Option-Implied Preferences Adjustments and Risk-Neutral Density Forecasts. The Evidence for the Ibex 35

Francisco Alonso

(Banco de España)

Roberto Blanco

(Banco de España)

and

Gonzalo Rubio

(Universidad del País Vasco)

June 2005

Gonzalo Rubio acknowledges the financial support from Fundación BBVA research grant 1-BBVA 00044.321-15466/2002. The contents of this paper are the sole responsibility of the authors.

ABSTRACT

The main objective of this paper is to analyse the value of information contained in prices of options on the IBEX 35 index at the Spanish Stock Exchange Market. The forward looking information is extracted using implied risk-neutral density functions estimated by a mixture of two-lognormals and three alternative risk-adjustments: the classic power and exponential utility functions and a habit-based specification that allows for a counter-cyclical variation of risk aversion. Our results show that at four-week horizon we can reject the hypothesis that between October 1996 and March 2000 the risk-neutral densities provide accurate predictions of the distributions of future realisations of the IBEX 35 index at a four-week horizon. When forecasting through risk-adjusted densities the performance of this period is statistically improved and we no longer reject that hypothesis. All risk-adjusted densities generate similar forecasting statistics. Then, at least for a horizon of four-weeks, the actual risk adjustment does not seem to be the issue. By contrast, at the one-week horizon risk-adjusted densities do not improve the forecasting ability of the risk-neutral counterparts.

Key words: risk-adjustment, option-implied densities, forecasting performance, risk aversion, Ibex 35.

JEL: G10, G12.

1. Introduction

Prices of European exchange-traded options on stock indices implicitly contain the risk-neutral density (RND hereafter) which is a key component for risk-neutral valuation. In this context, prices are the present value at the risk-free rate of their expected payoffs calculated under the RND. When the market is dynamically complete it is well known that the RND can be recovered from the corresponding option prices using the insights on Breeden and Litzenberger (1978). In particular, the RND is proportional to the second derivative of the option pricing function with respect to the exercise.

Noting that option prices should capture forward-looking distributions of the underlying assets, academic researchers and central banks have used implied RNDs to proxy the market expectations of the distribution of the underlying asset or to forecast future outcomes. They have the advantage relative to other historical time-series data that they are taken from a single point in time when looking toward expiration. Hence, they should be more responsive to changing expectations than competing alternatives. Unfortunately, in practice, there is no a continuum of exercise prices. Neither very low nor high exercises are available and, in any case, they are set at discrete intervals by market officials. This complicates the estimation of RND and, not surprisingly, numerous alternative parametric and non-parametric methods have been proposed in literature. Moreover, the existence of risk aversion means that RNDs will probably differ from the actual density from which realizations of returns are drawn.

Interestingly most of these papers are concerned with the estimation and the ex-post assessment of the alternative RNDs as a way of forecasting the actual realizations of the underlying asset at expiration. Bliss and Panigirtzoglou (2002) and Bondarenko (2003) compare several competing procedures and conclude that nonparametric methods based on either the smoothed (spline) implied volatility smile and the positive convolution approximation seem to dominate the two-lognormal approach and other parametric techniques when estimating RNDs. Moreover, Anagnou, Bedendo, Hodges and Tompkins (2003) for the UK option market, Craig, Glatzer, Keller and Scheicher (2003) for the German stock option data, Bliss and Panigirtzoglou (2004) for the US and the UK option data, and Alonso, Blanco and Rubio (2005) for the Spanish option prices conclude that the RND is not an unbiased estimator of actual probability density function. This may not be surprising given the risk-neutrality embedded in these estimates. In other words, these papers suggest that the forecasting differences arise from the risk aversion of the representative investor. In fact, by imposing a stationary utility function (a stationary risk aversion parameter), Anagnou, Bedendo, Hodges and Tompkins (2003) and Bliss and Panigirtzoglou (2004) test whether either power or exponential utility functions are improved forecasters of future values of the underlying. In general, they are not able to reject the null that implied risk-adjusted densities are unbiased forecasts of future outcomes. However, Bliss and Panigirtzoglou obtain a disturbing result in the sense that the implicit risk aversion parameter they estimate increases as market risk declines. This suggests a misspecification of the utility functions imposed in their paper and, as in the asset pricing literature, it points out towards alternative utility functions with habit persistence and where the risk aversion parameter is not theoretically linked to the elasticity of intertemporal substitution.

