Journal of Financial Markets 9 (2006) 180–198

Journal of FINANCIAL MARKETS

www.elsevier.com/locate/finmar

Option-implied risk preferences: An extension to wider classes of utility functions

Byung Jin Kang*, Tong Suk Kim

Graduate School of Management, Korea Advanced Institute of Science and Technology, 207-43 Cheongnyangni 2-dong, Dongdaemun-ku, Seoul, 130-722, Korea

Accepted 23 December 2005
Available online 27 March 2006

Abstract

Investors' risk aversion functions can be derived from the risk neutral probability density functions (RN-PDFs) and an assumed well-behaved functional form for the utility function. This paper extends the analysis to more general cases by assuming wider classes of utility functions. Using FTSE 100 index options, we evaluate the forecasting ability of RN-PDFs and subjective PDFs with five assumed utility functions and then derive the corresponding option-implied risk aversion functions. From our empirical analysis, we find that: (1) assuming more flexible utility functions generally increases the forecasting ability of the derived subjective PDFs; and (2) the measure of relative risk aversion is significantly different from zero and decreases across wealth. These results essentially hold regardless of forecast horizon. Out of sample tests also confirm the robustness of our findings.

JEL classification: G13; C12

Keywords: Implied risk aversion; Risk neutral probability density function; HARA utility function; Log plus power utility function; Linear plus exponential utility function; FTSE 100 index option

1. Introduction

Under a complete market economy that implies the existence of a representative agent, risk preference structure, asset return dynamics and state price densities are mutually related. As was noted by Leland (1980) and Bick (1990), the first two imply the third one.

*Corresponding author. Tel.: 82 2 958 3689; fax: 82 2 958 3618.
E-mail address: bjk@kaist.ac.kr (B.J. Kang).

1386-4181/$ - see front matter
doi:10.1016/j.finmar.2005.12.004

Applying these theoretical restrictions to an asset market, many authors have investigated the investors' risk aversion functions implied by option prices.

In this line of research, there have been developed two different methodologies to derive the investors' risk aversion functions. The first method is to derive the risk aversion functions by joint estimation of risk neutral PDFs (RN-PDFs) and subjective PDFs. Most research on option-implied risk preferences including Jackwerth (2000), Ait-Sahalia and Lo (2000), Coutant (1999), Luis (2001), Perignon and Villa (2002), Rosenberg and Engle (2002), Kliger and Levy (2002), Giamouridis (2003) and Ziegler (2003) falls into this category. When using this method, subjective PDFs are mostly estimated from historical time series of underlying asset returns under stationary assumptions or stochastic process assumptions. Combining them with the RN-PDFs that are estimated from cross sectional option prices, the risk aversion function across wealth could be derived. As this method does not need to assume any functional form for utility functions, it has the great advantage that the resulting risk aversion function could take any shape. However, it also has the disadvantage that the stationary assumptions for subjective PDFs become questionable, especially when faced with changing RN-PDFs.[1]

The second method, introduced by Bliss and Panigirtzoglou (2004), was developed to mitigate the aforementioned weak point. They make use of a specific class of utility functions to adjust RN-PDFs in order to produce subjective PDFs. That is, instead of making assumptions on underlying asset return distributions, they use a well-behaved functional form for the underlying utility functions to derive implied-subjective PDFs.[2] With this newly introduced method by Bliss and Panigirtzoglou (2004), the derived subjective PDF is time-varying like the RN-PDF. Furthermore, unlike the subjective PDF estimated from historical data, their subjective PDF has a forward looking property. However, as they only assume the functional form for utility functions as either a power or an exponential function, the resulting investors' risk preferences are restricted to be either constant RRA (relative risk aversion) or constant ARA (absolute risk aversion). Therefore, many previously reported findings with respect to the functional form for risk aversion such as 'U-shaped' or 'generally declining' across wealth cannot be explained by their method.[3]

In this paper, we accept the paradigm of Bliss and Panigirtzoglou (2004) and extend their method by relaxing the assumption of the power- and exponential- utility functions. We do not restrict the assumed functional form for utility functions to the power or the exponential function. Our assumed functional forms for utility functions are as follows: power function, exponential function, Hyperbolic Absolute Risk Aversion (HARA) function with two parameters, log plus power function and linear plus exponential function.[4] Obviously, the power utility function is a special case of the HARA and the log plus power utility, and the exponential utility function is a special case of the HARA and the linear plus exponential utility. While the RRA functions from the first two utility

[1] Bliss and Panigirtzoglou (2004) asserted that time-invariant subjective PDFs may be inconsistent with time-varying RN-PDFs.

[2] An implied-subjective PDF means the PDF which is adjusted from the RN-PDF under the assumption of a specific utility function.

[3] The assumption of the power or exponential utility functions has been widely accepted because of their well behaved property. However, the investors' relative (absolute) risk aversion from option prices has been reported not to be constant in most previous researches. For example, Jackwerth (2000) and Ait Sahalia and Lo (2000) reported 'U-shaped' risk aversion and Kliger and Levy (2002) reported 'decreasing' RRA.

[4] In the next section, we will provide the detailed structures for our assumed utility functions.

classes depend on only one parameter, those from the last three utility classes depend on two parameters.

The selection of candidate utility functions in this study mainly depends on their tractability and parsimoniousness. Some well-known utility functions other than our assumed functions were also examined in our preliminary test.[5] Most of them were found either to be outperformed by our assumed functions in overall performances or to fail to produce well-behaved RRA functions and thus were excluded from our regular analysis. Nevertheless, it should be kept in mind that there remain many other unexplored large classes of utility functions. Assumed utility functions in this study are not exhaustive and thus cannot always guarantee recovering the true risk attitudes of investors. However, by assuming more general classes of utility functions than simple power or exponential, we can avoid excessively restricting the shape of empirically driven RRA (or ARA) functions and can thus reduce the probability of mis-specifying the investors' risk preferences. That is, when faced with an unknown functional form for risk preferences, it is more likely that our extension, by assuming more flexible utility functions, can come close to providing the investors' true risk preferences. Furthermore, all the advantages of the method in Bliss and Panigirtzoglou (2004) which are previously discussed are still available in our study.

