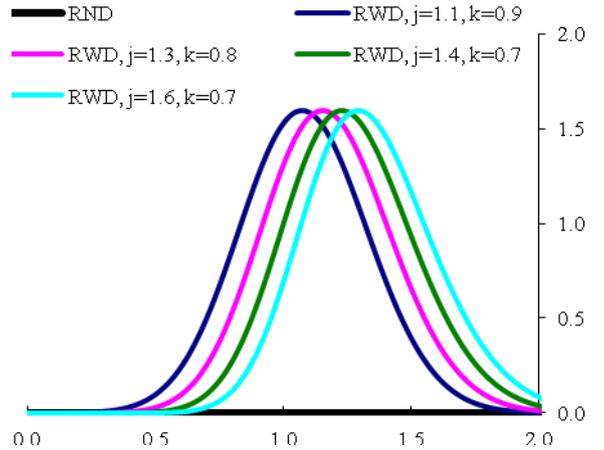
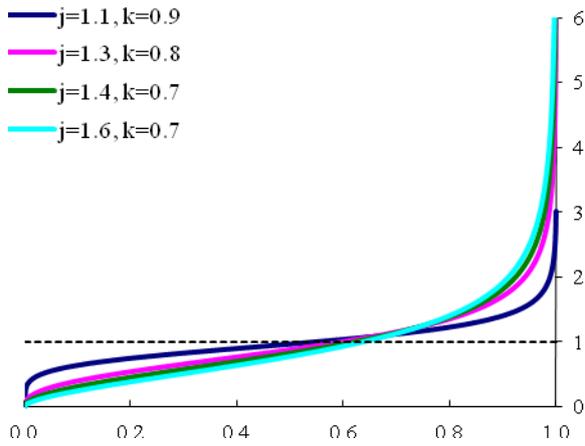
This appendix gives examples of the various transformations afforded by the PDF of the Beta distribution. Though it uses only two parameters, it nests many simple forms of transformation such as a mean shift, mean-preserving changes in variance, and changes involving mean, variance and skew.

These are illustrated in the charts of Figure 1. Charts in the left-hand column plot the beta CDF resulting from different combinations of its  $j$  and  $k$  parameters; the right-hand column contains plots of the corresponding resulting transformations of an example RND.

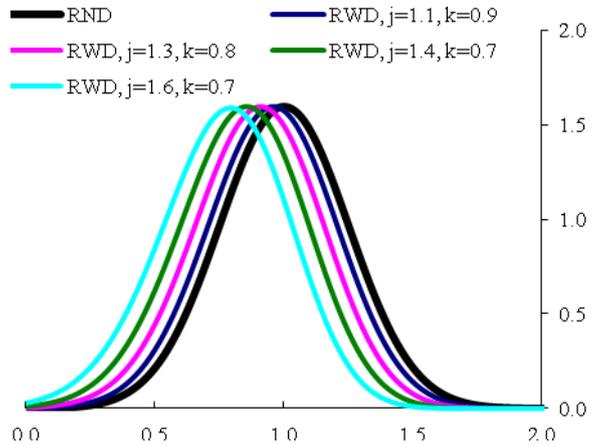
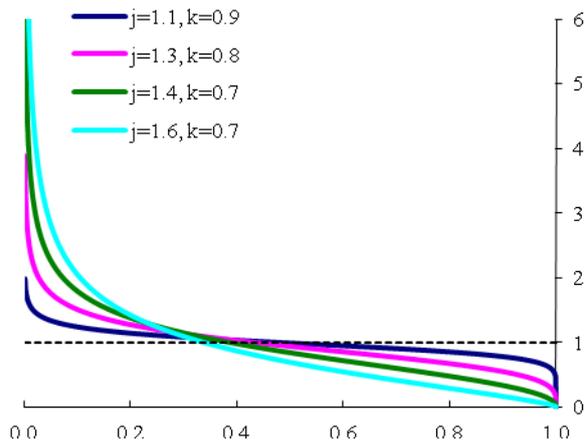
**Figure 1: The transformations afforded by the beta CDF function**