This paper investigates the forecasting power of RNDs for alternative horizons of one and four weeks using exchange data of the European future options contract on the Spanish IBEX-35 index. Given the previous evidence reported by Alonso, Blanco and Rubio (2005) in which the predicting ability of RNDs estimated either by a mixture of two log-normals or splines is indistinguishable, we only report results from the parametric log-normal case. Moreover, given the estimation of a RND from a cross-section of option prices with a given maturity, this work obtains the implicit risk adjustment that makes the subjective density forecasts of the agents to be the best assessment of the objective or physical densities from which the realizations are actually drawn. We employ three competing utility functions. As in the previous two existing papers, we assume a constant relative risk aversion power utility function and an exponential utility function, and derive the implicit parameters assuming that their value is stationary over the sample period but using, of course, the time-varying RNDs estimated with our option pricing data. Additionally, as our key contribution to literature, we also impose a utility function that incorporates the possibility that the price of risk varies counter-cyclically over time, a key characteristic to explain both the cross-section of asset prices and the counter-cyclical variation of the expected equity risk premium. This is accomplished in a model where the representative investor displays habit-formation which means that a positive effect of today´s consumption on tomorrow´s marginal utility of consumption exists.

The utility function is a power function to maintain the property of scale-invariant[1]. In particular, it is a power function of the difference between consumption and habit, where habit is a slow-moving linear (nonlinear) average of past aggregate consumption. Our utility specification is based on the model proposed by Campbell and Cochrane (1999) where habit is external and reacts only gradually to changes in consumption[2]. The key idea is that this utility makes agents more risk-averse in bad times, when consumption is low relative to its past history, than in good times, when consumption is high relative to its past history. To obtain this time-varying risk aversion, Campbell and Cochrane model utility as the difference between consumption and habit, rather than taking the ratio between the two. Thus, the behavior of market volatility is explained by a small consumption risk, amplified by variable risk aversion, while the equity premium is explained by high market volatility, together with a high average level of risk aversion. Interestingly, their habit model is also able to keep low both the long-term mean and the volatility of interest rates. Following the literature on option implied risk-adjustments we proxy consumption by the stock index level.

The results of this paper show that between 1996 and 2004 we cannot reject the hypothesis that the RNDs provide accurate predictions of the distributions of future realizations of the IBEX 35 index at four-week horizon. However, tests based on the tails of the distribution show that RNDs significantly understated the right tail of the distribution. Moreover, this predicting ability is not robust by subperiods. More specifically, and using a four-week horizon, we find that RNDs are not able to consistently predict the realizations of returns from October 1996 to February 2000. These results suggest that the ability of RNDs to forecasts future realizations might possibly be improved if risk preference adjustments were introduced. Indeed, this is the case. Once risk-adjustments are made we are not able to reject the hypothesis that risk-adjusted densities contain good predictions of the distributions of future realizations of the underlying index. Moreover, for the whole period and contrary to the RNDs, the risk-adjusted densities do not understate the right-tails of the distribution. Interestingly, this latter result is not maintained during the first sub-period. However, despite the fact

[1] This basically means that the risk premium does not change over time as aggregate wealth and the scale of the economy increase.
[2] In the external habit models, habit depends on aggregate consumption which is unaffected by one agent´s decisions.

that the risk aversion estimates from the habit-based utility function is clearly counter-cyclical and consistently higher than estimates from either power or exponential utility, the predicting ability of all three specifications are quite similar.

The results using a one-week horizon are practically the same independently either employing RNDs or risk-adjusted densities. This probably makes sense. At high frequencies, it is well known from the first order condition of the optimization problem of the representative agent that pricing of risky assets may easily be consistent with a linear utility function or risk-neutralilty. Since consumption and risk aversion do not change much from week to week, we might expect that prices are well approximated as random walks. Of course, this result is very different once we allow for a longer investment horizon, where payoffs are scaled by the marginal rate of substitution of consumption. This is the intuition probably reflected in our empirical results.

This paper is organized as follows. Section 2 briefly discusses how we estimate RNDs, while in Section 3 we present the testing procedures to assess the forecasting ability of our densities to check if they conform to the actual densities from which realisations are drawn. Section 4 discusses option-implied preferences adjustments, Section 5 contains the description of the data set used in the paper, and Section 6 reports the empirical results using alternative risk-neutral and risk-adjusted densities. Conclusions follow in Section 7.