Using FTSE 100 index option data from April 1992 through July 2000, we examined the option-implied investors' risk preferences. In order to precisely assess the incremental performance of our extended method, we examine data similar to those that were used by Bliss and Panigirtzoglou (2004).[6] From our empirical analysis, we find that assumed more flexible utility functions result in the better forecasting ability of the risk-adjusted subjective PDFs. That is, the flexible utility functions such as HARA, log plus power and linear plus exponential outperform in forecasting ability both power and exponential utility not only from an in-sample test, but also from an out-of-sample test. We also find that investors' risk aversion is significantly different from zero and exhibits 'decreasing RRA' across wealth. All these findings suggest that either the power or the exponential utility function would not be sufficient to characterize the investors' true risk preferences. We need more flexible utility functions to characterize investors' true risk preferences, and the three alternative utility functions examined in this study can be candidates for them.

The remainder of this paper is organized as follows. Section 2 presents the model for deriving the risk aversion from option prices and the method for estimating the RN-PDFs from option prices. Section 3 describes the data. Section 4 presents the empirical results. Finally, Section 5 concludes with a brief summary.

2. Model and estimation method

2.1. Introducing various classes of utility functions

In this study, we assume five classes of utility functions to produce risk-adjusted subjective PDFs from RN-PDFs. Each of them can be expressed as follows:

[5] Log, quadratic, special exponential, arc-tangent (Kallberg and Ziemba 1983), exponential-power (Torkamani and Haji-Rahimi, 2001), sum-exponential and sum-power utility functions (Bell, 1988; Pedersen and Satchell, 2001 Gelles and Mitchell, 1999) are examined in our preliminary test.

[6] The sample period of FTSE 100 index options examined by Bliss and Panigirtzoglou (2004) extends from June 1992 through March 16, 2001.

Power utility:

$$U(W) = \frac{1}{1-\alpha} W^{1-\alpha}, \quad RRA(W) = \alpha, \tag{1}$$

Exponential utility:

$$U(W) = -\frac{\mathrm{e}^{-\alpha W}}{\alpha}, \quad RRA(W) = \alpha W, \tag{2}$$

HARA utility:

$$U(W) = \frac{1}{\alpha - 1}[\alpha W + \beta]^{1-1/\alpha}, \quad RRA(W) = \frac{W}{\alpha W + \beta}, \tag{3}$$

Log plus power utility:

$$U(W) = \alpha \log W + \frac{1}{\beta} W^{\beta}, \quad RRA(W) = \frac{\alpha + (1-\beta)W^{\beta}}{\alpha + W^{\beta}}, \tag{4}$$

Linear plus exponential utility:

$$U(W) = \alpha W - \frac{1}{\beta}\mathrm{e}^{-\beta W}, \quad RRA(W) = \frac{\beta W \mathrm{e}^{-\beta W}}{\alpha + \mathrm{e}^{-\beta W}} \tag{5}$$

where W denotes the level of wealth.

The power and exponential utility functions in Eqs. (1) and (2) are two of the most widely used functions in finance and economics. These two utility functions contain only one parameter and thus the corresponding risk aversion functions are entirely characterized by the estimate of the single parameter. The HARA (or Linear Risk Tolerance) utility functions are a class of utility functions whose absolute risk aversions are hyperbolic and cover a large body of commonly used utility functions. Both the power and exponential utility functions are special cases of the HARA utility functions. In this study, we assume the linear risk tolerance ($T(W) = \alpha W + \beta$) and then express the corresponding utility function in equation (3).[7] The last two utility functions in Eqs. (4) and (5) have been widely used to explain a switching problem between gambles.[8] To the extent we know, Bell (1988) is the first to explore large classes of preferences guaranteeing the switching behavior. He suggested the exhaustive classes of utility functions that allow one-switch for additive gambles. The linear plus exponential utility function in Eq. (5) was proven to be one of them. Pedersen and Satchell (2001) and Gelles and Mitchell (1999) also suggested the exhaustive classes of utility functions that allow one-switch for multiplicative gambles.

[7]The HARA utility function in this study is suggested as an example of general HARA utility functions by Bick (1990).

[8]The switching means that preference of one gamble over another gamble changes as wealth increases. According to the number of switches between gambles over the entire range of initial wealth, we can characterize the corresponding functional form for utility functions. For example, the power utility function is the zero-switch utility function for multiplicative gambles and the exponential utility is the zero-switch utility function for additive gambles.

The log plus power utility function in Eq. (4) was proven to be one of them. As an extreme case, the log plus power utility becomes equivalent to the power utility and the linear plus exponential utility becomes equivalent to the exponential utility with zero value of α. The utility functions from equation (3) to (5) contain two parameters and thus can allow a more flexible shape of risk aversion functions.[9]

2.2. Deriving the risk aversion functions from option prices

Given an assumed utility function and an estimated RN-PDF, we can derive an implied subjective PDF. As Jackwerth (2000) notes, in equilibrium the implied risk aversion can be solved from the RN- and subjective PDFs:

$$ARA(S_T) = -\frac{U''(S_T)}{U'(S_T)} = \frac{P'(S_T)}{P(S_T)} - \frac{Q'(S_T)}{Q(S_T)}, \tag{6}$$

$$RRA(S_T) = -\frac{S_T U''(S_T)}{U'(S_T)} = S_T\left(\frac{P'(S_T)}{P(S_T)} - \frac{Q'(S_T)}{Q(S_T)}\right), \tag{7}$$

where ARA and RRA are the absolute and relative risk aversion functions, respectively, U is a state-independent utility function across states, P is a subjective PDF, Q is a RN-PDF, S is the return on the market portfolio across states and T is an option expiration date.

By transforming Eq. (7), the implied subjective PDF can be solved:

$$P(S_T) = \frac{Q(S_T)/U'(S_T)}{\int \left(Q(x)/U'(x)\right)\, \mathrm{d}x}, \tag{8}$$

The denominator in Eq. (8) is added to normalize the resulting subjective PDF to integrate to one. As can be seen easily in Eq. (8), the subjective PDF can be derived from the RN-PDF and the assumed utility function as follows:

Power utility:

$$P_t(S_T) = \frac{S_T^{\alpha} Q_t(S_T)}{\int x^{\alpha} Q_t(x) \mathrm{d}x}, \tag{9}$$

Exponential utility:

$$P_t(S_T) = \frac{\mathrm{e}^{\alpha S_T} Q_t(S_T)}{\int \mathrm{e}^{\alpha x} Q_t(x)\, \mathrm{d}x}, \tag{10}$$

HARA utility:

$$P_t(S_T) = \frac{S_T(\alpha S_T + \beta)^{1/\alpha} Q_t(S_T)}{\int x(\alpha x + \beta)^{1/\alpha} Q_t(x)\, \mathrm{d}x}, \tag{11}$$

[9] They can produce decreasing or increasing RRA (or ARA) functions as well as constant RRA (or ARA) functions.