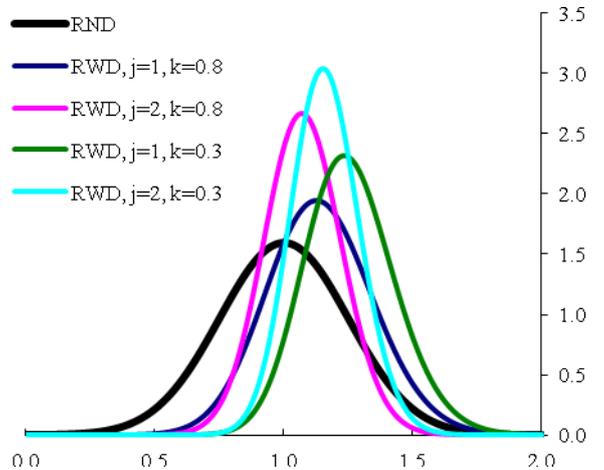
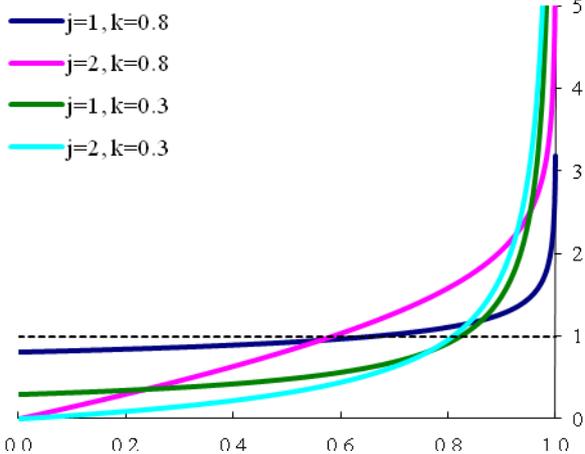
**$j = 1, k < 1$ ; mean shift (increasing)**



**Mean shift (decreasing)**



**Various**



## References

**Alonso, F, Blanco, R and Rubio, G (2006)**, 'Testing the forecasting performance of IBEX 35 option-implied risk-neutral densities', *Banco de España Working Paper No. 0630*.

**Anagnou-Basioudis, I, Bedendo, M and Hodges, S (2005)**, 'Forecasting the accuracy of implied and GARCH-based probability density functions', *Review of Futures Markets*, Vol. 11, pages 31-66.

**Anderson, T (1962)**, 'On the distribution for the two sample Cramer-von Mises Criterion', *The Annals of Mathematical Statistics*, Vol. 33, pages 1,148-59.

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

**Breeden, D and Litzenberger, R (1978)**, 'Prices of state-contingent claims implicit in option prices', *Journal of Business*, Vol. 51, pages 621-51.

**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.

**de Vincent-Humphreys, R and Puigvert, J (2012)**, 'A quantitative mirror on the Euribor market using implied probability density functions', *Eurasian Economic Review*, Vol. 2 No. 1, pages 1-31.

**Fackler, P and King, R (1990)**, 'Calibration of option-based probability assessments in agricultural commodity markets', *American Journal of Agricultural Economics*, Vol. 72, pages 73-83.

**Gai, P and Vause, N (2005)**, 'Measuring investors' risk appetite', *Bank of England Working Paper No. 283*.

**Girsonav, I (1960)**, 'On transforming a certain class of stochastic processes by absolutely continuous substitution of measures', *Theory of Probability and its Applications*, Vol. 5, pages 285-301.

**Heston, S (1993)**, 'A closed-form solution for options with stochastic volatility with applications to bond and currency options', *Review of financial studies*, Vol. 6, pages 327-43.

**Inkinen, M, Stringa, M and Voutsinou, K (2010)**, 'Interpreting equity price movements since the start of the financial crisis', *Bank of England Quarterly Bulletin*, Vol. 50, No. 1, pages 24-33.



**Liu, X, Shackleton, M, Taylor, S and Xu, X (2007)**, ‘Closed-form transformations from risk-neutral to real-world distributions’, *Journal of Banking and Finance*, Vol. 31, pages 1,501-20.

**Malz, A (1997)**, ‘Option-based estimates of the probability distribution of exchange rates and currency excess returns’, *Federal Reserve Bank of New York*.

**Rosenblatt, M (1952)**, ‘Remarks on a multivariate transformation’, *Annals of Mathematical Statistics*, Vol. 23, pages 470-72.

**Shackleton, M, Taylor, S and Yu, P (2010)**, ‘A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices’, *Journal of Banking and Finance*, Vol. 34, pages 2,678-93.

**Weinberg, S (2001)**, ‘Interpreting the volatility smile: an examination of the information content of options prices’, Federal Reserve Board, *International Finance Discussion Paper No. 706*.