2. Estimating Risk-Neutral Densities

Prices of European call options at time t on the underlying asset P with expiration at $t+\tau$ and strike prices K are given by the well known expression[3]:

$$c(t,\tau,K) = e^{-r\tau} \int_{K}^{\infty} q_{t,\tau}\left(P_{t+\tau}\right)\left(P_{t+\tau} - K\right) dP_{t+\tau} \tag{1}$$

where $q_{t,\tau}\left(P_{t+\tau}\right)$ is the risk-neutral probability density function for the value of the underlying asset at time $t+\tau$. As pointed out by Breeden and Litzenberger (1978), if

[3] The same reasoning can of course be done in terms of put options.

we differentiate twice (1) with respect to K we obtain the risk-neutral probability density function

$$\frac{\partial^2 c(t,\tau,K)}{\partial K^2} = e^{-r\tau} q_{t,\tau}\left(P_{t+\tau}\right) \tag{2}$$

Given the similarities found in our own previous empirical evidence with the Ibex 35 index between parametric and non-parametric estimation methods, this paper employs the two-lognormal mixtures of Melick and Thomas (1997) as the way to estimate RNDs,

$$q_{t,\tau}\left(P_{t+\tau}\right) = \theta \log N\left(\alpha_1,\beta_1;P_{t+\tau}\right) + \left(1-\theta\right) \log N\left(\alpha_2,\beta_2;P_{t+\tau}\right) \tag{3}$$

where $\log N\left(\alpha_i,\beta_i;P_{t+\tau}\right)$ is the i^{th} lognormal density with parameters α_i and β_i:

$$\alpha_i = \ln P_t + \left(\mu_i - \frac{1}{2}\sigma_i^2\right)\tau \quad ; \quad \beta_i = \sigma_i\sqrt{\tau} \quad ; \quad i = 1,2 \tag{4}$$

and where μ_i and σ_i are, respectively, the mean and standard deviation of associated normal distributions, and the stochastic process is based on two states with different first and second moments, governed by the weights θ and $1-\theta$ for $0 \le \theta \le 1$. Thus, this is a flexible specification for the RND that is able to capture skewness and excess kurtosis and allows for a rich and wide range of shapes including bi-modal distributions, which would appear if, for example, market participants are placing a high weight on an extreme move in the underlying price but are unsure of its direction.

Plugging this mixture of two-lognormals into equation (1) we can obtain theoretical prices for both calls and puts. Then, the numerical estimation of the five parameters, $\alpha_1,\beta_1,\alpha_2,\beta_2,\theta$, is obtained by minimizing the squared pricing error as defined by the difference between the theoretical and observed option prices:

$$\min_{\{\alpha_1,\alpha_2,\beta_1,\beta_2,\theta\}} \left\{ \sum_{j=i}^{N_j} \left[c_j\left(t,\tau,K_j\right) - c_j^m \right]^2 + \sum_{h=i}^{N_h} \left[p_h\left(t,\tau,K_h\right) - p_h^m \right]^2 \right\} \tag{5}$$

subject to $\beta_1,\beta_2 > 0$ *and* $0 \le \theta \le 1$, and where N_j, N_h, c_j^m and p_h^m stand respectively for number of calls, number of puts, market price of call *j* and market price of put *h*[4].

3. Testing the Forecasting Performance of Probability Density Functions

To study the predicting ability of estimated probability density functions (PDFs hereafter), both risk-neutral PDFs and risk-adjusted PDFs, we first employ a method based on the relationship between the data generating process (the true density function), $f_{t,\tau}\left(P_{t+\tau}\right)$, and the estimated sequence of density forecasts, $q_{t,\tau}\left(P_{t+\tau}\right)$, as related through the probability integral transform, $z_{t,\tau}$, of the realization of the process taken with respect to the density forecast, where τ represents the forecasting horizon. In other words, each cross-section of options at time *t* for a given time-to-expiration τ produces an estimated PDF, $q_{t,\tau}\left(P_{t+\tau}\right)$. We want to test the hypothesis that our estimated $q_{t,\tau}\left(P_{t+\tau}\right)$ are equal to $f_{t,\tau}\left(P_{t+\tau}\right)$. Note of course that we have an estimated PDF for a given expiration and only one realization, $P_{t+\tau}$, is available on a given date and for that particular expiration. The probability integral transform is defined as

$$z_{t,\tau} = \int_{-\infty}^{P_{t+\tau}} q_{t,\tau}\left(u\right) du = Q_{t,\tau}\left(P_{t+\tau}\right) \tag{6}$$

Hence, $z_{t,\tau}$ is equal to the probability value of the estimated cumulative density function, $Q_{t,\tau}(.)$, τ days ahead at the realization of the underlying on day $t+\tau$, $P_{t+\tau}$.