Log plus power utility:

$$P_t(S_T) = \frac{\left(S_T/\alpha + S_T^{\beta}\right) Q_t(S_T)}{\int \left(x/\alpha + x^{\beta}\right) Q_t(x)\, \mathrm{d}x}, \tag{12}$$

Linear plus exponential utility:

$$P_t(S_T) = \frac{\left(1/\alpha + \mathrm{e}^{-\beta S_T}\right) Q_t(S_T)}{\int \left(1/\alpha + \mathrm{e}^{-\beta x}\right) Q_t(x)\ \mathrm{d}x}. \tag{13}$$

To estimate the parameters in the above equations, we need to: (a) compute risk neutral PDFs from option prices; and (b) test the forecasting ability of derived subjective PDFs. In the next two sub-sections, we present the statistical methods which will be applied in each step.

2.3. Estimating risk neutral probability density functions (RN-PDFs)

Methods of estimating RN-PDFs can either be parametric or non-parametric. The parametric approach can be further divided into expansion methods, generalized distribution methods, and mixture methods. The non-parametric approach can also be divided into kernel methods, maximum entropy methods, and curve fitting methods.[10] In this paper, we use the smoothed implied volatility smile (SML) method of curve fitting methods.[11]

The smoothed implied volatility smile (SML) method, originally introduced by Shimko (1993), fits a function through observed implied volatilities.[12] The fitted implied volatility function is translated into an option price function and then the implied PDFs are recovered using the Breeden and Litzenberger (1978) result:

$$q(S_T) = \mathrm{e}^{rt} \frac{\partial^2 C(K,t)}{\partial K^2}\bigg|_{K=S_T} \tag{14}$$

As noted by Bahra (1997), Syrdal (2002), and Brunner and Hafner (2003), two related issues have to be considered in the SML method. First, we cannot know the implied volatility function beyond the range of traded strike prices. Therefore, we must extrapolate or model the tails of distributions to obtain a well-defined PDF. Second, the estimated implied volatility function should not allow for arbitrage. To guarantee this, implied probabilities should be non-negative, integrate to one and satisfy the martingale restriction, which means that discounted asset prices have to be martingales with respect to these probabilities.

Shimko (1993) assumed that implied volatility is a quadratic function of the strike price, which is equivalent to the assumption of the log-normally distributed tails. His method

[10] Jackwerth (1999) extensively discussed various methods of estimating RN-PDFs from option prices.

[11] As was discussed by Cooper (1999), Bliss and Panigirtzoglou (2002), the smoothed implied volatility smile method generally tends to produce more stable RN-PDF than the double log-normal mixture method.

[12] As pointed out by Bates (1991), option prices are not suitable for fitting because of their substantial amount of non-linearity.

cannot, however, always satisfy the martingale restriction. Thus, Malz (1997) assumed that implied volatility is a quadratic function with respect to the option's delta. Fitting the implied volatility function with respect to the option's delta has an advantage in that the corresponding implied density integrates to one and satisfies the martingale restriction. Campa et al. (1998), and Bliss and Panigirtzoglou (2002), furthermore, used a smoothing cubic spline method to fit the implied volatilities. Campa et al. (1998) fitted the implied volatilities as a function of strike prices, while Bliss and Panigirtzoglou (2002) fitted them with respect to the option's delta.

In this study, we use the cubic spline method on the option's delta spaces. To model the tails of implied distributions, we should also extrapolate the spline function outside the range of observed option data. In this study, we extend the spline function in a horizontal manner outside the range of option data, which is suggested by Bliss and Panigirtzoglou (2004). Our method can be constructed as the following optimization problem:[13]

$$\min_{\Psi} \quad \left[\lambda \sum_{i=1}^{n} \varpi_i \left[y_i - f(x_i; \Psi) \right]^2 + (1-\lambda) \int f''(x; \Psi)^2 \, \mathrm{d}x \right], \tag{15}$$

where x_i and y_i are option deltas and implied volatilities, respectively, $f(x, \Psi)$ is a spline function, Ψ is a parameter matrix of the spline function, ϖ_i is a weighting parameter for observation i and λ is the smoothing parameter.[14]

In Eq. (15), the first part relates to the goodness of fit (GOF) of the spline function. The second part relates to the smoothness of the spline function. λ, known as the smoothing parameter, determines the relative weight between two parts. A large value of λ increases the GOF and a resulting density function may show too much oscillation. With a low value of λ, by contrast, a spline function is too smooth and does not fit the data well. At the extreme, the spline function would accurately interpolate the observed data with λ equal to one and minimize the curvature with λ equal to zero.

In fact, choosing a suitable smoothing parameter is one of the very important issues and several methods have been proposed to determine an optimal smoothing parameter. The cross validation score method (Craven and Wahba, 1979; Hastie and Tibshirani, 1990) is one well-known method. In this paper, we check the sensitivity of our empirical results to the values of λ ranging from 0.99 to 0.999 in our preliminary tests on a sub-sample as in Bliss and Panigirtzoglou (2004). As the results are insensitive to the choice of λ, we therefore apply the SML method with λ of 0.99 to extract RN-PDFs in this paper.

2.4. Testing the forecast ability of derived subjective PDFs

Once a RN-PDF is produced from each option's cross-sectional prices, then a corresponding subjective PDF for an assumed specific utility function can be derived from Eqs. (9)–(13). As in Bliss and Panigirtzoglou (2004), the next step we should take is to test that the estimated subjective PDFs are equal to the true PDFs.

Under the null hypothesis that the estimated PDFs are the true PDFs and that realized returns at each option expiry date are independent, the inverse probability transformations

[13] For more concrete analyses, see Syrdal (2002) and Kim and Kim (2003).

[14] In this study, we use the equal weight between observations. From our preliminary tests, the results were found to be robust to the choice of alternative weighting schemes between observations, such as vega weighting.

of the realized returns will be independently and uniformly distributed.

$$y_t = \int_{-\infty}^{X_t} \hat{P}_t(s)\,\mathrm{d}s \sim i.i.d.\mathrm{U}(0,1), \tag{16}$$

where X_t is a realized return at an option expiry date and $\hat{P}_t(\cdot)$ is an estimated subjective PDF.

In this paper, we use the Berkowitz (2001) statistics to test uniformity and independence jointly. Berkowitz (2001) tests the independence and normality of the Y_t by estimating the following AR(1) model and using a likelihood ratio test.

$$Z_t - \mu = \rho(z_{t-1} - \mu) + \varepsilon_t, \tag{17}$$

where $z_t = \Phi^{-1}(y_t)$ and Φ is the standard normal cumulative PDF.

Under the null hypothesis, the parameters in Eq. (17) should be: $\mu = 0$, $\rho = 0$, $\sigma_{\varepsilon_t} = 1$. Our algorithms for determining the parameters of utility functions are characterized by the following optimization problem.[15]

Power/exponential utility:

$$\min_{\alpha} LR3 = -2\left[L(0,1,0) - L\left(\hat{\mu}, \hat{\sigma}^2, \hat{\rho}\right)\right] \tag{18}$$

HARA/Log plus power/linear plus exponential utility:

$$\min_{\alpha,\beta} LR3 = -2\left[L(0,1,0) - L\left(\hat{\mu}, \hat{\sigma}^2, \hat{\rho}\right)\right], \tag{19}$$

where $L(\mu, \sigma^2, \rho)$ is the log-likelihood function from Eq. (17).

As noted by Berkowitz (2001) and Bliss and Panigirtzoglou (2004), the above $LR3$-statistics jointly test the uniformity and the independence of y_ts. Therefore, to examine the independence of y_ts separately, we also calculate the $LR1(= -2[L(\hat{\mu}, \hat{\sigma}^2, 0) - L(\hat{\mu}, \hat{\sigma}^2, \hat{\rho})])$ -statistics in each case. The statistical tests using $LR3$- and $LR1$-statistic are conducted based on their distributions under the respective null hypothesis. The Berkowitz $LR3$-statistic has a χ^2_3 distribution and $LR1$-statistic has a χ^2_1 distribution under the null hypothesis for fixed parameters, respectively. However, the process of searching for optimal $LR3$ values over parameter space distorts the distribution of the resulting $LR3$-statistic. In this paper, to investigate the statistical properties of our estimates under the null hypothesis, we use the Monte Carlo tests that were suggested by Bliss and Panigirtzoglou (2004). The results of the Monte Carlo tests will be discussed in Section 4.

By examining $LR3$-statistic and $LR1$-statistic, we can test whether the estimated PDFs from our models provide accurate forecasts of the true distributions. If $LR3$-statistic rejects the hypothesis that $z_t \sim i.i.d.$ N(0,1) and $LR1$-statistic fails to reject the hypothesis of serially uncorrelated z_t, it means that the estimated PDFs do not provide accurate forecasts of the true distributions.[16] If both statistics reject the hypothesis, we cannot determine whether estimated PDFs accurately forecast the true distributions or not. That is, we cannot determine whether the rejection is due to a lack of forecast ability or serial correlation. If both statistics fail to reject the hypothesis, we cannot reject the hypothesis that the estimated PDFs are good forecasts for the true distributions.

[15] Berkowitz (2001) and Bliss and Panigirtzoglou (2004) developed and suggested the algorithm in detail.

[16] Hypothesis that $z_t \sim i.i.d.$N(0,1) is equivalent to the hypothesis that $y_t \sim i.i.d.$U(0,1).

3. Data

For our empirical analysis, we use FTSE 100 index options traded on the London International Financial Futures Exchange (LIFFE), which are European-style index options and one of the most similar to the market portfolio.

The sample period for FTSE 100 index options extends from April 3, 1992 through July 7, 2000. The FTSE 100 index options expire on the third Friday of the expiration month. In terms of time to maturity, they are also listed for the quarterly cycle of March, June, September, December and the three additional consecutive near-term delivery months. The options prices used in this study are the LIFFE-reported settlement prices. For corresponding index levels, we use the actual or theoretical futures price reported by LIFFE for that contract.[17]

The risk-free interest rate used in this study is calculated from a collection of continuously compound zero-coupon interest rates at various maturities. Each zero rate is extracted from the zero curve, which is derived from BBA (British Bankers Association) Euro-dollar LIBOR rates.[18] For a given option, the appropriate risk-free interest rate that has a maturity equal to the option's expiration is obtained by linearly interpolating between the two closest zero rates on the zero curve.

For each delivery month, we select the cross sections corresponding to two weeks, four weeks and six weeks before the option's expiry. If no options are traded on the target observation date, we use options on the nearest trading date. For each target observation date, following filtering rules are used to construct our final samples. As the SML method to derive RN-PDFs cannot be reconciled with duplicate strikes in the data, ITM options that are relatively illiquid are discarded.[19] Options which don't satisfy general arbitrage restrictions and with have implied volatility of greater than 100 percent are excluded. To mitigate the impact of price discreteness on the option value, options with lower prices than two ticks are excluded from the sample. Finally, if the remaining option cross sections are less than five strikes, the entire cross section is excluded. Table 1 briefly shows the summary statistics of our sample filtered by the above rules.

4. Empirical results

4.1. Results for Monte Carlo simulations

While the Berkowitz *LR*3 statistic for fixed parameters is distributed $\chi^2(3)$, its distribution is distorted when searching over the parameter space to minimize it.[20] We therefore use the Monte Carlo simulations to derive empirical distributions for the *LR*3 statistics. As the procedure for simulations used in this study was fully introduced by Bliss and Panigirtzoglou (2004), it will only be briefly discussed. In the Monte Carlo tests, actual realizations of underlying asset prices are replaced with pseudo-realizations repeatedly

[17]LIFFE reports the actual or theoretical futures prices as the value of the underlying asset in their options data. As suggested by Bliss and Panigirtzoglou (2004), we use them as the corresponding index levels to compute the implied volatilities.

[18]To construct the zero curve, the BBA LIBOR rates for maturities 1 week, 1~2months are used. The curve for the longer term is not needed because the maturity of options in our sample is not more than 6 weeks.

[19]The liquidity of FTSE 100 index options is generally concentrated in ATM and OTM strikes.

[20]For more concrete analyses, see Bliss and Panigirtzoglou (2004).

Table 1
Summary statistics of option cross-sectional samples
The sample period for FTSE 100 index options extends from April 3, 1992 through July 7, 2000. Target observation dates are determined for horizons of two, four and six weeks before the option's expiry. To avoid duplicate strike prices, we only include the OTM call/put options, which are more frequently traded than ITM options. To mitigate the effects of price discreteness, we also exclude the options whose prices are not greater than two ticks.

Forecast horizon	Call/Put	# of cross-sections	Strikes per cross-sections		Average price
			Average	Range	
2 weeks	Call options	100	7.93	2–24	18.05
	Put options		10.50	3–26	19.23
	Total		18.43	6–45	18.72
4 weeks	Call options	100	9.89	3–24	29.06
	Put options		11.30	4–30	38.55
	Total		21.19	7–54	34.12
6 weeks	Call options	100	11.00	3–30	39.53
	Put options		11.42	5–28	55.54
	Total		22.42	9–58	47.68

drawn from the RN-PDFs. For each simulated set of outcomes, we estimate the parameters by minimizing the $LR3$ statistic as described in Section 2. One thousand replications were conducted and the resulting distributions of p-values and estimated RRA functions were examined.