[4] Note that the mean of a RND is the future price. Some papers include in equation (5) the difference between the future price and the expected value of the underlying asset at t+τ. In our sample the impact on the estimated parameters of the introduction of this additional term is negligible.

We basically integrate up to the realization of the underlying at each date and statistically check if the resulting probabilities are drawn from the estimated PDFs.

As shown by Diebold, Gunther and Tay (1998), under independence and if the forecasts and the true densities coincide, then the sequence of the probability integral transforms, $z_{t,\tau}$, is uniformly distributed as $U(0,1)$. Berkowitz (2001) proposes a parametric approach for jointly testing uniformity and independence. In particular, a further transformation, $x_{t,\tau}$, of the inverse probability transform, $z_{t,\tau}$, is defined using the inverse of the standard normal cumulative density function, $N(.)$:

$$x_{t,\tau} = N^{-1}\left(z_{t,\tau}\right) = N^{-1}\left(\int_{-\infty}^{P_{t+\tau}} q_{t,\tau}\left(u\right) du \right) \tag{7}$$

under the null, $q_{t,\tau}\left(P_{t+\tau}\right) = f_{t,\tau}\left(P_{t+\tau}\right)$, $x_{t,\tau}$ has an independent and identically distributed N(0,1). In order to estimate the independence and standard normality of the $x_{t,\tau}$, Berkowitz suggests the following autoregressive model[5]

$$x_{t,\tau} - \mu = \rho\left(x_{t-1,\tau} - \mu\right) + \varepsilon_{t,\tau} \tag{8}$$

which is estimated using maximum likelihood and then testing the corresponding restrictions by a likelihood ratio test. The log-likelihood function, $L\left(\mu,\sigma^2,\rho\right)$, associated with the model in (12) is given by

$$L\left(\mu,\sigma^2,\rho\right) = -\frac{1}{2}log\left(2\pi\right) - \frac{1}{2}log\left[\frac{\sigma^2}{\left(1-\rho^2\right)}\right] - \frac{\left[x_{t,\tau} - \mu/(1-\rho)\right]^2}{2\sigma^2/\left(1-\rho^2\right)} - \frac{n-1}{2}log\left(2\pi\right) - \frac{n-1}{2}log\left(\sigma^2\right) - \sum_{t=2}^{T}\left[\frac{\left(x_{t,\tau} - \mu(1-\rho) - \rho x_{t-1,\tau}\right)^2}{2\sigma^2}\right] \tag{9}$$

[5] Berkowitz (2001) shows that higher order autoregressive processes results in increasing the number of parameters and reduced power. Also, Bliss and Panigirtzoglou (2004) compares alternative tests and conclude that the Berkowitz tests is more reliable in small samples.

Note that, under the assumptions of the model, the parameters should be equal to $\mu = \rho = 0$ and $\sigma^2(\varepsilon_{t,\tau}) = 1$. Then, the likelihood ratio statistic, $LR = -2\left[L(0,1,0) - L\left(\hat{\mu}, \hat{\sigma}^2, \hat{\rho}\right)\right]$, is distributed as $\chi^2(3)$ under the null hypothesis.

When the available data implies that we have to test overlapping forecasts, a potential rejection may be due to the overlapping nature of the data, which may produce autocorrelation. Berkowitz also proposes to test the independence assumption separately by the alternative likelihood ratio statistic given by $LR(i) = -2\left[L\left(\hat{\mu}, \hat{\sigma}^2, 0\right) - L\left(\hat{\mu}, \hat{\sigma}^2, \hat{\rho}\right)\right]$ which is distributed as $\chi^2(1)$ under the null hypothesis.

As explained by Bliss and Panigirtzoglou (2004), if *LR* rejects the hypothesis, failure to reject *LR(i)* provides evidence that the estimated PDFs are not producing accurate forecasts of the true density. However, if both *LR* and *LR(i)* reject, it is not possible to conclude if there is lack of predicting ability or serial correlation. Finally, failure to reject both *LR* and *LR(i)* would be consistent with forecasting capacity.