Table 2 shows the mean, median, lower and upper quartiles of the $LR3$ p-values from the Monte Carlo test. From the table, we can easily observe that the distributions of the $LR3$ p-values are distorted. If our optimizing process does not affect the distributions, then the mean and the median of the resulting $LR3$ p-values should be 0.5. We can, however, find that the optimizing process overstates the $LR3$ p-values in our sample, regardless of the assumed functional form for utility functions. The $LR3$ p-values for the utility functions with two parameters are more overstated than those for the utility functions with one parameter. That is, the utility functions with more parameter are found to distort the distributions more severely. This finding means that even for the smaller $LR3$-values for the HARA, log plus power or linear plus exponential utility functions do not always guarantee higher $LR3$ p-values.[21] Fig. 1 plots the (5th, 10th, 25th, 50th, 75th, 90th, 95th) percentiles for the distributions of the estimated $LR3$ p-values and confirms the results of Table 2.

4.2. Explanatory power of the estimated PDFs

In Table 3, the Berkowitz p-values are presented. The corresponding p-values are adjusted by the empirical distributions from the Monte Carlo simulations in the previous section, and then reported in the table. As we expected, assuming more flexible utility

[21]Obviously, the more parameters we use, the lower $LR3$-values will be achieved in-sample. However, even if the tests are in-sample, adding more parameters does not necessarily produce the higher $LR3$ p-values.

Table 2
Results for estimated p-values from Monte Carlo simulations
Pseudo realizations are drawn from the actual FTSE 100- risk neutral PDFs and the p-values are obtained by minimizing Berkowitz likelihood ratio statistics over parameter space. Simulations are repeated 1,000 times for each horizon/utility function.

Forecast horizon	Assumed functional form for utility	Mean	Median	Quartile	
				Lower	Upper
2 weeks	Power	0.652	0.707	0.414	0.896
	Exponential	0.651	0.709	0.428	0.887
	HARA	0.690	0.743	0.485	0.918
	Log-power	0.699	0.752	0.494	0.932
	Linear-exponential	0.693	0.755	0.493	0.922
4 weeks	Power	0.632	0.674	0.388	0.880
	Exponential	0.631	0.673	0.398	0.880
	HARA	0.667	0.708	0.443	0.901
	Log-power	0.673	0.724	0.453	0.913
	Linear-exponential	0.669	0.714	0.452	0.913
6 weeks	Power	0.636	0.683	0.388	0.889
	Exponential	0.635	0.690	0.398	0.890
	HARA	0.672	0.741	0.456	0.908
	Log-power	0.684	0.758	0.462	0.915
	Linear-exponential	0.679	0.752	0.469	0.912

seems generally to improve the forecasting ability of the corresponding subjective PDFs. The forecasting ability of the RN-PDFs is found to be insignificant, regardless of forecast horizons. Moreover, the power and the exponential utility functions fail to produce good forecasts of the distributions of future realizations at both the 2 weeks and the 6 weeks horizon. Only the risk adjusted PDFs by the HARA, the log plus power and the linear plus exponential utility are found to be good forecasts of the distributions of future realizations at those horizons. At the 4 weeks horizon, while all the PDFs except for the RN-PDF are found to be good forecasts of the distributions, the PDFs by more flexible utility functions generally outperform those by the power and the exponential. To sum up, for the FTSE 100 contracts, our extension to the more flexible utility functions seems to succeed in providing better forecasts of future realizations and more reliable risk aversion functions.

Next, to show whether the incremental performance by risk-adjustment is significant, we conducted an out-of-sample test. The procedures for the out-of-sample test are as follows: First, we randomly divide our sample into two groups of the same number of observations, in-sample and test-sample. As our sample for each forecast horizon has 100 cross sections, each in-sample and test-sample should contain 50 cross sections. Second, from each in-sample, the parameters of the assumed utility functions are estimated by Eqs. (18) and (19). Third, the resulting utility functions are used to produce the corresponding subjective PDFs for the test-sample by Eqs. (9)–(13). Finally, we assess the forecasting ability of the PDFs by calculating the Berkowitz $LR3$ p-values for the test-sample. In-sample and test-sample are newly constructed per each test and one thousand replications are employed.

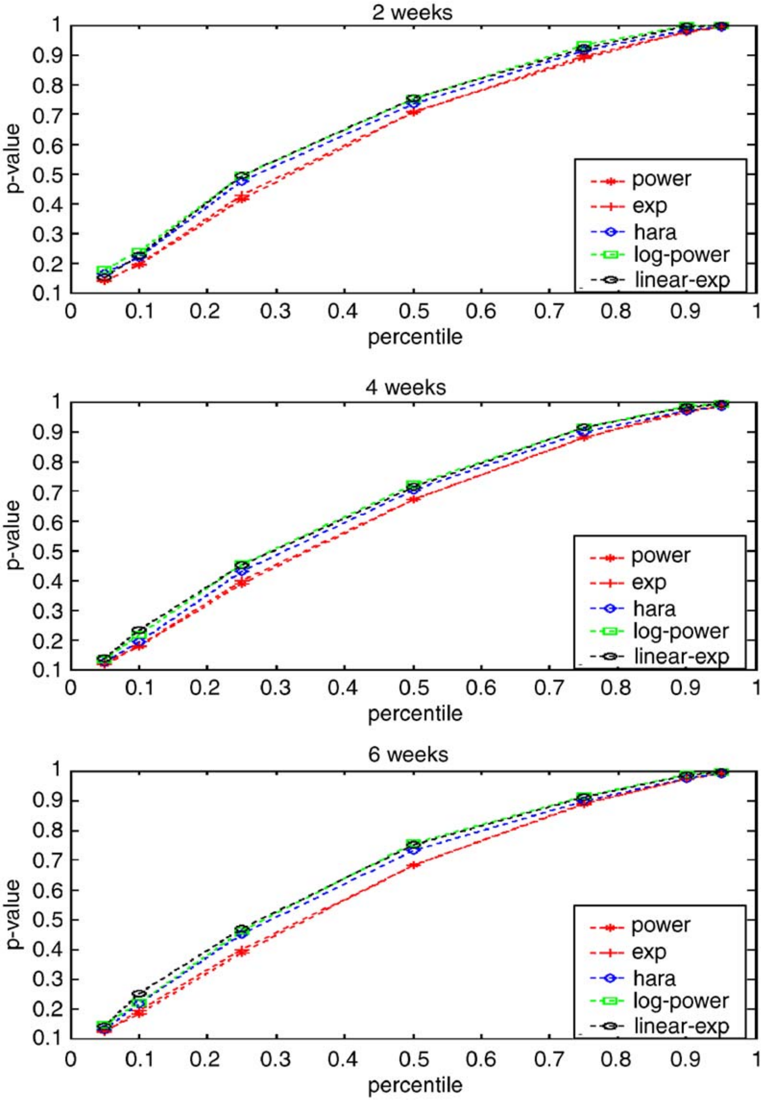

Fig. 1. Distribution of estimated p-values from Monte Carlo simulations. (5th, 10th, 25th, 50th, 75th, 90th, 95th) percentiles for the distributions of the estimated $LR3$ p-values from MC simulations are plotted. The lines with asterisks, crosses, diamonds, squares and circles are the lines corresponding to the power, exponential, HARA, log plus power and linear plus power utility functions, respectively.

Table 4 shows the results for the out of sample tests. We present the pair-wise differences in out-of-sample performance between (a) the RN-PDFs and the power utility-adjusted PDFs, (b) the RN-PDFs and the exponential utility-adjusted PDFs, (c) the power utility-adjusted PDFs and the HARA utility-adjusted PDFs, (d) the exponential utility-adjusted PDFs and the HARA utility-adjusted PDFs, (e) the power utility-adjusted

Table 3
Berkowitz statistic *p*-values

Using the modified Berkowitz tests, we test the ability of the risk neutral and subjective PDFs to forecast the future distributions of the underlying asset prices. The reported values are the *p*-values of Berkowitz likelihood ratio tests. All of the values, except for the case of risk neutral PDFs, have been adjusted for the bias resulting from optimizing the Berkowitz LR3-statistics over parameter space.

Forecast horizon	PDF	LR3 *p*-value	LR1 *p*-value
2 weeks	Risk neutral	0.016	0.553
	Power	0.025	0.591
	Exponential	0.024	0.595
	HARA	0.065	0.584
	Log-power	0.083	0.573
	Linear-exponential	0.083	0.579
4 weeks	Risk neutral	0.039	0.330
	Power	0.120	0.487
	Exponential	0.116	0.479
	HARA	0.205	0.500
	Log-power	0.212	0.508
	Linear-exponential	0.209	0.493
6 weeks	Risk neutral	0.034	0.259
	Power	0.042	0.159
	Exponential	0.042	0.162
	HARA	0.093	0.204
	Log-power	0.111	0.195
	Linear-exponential	0.118	0.205

PDFs and the log plus power utility-adjusted PDFs and (f) the exponential utility-adjusted PDFs and the linear plus exponential utility-adjusted PDFs. The first and the second column of the table present the mean and the standard deviations of the signed differences, respectively. If one PDF outperforms another PDF in forecasting ability, we expect a significantly non-zero mean of the signed differences. The third column of the table presents the number of experiments with positive differences. As the total number of experiments is 1000, we expect that the closer the number to 0 or 1000, the larger the differences between them. One thing kept in mind is that each test-sample is not independent of the others because of overlapping data. We cannot, therefore, carry out a rigorous statistical test to assess the differences in performances. However, the results reported in Table 4 cannot only give the differences in performance between the PDFs approximately, but also measure the relative improvement by each assumed utility function.

From Table 4, we can find that the risk-adjusted PDFs consistently outperform the RN-PDFs, regardless of assumed utility functions and forecast horizons. Especially, in terms of the degree of improvement, the difference in performance between the power (or exponential) utility- and the more flexible utility-adjusted PDFs is found to be greater than that between the RN-PDFs and the power (or exponential) utility-adjusted PDFs. These findings from the out of sample tests confirm that we can significantly improve the forecasting ability by assuming more flexible utility functions.

Table 4
Results for out of sample tests

Our sample data were randomly divided into two groups of the same number of observations, in-sample and test-sample. For each in-sample, the parameters of the assumed utility functions are estimated by minimizing the Berkowtiz LR3-statistic. After deriving the RN-PDFs and the subjective PDFs with the estimated parameters and assumed models, we assess the forecasting ability of the PDFs for the test-sample. Tests are repeated 1,000 times by randomly re-constructing the in- and test-samples. The reported values under the table are the means and the standard deviations of signed differences between the Berkowitz LR3 p-values for the respective PDFs. Signed differences are defined by $(a_i–b_i)$, $i = 1, 2, \ldots, 1000$, where a_i and b_i is the Berkowitz LR3 p-value by the corresponding PDF, respectively. We also report the number of tests with positive differences. For example, in the case of 2 weeks / 'Power-RND', the power utility-adjusted PDFs outperform the RN-PDFs, 628 times out of 1,000.

Forecast horizon	Difference between PDFs	Mean	Standard deviation	# (>0)
2 weeks	Power - RND	0.044	0.123	628
	Exp - RND	0.040	0.120	611
	HARA - Power	0.026	0.040	813
	HARA - Exp	0.034	0.052	734
	Log.Power - Power	0.028	0.034	822
	Linear.Exp - Exp	0.037	0.060	745
4 weeks	Power - RND	0.122	0.174	806
	Exp - RND	0.110	0.167	798
	HARA - Power	0.049	0.028	980
	HARA - Exp	0.053	0.034	987
	Log.Power - Power	0.056	0.040	963
	Linear.Exp - Exp	0.072	0.043	972
6 weeks	Power - RND	0.108	0.205	761
	Exp - RND	0.095	0.195	762
	HARA - Power	0.058	0.046	944
	HARA - Exp	0.084	0.055	979
	Log.Power - Power	0.063	0.057	912
	Linear.Exp - Exp	0.090	0.065	975

4.3. The Arrow-Pratt relative risk aversion (RRA) functions and equity risk premium

Table 5 reports the mean and the variations of the Arrow-Pratt RRAs. The Arrow-Pratt RRA values for the exponential, HARA, log plus power and linear plus exponential utility functions are computed using the RRA values on realized return states at each expiry. The mean of RRA values is found to be much greater than zero, regardless of assumed utility functions. In their variations, the RRA values for our flexible utility functions are much more variable than those for either the power or exponential utility functions.

Fig. 2 plots the Arrow-Pratt RRA functions across gross return states. We present the corresponding RRA functions for each forecast horizon in each panel of the figure, respectively. The dotted lines are the RRA functions derived under the assumption of the power and the exponential utility. The lines with diamond, square and circle denote the RRA functions derived under the assumption of the HARA, the log plus power and the linear plus exponential utility, respectively. As expected by the results of forecasting ability from the previous section, the Arrow-Pratt RRA (or ARA) functions after the restriction on the utility function is relaxed are found not to be constant across wealth.