Unlike most previous papers testing the forecasting ability of PDFs, we not only want to test the performance of the whole body of the distribution, but also analyze the performance of the tails of the distribution. We follow Anagnou, Bedendo, Hodges and Tompkins (2003) in employing the scoring rules based on the distance between the forecasted probability mass, $q_{t,\tau}^{tail}$, in a given tail and a binary variable, $R_{t,\tau}$, which takes the value of 1 if the actual realization of the underlying falls in the tail, and 0 otherwise. The so called Brier score is given by

$$B = \frac{1}{T} \sum_{t=1}^{T} 2\left(q_{t,\tau}^{tail} - R_{t,\tau}\right)^2 \tag{10}$$

which takes values between 0 and 2 and a better performance is captured by smaller values for the score. To test if it departs from its expected value, $\sum_{t=1}^{T} q_{t,\tau}^{tail}\left(1 - q_{t,\tau}^{tail}\right)$, the following statistic, suggested by Seillier-Moiseiwisch and Dawid (1993) is employed:

$$ASN = \frac{\sum_{t=1}^{T}\left(1 - 2q_{t,\tau}^{tail}\right)\left(R_{t,\tau} - q_{t,\tau}^{tail}\right)}{\left[\sum_{t=1}^{T}\left(1 - 2q_{t,\tau}^{tail}\right)^{2} q_{t,\tau}^{tail}\left(1 - q_{t,\tau}^{tail}\right)\right]^{\frac{1}{2}}} \tag{11}$$

which is asymptotically distributed as a standard normal.

4. Option-Implied Preferences Adjustments

Given the lack of consistently adequate forecasting ability of RNDs throughout the sub-periods analyzed by Alonso, Blanco and Rubio (2005), and in order to incorporate risk aversion into the analysis along the lines of Bliss and Panigirtzoglou (2004), we first assume that the risk aversion function is characterized by either a power or an exponential utility functions.

Given an estimation of a RND from a cross-section of option prices on the index, we look for the implied preferences that force the preference-adjusted density to be as close as possible to the distribution of realizations of the underlying as defined by the Berkowitz test statistic. In other words, we choose the preference parameters from the assumed utility function to maximize the predicting ability of the estimated density by maximizing over the p-value of the Berkowitz *LR* statistic.

In contrast with previous papers, we also assume a habit-based utility function to extend the simple power and exponential utility functions employed in literature. In particular, we employ a utility function that allows us to separate the coefficient of relative risk aversion and the elasticity of intertemporal substitution. Moreover, and probably more relevant in our case, this utility function incorporates the possibility that the price of risk varies over time. This is accomplished in a model where agents displays habit-formation making investors more risk-averse in bad times, when consumption is low relative to its past history, than in good times, when consumption is high relative to its past history.

Under dynamically complete and frictionless markets, the objective density function, $f_{t,\tau}\left(P_{t+\tau}\right)$, is related to the risk-neutral density function, $q_{t,\tau}\left(P_{t+\tau}\right)$, by the marginal rate of substitution of the representative investor as discussed by Jackwerth (2000) and Aït-Sahalia and Lo (2000):

$$\frac{f_{t,\tau}\left(P_{t+\tau}\right)}{q_{t,\tau}\left(P_{t+\tau}\right)} = \lambda \frac{U'\left(P_{t+\tau}\right)}{U'\left(P_t\right)} \equiv \zeta\left(P_{t+\tau};P_t\right) \tag{12}$$

where λ is a constant proportionality factor, and $\zeta\left(P_{t+\tau}\right)$ is the pricing kernel. This is a very useful result because, given any of the two functions, we may infer the third one.

Given a utility function and some estimated RND, Bliss and Panigirtzoglou (2004) modify slightly equation (12) to solve for the implied subjective density function once is normalized to integrate to one:

$$f_{t,\tau}\left(P_{t+\tau}\right) = \frac{\dfrac{q_{t,\tau}\left(P_{t+\tau}\right)}{\zeta\left(P_{t+\tau};P_t\right)}}{\displaystyle\int \frac{q\left(y\right)}{\zeta\left(y;P_t\right)}dy} = \frac{\dfrac{U'\left(P_t\right)}{\lambda U'\left(P_{t+\tau}\right)}q_{t,\tau}\left(P_{t+\tau}\right)}{\displaystyle\int \frac{U'\left(P_t\right)}{\lambda U'\left(y\right)}q\left(y\right)dy} = \frac{\dfrac{q_{t,\tau}\left(P_{t+\tau}\right)}{U'\left(P_{t+\tau}\right)}}{\displaystyle\int \frac{q\left(y\right)}{U'\left(y\right)}dy} \tag{13}$$

From the specific functional form of marginal utility, and given the estimated RND, we may estimate the preference parameters that maximize the forecasting ability of the subjective density. This also allows us to analyze the behaviour of the implied risk aversion estimates over time.

In the empirical exercise below we assume the well known power utility function

$$U\left(P_{t+\tau}\right) = \frac{P_{t+\tau}^{1-\gamma} - 1}{1-\gamma} \tag{14}$$

with marginal utility given by $U'\left(P_{t+\tau}\right) = P_{t+\tau}^{-\gamma}$, where γ is the constant relative risk aversion coefficient.

This utility function, although convenient, it is known to have serious problems in explaining both the temporal and cross-sectional behaviour of asset prices. Among other things, it seems that the assumption about the time-invariant behaviour of the risk aversion coefficient is not empirical reasonable. It is not surprising then, that the empirical evidence of implied risk aversion estimates of Bliss and Panigirtzoglou (2004) and Jackwerth (2000) is very controversial. As an example, Bliss and Panigirtzoglou report that the degree of risk aversion declines with the forecast horizon and is lower during periods of high market volatility.

We also assume the exponential utility function

$$-\frac{e^{-\gamma P_{t+\tau}}}{\gamma} \tag{15}$$

with marginal utility $e^{-\gamma P_{t+\tau}}$ and relative risk aversion given by $\gamma P_{t+\tau}$. It should be noted that this function has increasing relative risk aversion and constant absolute risk aversion. Together with the assumption of normal stock returns this is also a tremendously popular specification of preferences. However, as before, risk aversion seems to increase in periods of low market volatility and short horizons.

Surprisingly, the option-implied risk aversion estimates have never been obtained by imposing utility functions which are known to be relevant in explaining the behaviour of asset prices. This may be a serious drawback of this literature and casts doubts on the implied parameters estimated implicitly through option data.

In this study we also assume a habit-based utility function as proposed by Campbell and Cochrane (1999):

$$U\left(P_{t+\tau}, X_{t+\tau}\right) = \frac{\left(P_{t+\tau} - X_{t+\tau}\right)^{1-\gamma} - 1}{1-\gamma} \tag{16}$$

where $X_{t+\tau}$ is the level of habit, and the power coefficient γ is not the relative risk aversion coefficient as in previous specifications. Utility is only defined when financial wealth, as represented by the stock market index level, exceeds habit.

Our problem to incorporate a counter-cyclical time-varying risk aversion is to define a reasonable level of habit which should be by construction lower than the current level of the stock exchange index. The habit-based pricing models employ the so called surplus consumption ratio which is a recession variable defined as

$$S_{t+\tau} = \frac{P_{t+\tau} - X_{t+\tau}}{P_{t+\tau}} \tag{17}$$

Then, under an external specification of habit, marginal utility is given by

$$U'\left(P_{t+\tau}, X_{t+\tau}\right) = \left(P_{t+\tau} - X_{t+\tau}\right)^{-\gamma} = P_{t+\tau}^{-\gamma} S_{t+\tau}^{-\gamma} \tag{18}$$

and the relative risk aversion is

$$RRA_{t+\tau} = \frac{\gamma}{S_{t+\tau}} \tag{19}$$

It is now the case that the local curvature of the utility function depends on how far the current market level is above the habit, as well as the power γ. It is clear that $S_{t+\tau}$ should be always positive independently of the level of the stock index. Hence, when habit is close to the actual level of financial wealth, which may be an indication of a bear market, investors become more risk averse. This is precisely the time-varying behaviour we need to capture in any reasonable utility function specification. In particular, we obtain a counter-cyclical behaviour of risk aversion over time. In fact, a low power coefficient γ can still mean a high (and time-varying) risk aversion.

In order to estimate the subjective density function, as given by expression (13) under the habit-based utility function of equation (16), we follow a paper by Chen and

Ludvigson (2004) on testing pricing models with consumption and habit. We may note first that habit seems reasonable to depend upon the past levels of financial wealth:

$$X_{t+\tau} = h\left(P_{t+\tau}, P_{t+\tau-1}, \ldots, P_{t+\tau-L}\right) \tag{20}$$

and it is also the case that the level of the stock exchange index is trending, so it is necessary to transform the model to use stationary observation on the stock index, such as observations on stock level growth. We assume that the unknown function h is homogeneous of degree one, and this allows us to write habit as

$$X_{t+\tau} = P_{t+\tau} h\left(1, \frac{P_{t+\tau-1}}{P_{t+\tau}}, \ldots, \frac{P_{t+\tau-L}}{P_{t+\tau}}\right) \tag{21}$$

which can be redefined as

$$X_{t+\tau} = P_{t+\tau} g\left(\frac{P_{t+\tau-1}}{P_{t+\tau}}, \ldots, \frac{P_{t+\tau-L}}{P_{t+\tau}}\right) \tag{22}$$

According to our previous reasoning we need to ensure that $X_{t+\tau} < P_{t+\tau}$. A reasonable function would be the following:

$$\psi(x) = \left(1 + e^{-x}\right)^{-1} \tag{23}$$

where $x = \left(\delta \frac{P_{t+\tau-1}}{P_{t+\tau}} + \delta^2 \frac{P_{t+\tau-2}}{P_{t+\tau}} + \ldots + \delta^L \frac{P_{t+\tau-L}}{P_{t+\tau}}\right)$, and $0 \leq \psi(x) \leq 1$. Therefore, independently of δ we have that $X_{t+\tau} < P_{t+\tau}$.