Table 5
Estimates for relative risk aversion
For the power utility function, reported values are the coefficients of RRA and do not vary across wealth. However, the RRAs for other utility functions vary with the level of wealth and thus the mean and the variation of the RRAs on realized return states for each sample are reported.

Forecast horizon	Assumed functional form for utility	Mean	Variation	
			Standard deviation	Range
2 weeks	Power	2.583	—	—
	Exponential	2.467	0.066	2.337~2.732
	HARA	2.768	0.121	2.374~3.044
	Log-power	3.068	0.364	1.943~3.907
	Linear-exponential	2.872	0.454	1.320~3.836
4 weeks	Power	4.090	—	—
	Exponential	4.033	0.154	3.681~4.415
	HARA	4.135	0.221	3.676~4.739
	Log-power	3.990	0.669	2.591~5.719
	Linear-exponential	3.829	0.744	2.086~5.500
6 weeks	Power	2.699	—	—
	Exponential	2.608	0.128	2.314~3.038
	HARA	2.773	0.155	2.379~3.215
	Log-power	2.937	0.659	1.551~4.874
	Linear-exponential	2.766	0.817	0.750~4.817

When we assume either the power or exponential utility functions, only the constant RRA or constant ARA functions are derived. However, we obtain the 'declining RRA functions' by assuming more flexible utility functions with two parameters. The decreasing RRA functions presented in Fig. 2 are observed regardless of the assumed utility functions with two parameters and the forecast horizons.[22] These findings are consistent with the results in the forecasting ability of the estimated PDFs, which are presented in the previous section. In the previous section, we observed that the subjective PDFs by more flexible utility functions with two parameters outperform those by the power and exponential utility functions as well as the RN-PDFs, regardless of the forecast horizon. The RN-PDFs and the subjective PDFs by the power and exponential utility functions generally fail to provide an accurate forecast of the distributions of future realizations, but the subjective PDFs by more flexible utility functions with two parameters were proven to forecast them accurately. Moreover, the results of the out-of-sample test presented in Table 4 also give strong evidence for the reliability of the recovered 'declining RRA functions' from this market. Taking all these results into account, we can conclude that the Arrow-Pratt RRA functions for the FTSE 100 are 'declining' rather than 'constant' or 'inclining'.

Compared to the results in previous works such as Jackwerth (2000) and Ait-Sahalia and Lo (2000), we cannot observe empirical anomalies such as the 'U-shaped' RRA functions ('implied risk aversion smile') in the FTSE 100 index options market. We only find that the investors trading the FTSE 100 index options exhibit a 'declining RRA' across wealth.

[22]The specific restrictions on each utility functions with two parameters make only a little difference in the slope and the curvature of the derived RRA functions.

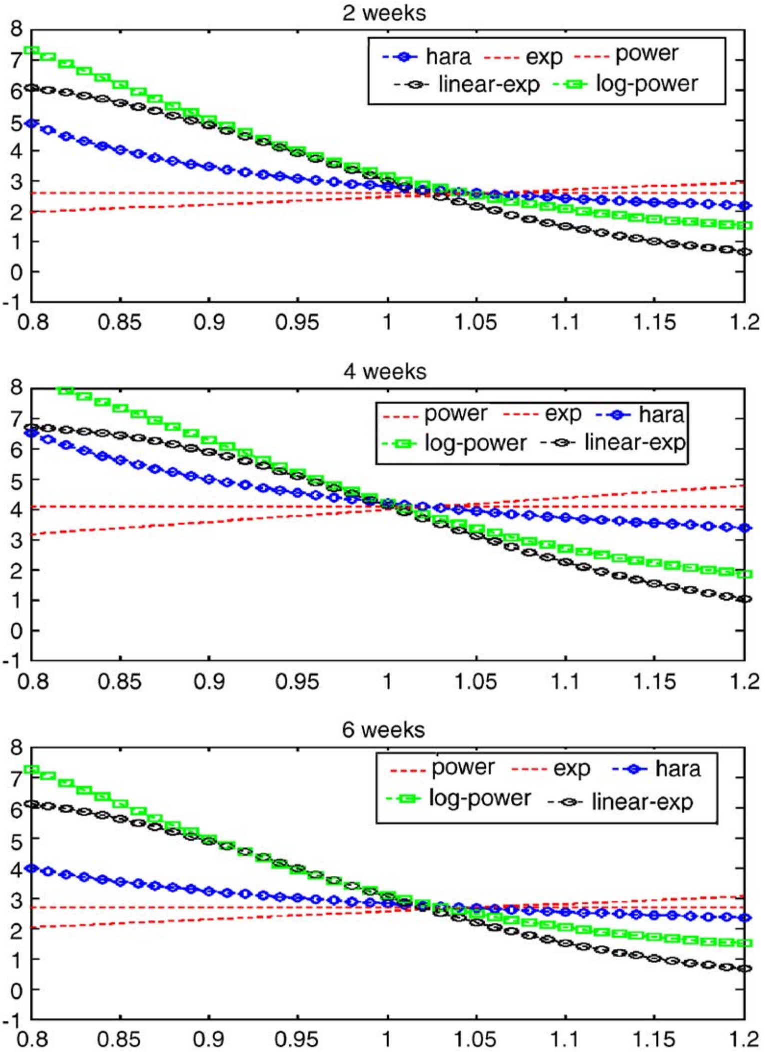

Fig. 2. Option-implied relative risk aversion. The Arrow-Pratt measures of relative risk aversion (RRA) on gross return states of from -20% to 20% are plotted. The dotted lines are the RRA functions for the power and exponential utility functions. The lines with diamonds, squares and circles are the RRA functions for the HARA, log plus power and linear plus exponential utility functions, respectively.

However, the relatively steep slope of RRA functions in our work leave some room for further discussions.

Table 6 shows the mean and the variations of estimated equity risk premiums from our samples. The equity risk premium in our work is approximately measured by the difference between the mean of the subjective and risk neutral PDFs. The mean and the variability of the equity risk premiums are found to be generally higher in the case of flexible utility functions than in the case of power (or exponential) utility functions. From the time series of derived equity risk premiums, we also find that extremely high equity risk premium

Table 6
Estimates for expected equity risk premium
Expected equity risk premium is approximately measured by the difference between the mean of the risk neutral PDF and the corresponding subjective PDF. The reported risk premium here is the annualized value.