To summarize, the habit specification is given by

$$X_{t+\tau} = P_{t+\tau}\left(1+e^{-\left(\delta\frac{P_{t+\tau-1}}{P_{t+\tau}}+\delta^2\frac{P_{t+\tau-2}}{P_{t+\tau}}+....+\delta^L\frac{P_{t+\tau-L}}{P_{t+\tau}}\right)}\right)^{-1} \tag{24}$$

In the empirical exercise we impose $L = 12$, but we jointly estimate γ and δ by maximizing over the p-value of the Berkowitz *LR* statistic.

5. Data

In this research, we employ the European-style Spanish equity option contract on the IBEX-35 futures which is one of the largest options equity market within the euro area. The Spanish IBEX-35 index is a value-weighted index comprising the 35 most liquid Spanish stocks traded in the continuous auction market system. The official derivative market for risky assets, which is known as MEFF, trades a futures contract on the IBEX-35, the corresponding option on the IBEX-35 futures contracts for calls and puts, and individual futures and option contracts for blue-chip stocks. The option contract on the IBEX-35 futures is a cash settled European option with trading over the three nearest consecutive months and the other three months of the March-June-September-December cycle. The expiration day is the third Friday of the contract month. The multiplier is 1 € and the exercise prices are given by 50 index point intervals. Our database is comprised of settlement IBEX-35 index futures prices, the associated settlement prices of all call and put options traded on each day, and the implied volatility for each option. Moreover, for each option we also have the expiration date and the associated strike. At expiration, the options settle to the exchange delivery settlement future price determined by MEFF by calculating the arithmetic average between 16:15 and 16:45 taking an index value per minute. This series is employed to compute the payoffs of the future in this work.

The options prices employed throughout our research are the MEFF-reported settlement prices. The implied volatility for all at-the-money options reflects the closing market price of each option. For the rest of strikes, MEFF linearly approximates the implied volatility by two segments. Two different slopes are employed for strikes corresponding to options in-the-money and out-of-the-money. The slopes are obtained according to the closing market conditions of the market on each Friday which will be the day from which forecasts are made in our study. The settlement prices are calculated using Black´s (1976) formula, the underlying settlement price and the previous volatilities. Therefore, by construction all options reflect closing market conditions and are synchronous with the underlying asset price. The data cover the period from October 1996 through December 2004; i.e., 99 months[6].

Option settlement prices are available for expirations from one week to one year. It is very important to point out that a target observation date in the study is determined one (four) weeks before every option expiration. The number of strikes ranges between 28 (23) and 211 (211) with an average of 105 (103). Options with expires of less than three months, expire at monthly intervals. Hence, forecasts and realizations for horizons less than or equal to one month may be expected to be independent. The number of cross-sections is 99 for forecasts horizons of either one or four weeks. This is similar to the cross-sections employed by Bliss and Panigirtzoglou (2004) in the case of their data on FTSE 100, and slightly more than half of the available cross-sections for options on the S&P 500.

6. Empirical Results

6.1 The Empirical Performance of Risk-Neutral Densities

In order to determine whether there is evidence that the RNDs adequately forecast the distribution of ex-post realizations of the underlying index, we employ the Berkowitz test statistics discussed in Section 3. Table 1 shows the empirical results using the mixture of two lognormals as the estimation of RNDs for the Spanish stock IBEX-35 index. For the whole sample period, we cannot reject the hypothesis that the RNDs provide accurate predictions of the distributions of future realizations of the IBEX-35 index at the four-week horizon. On the one hand, with a p-value of 0.27 for the *LR* test statistic, we do not support that the RND forecasts poorly the actual realizations. At the same time, by looking at the *LR(i)* statistics we cannot reject the hypothesis that the probability integral transforms are uncorrelated. Hence, we divide the whole sample

[6] Before October 1996 MEFF computed settlement prices using constant implied volatilities.