Forecast horizon	Assumed functional form for utility	Mean (%)	Variation (%)	
			Standard deviation	Range
2 weeks	Power	9.291	7.987	2.840~61.971
	Exponential	8.702	7.739	2.688~57.193
	HARA	10.568	9.612	3.090~73.619
	Log-power	12.177	11.190	3.127~83.381
	Linear-exponential	11.372	10.322	2.534~77.861
4 weeks	Power	13.931	9.926	3.795~61.205
	Exponential	13.545	9.604	3.693~59.107
	HARA	14.482	10.516	3.902~65.237
	Log-power	14.537	10.575	3.911~64.928
	Linear-exponential	13.917	9.997	3.805~61.381
6 weeks	Power	9.389	7.015	2.759~42.560
	Exponential	8.890	6.589	2.620~39.956
	HARA	9.949	7.642	2.884~46.661
	Log-power	11.156	8.616	3.197~51.321
	Linear-exponential	10.596	8.026	3.109~47.692

amounting to 40% to about 70% prevailed from the mid-1998 through the beginning of 1999.

5. Conclusion

This paper has extended the methodology for deriving investors' risk preferences from option prices. Instead of assuming a well-behaved functional form for the underlying utility functions such as the power or the exponential, we assume wider classes of utility functions: the HARA, the log plus power and the linear plus exponential utility.

The main advantage of our approach is its generality. Under our extension to the more flexible functional form for the utility functions, various advantages from previous research on deriving investors' risk preferences can be obtained. As in the methodology assuming the simple functional form for utility functions, subjective PDFs obtained by our extended models can be time-varying and have a forward-looking property. Furthermore, as in the methodology assuming the stationarity of subjective PDFs, newly introduced utility functions in this study can produce a more flexible structure for the risk preferences across wealth than the previously assumed narrow classes of utility functions. On the other hand, one of the limitations we should be aware of is that the selection of candidate utility functions in our approach is not supported by an underlying theory and thus is arbitrary to some extent.

To estimate the risk neutral PDFs from option prices, the non-parametric smoothed implied volatility smile (SML) method is used. To test the explanatory power of derived PDFs, we use the Berkowitz LR test statistics. For our empirical analysis, investors' risk preferences implied by FTSE 100 index options in UK market are examined.

To summarize our empirical results, assuming more flexible utility functions increases the forecasting ability of the estimated PDFs and thus provides more reliable Arrow-Pratt RRA functions than a simple power or exponential utility function. The RRA functions in this market are found to be declining across wealth rather than constant or inclining, regardless of forecast horizons. The out-of-sample tests also confirm the robustness of our findings.

References

Ait-Sahalia, Y., Lo, A., 2000. Nonparametric risk management and implied risk aversion. Journal of Econometrics 94, 9–51.

Bahra, B., 1997. Implied risk-neutral probability density functions from option prices: theory and application. Working paper, Bank of England.

Bates, D.D., 1991. The Crash of '87: was it expected? The evidence from option markets. Journal of Finance 46, 1009–1044.

Bell, D.E., 1988. One-switch utility functions and a measure of risk. Management Science 34, 1416–1424.

Berkowitz, J., 2001. Testing density forecasts with applications to risk management. Journal of Business and Economic Statistics 19, 465–474.

Bick, A., 1990. On viable diffusion price processes of the market portfolio. Journal of Finance 45, 673–689.

Bliss, R.R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. Journal of Banking and Finance 26, 381–422.

Bliss, R.R., Panigirtzoglou, N., 2004. Option implied risk aversion estimate. Journal of Finance 59, 407–446.

Brunner, B., Hafner, R., 2003. Arbitrage-free estimation of the risk-neutral density from the implied volatility smile. Journal of Computational Finance 7, 75–106.

Campa, J.M., Chang, P.H.K., Reider, R.L., 1998. Implied exchange rate distributions: evidence from OTC option markets. Journal of International Money and Finance 17, 855–880.

Cooper, N., 1999. Testing techniques for estimating implied RNDs from the prices of european-style options. Working paper, Bank of England.

Coutant, S., 1999. Implied risk aversion in option prices using Hermite polynomials. Article presented at the BIS workshop in London.

Craven, P., Wahba, G., 1979. Smoothing noisy data with spline function. Numerishe Matematik 31, 377–403.

Gelles, G.M., Mitchell, D.W., 1999. Switching, mean-seeking, and relative risk with two or more risky assets. Working paper, University of Missouri.

Giamouridis, D., 2003. Inferring investors' risk preferences from option prices and asset returns. Working paper, Athens University.

Hastie, T.J., Tibshirani, R.J., 1990. Generalized Additive Models. Chapman & Hall, New York.

Jackwerth, J.C., 1999. Option-implied risk-neutral distributions and implied binomial trees: a literature review. Journal of Derivatives Winter, 66–82.

Jackwerth, J.C., 2000. Recovering risk aversion from option prices and realized returns. Review of Financial Studies 13, 433–451.

Kallberg, J.G., Ziemba, W.T., 1983. Comparison of alternative utility functions in portfolio selection problems. Management Science 29, 1257–1276.

Kim, I.J., Kim, S., 2003. On the usefulness of implied risk neutral distributions—evidence from Korean KOSPI 200 index options market. Journal of Risk 6.

Kliger, D., Levy, O., 2002. Risk preferences heterogeneity: evidence from asset markets. European Finance Review 6, 277–290.

Leland, H., 1980. Who should buy portfolio insurance. Journal of Finance 35, 581–594.

Luis, J.B., 2001. Expectation and risk aversion measures implicit in prices of options on US and European stock indexes. Article presented at the Forecasting Financial Market 2001 Conference in London.

Pedersen, C.S., Satchell, S.E., 2001. On the characterization of investor preferences by changes in wealth. The Geneva Papers on Risk and Insurance Theory 26, 175–193.

Perignon, C., Villa, C., 2002. Extracting information from options markets: smiles, state-price densities and risk aversion. European Financial Management 8, 495–513.

Rosenberg, J.V., Engle, R.F., 2002. Empirical pricing kernels. Journal of Financial Economics 64, 341–372.

Shimko, D., 1993. Bounds of probability. Journal of Risk 6, 34–37.

Syrdal, S.A., 2002. A study of implied risk-neutral density functions in the Norwegian option market. Working paper, Norges Bank.

Torkamani, J., Haji-Rahimi, M., 2001. Evaluation of farmer's risk attitudes using alternative utility functional forms. Journal of Agricultural Science and Technology 3, 243–248.

Ziegler, A., 2003. Why does implied risk aversion smile? Working paper, University of Lausanne.