period into two non-overlapping sub-periods from October 1996 to March 2000, and from April 2000 to December 2004. This allows us to check the robustness of the surprising results found for the complete sample. As reported in Table 1, in the first sub-period the Berkowitz test rejects the hypothesis that the RNDs are good forecasts of future realizations of the IBEX-35 index. Moreover, the *LR(i)* shows that the reason for rejecting is not the violation of the independence assumption underlying the test statistic. This result is consistent with the intuition that RNDs are very unlikely to adequately capture the future behavior of equity prices. It seems reasonable to expect that the stock market prices risks[7]. This result also confirms that the Berkowitz test has sufficient power to reject the null hypothesis. Finally, as in the case of the complete sample, the *LR* statistic is not able to reject the good predictive performance of the RNDs during the second sub-period. This is interesting since the years of the second sub-period coincide with a continuous negative performance of the stock market, and the opposite occurred from October 1996 to March 2000. It seems that during the first sub-period, the high levels reached by the stock market make RNDs unable to place enough probability on the right tail of the distribution relative to actual realizations. This would explain the poor performance of the RNDs during the first half of the sample. We will investigate this potential explanation by analyzing the behaviour of the tails.

As before, the results using a one-week horizon are not robust to alternative sub-periods. Surprisingly, however, the rejection of the null is now associated with the second sub-period. With a very short forecasting horizon noise may be playing a distorting impact on the results.

Table 2 contains the results from the tests designed to analyze the misspecification of the estimated RNDs and risk-adjusted densities on the tails of the distribution. For the right and left tails we compare the frequency with which realizations lie on those areas with the probability mass assigned by the estimated densities. We also report the test statistic given by equation (11). The tests indicate that, using the risk-neutral specification and for the four-week horizon, the probability mass assigned by our estimated RND to the right tail significantly underestimates the frequency of actual

[7] See the recent evidence of Ghysels, Santa-Clara and Valkanov (2004).

realizations. In other words, the strikingly good performance of the stock market during the first sub-period is not adequately forecasted by the RNDs estimated from option prices. For the right tail of the distribution the same results is observed for the full sample period. These results are consistent with the evidence reported for the full body of the implied RNDs. There seems to be a good performance of RNDs from April 2000 to December 2004, and a relatively bad performance of our estimated RND during the bull market of the first sub-period which may be explained by the high frequency of realizations on the right tail of the distribution. Of course, this would be expected since those differences may be arising from the risk aversion of the representative investor.

By contrast, at the one-week horizon the probability mass assigned by our estimated RND to the right tail does not seem to understate the frequency of actual realizations, especially during the second sub-period. Interestingly, it seems that the probability with which realizations lie on both tails is overestimated. This evidence suggests that rejection that RNDs provide good forecasts at the one-week horizon during the second sub-period does not seem to be related to the implied risk-neutrality assumption.

All in all the evidence presented in Tables 1 and 2, which of course is consistent to that reported in Alonso, Blanco and Rubio (2005), suggests that a risk premium adjustment might be needed at the four-week horizon to adequately forecast future outcomes but this does not seem to be the case for the one-week horizon. Section 6.2 investigates this.

6.2 Risk-Adjusted Densities

We first implicitly estimate the parameters and risk aversion of the three alternative utility functions. We employ a process of searching for the optimal level of γ (for the power and exponential cases) and the optimal level of γ and δ (for the habit-based case) to maximize the predicting ability of the resulting risk-adjusted densities by maximizing over γ (and δ when appropriate) the p-value of the Berkowitz statistics[8].

Panel A of Table 3 contains the estimated parameters for the three utility specifications employed. In all cases, the estimated parameters are lower when using a four-week

[8] This process does not provide a measure of whether the resulting parameters are significantly different from zero. A Monte Carlo simulation must be employed to obtain the distribution of generated p-values.

horizon, a disturbing result also obtained by Bliss and Panigirtzoglou (2004). The optimal δ coefficient in the habit-based preferences suggests that past levels of the underlying index tend to have a low impact on the current habit level. It seems to indicate a very low memory in establishing current levels of habit. As expected, the coefficient is higher for the shorter predicting horizon. More interesting are the risk aversion estimates reported in Panel B of Table 3. The mean (median) of both power and exponential utility functions over the whole sample period are very similar. The implied risk aversion estimates for a one-week horizon are around 3.5 while they are close to 1.7 when a four-week horizon is imposed. Of course, as pointed out above, the relative risk aversion under the exponential utility function is $\gamma P_{t+\tau}$. For this reason we provide a range for the risk aversion estimate over the sample period that depends on the associated values of the underlying.

On the other hand, the average risk aversion estimates for the habit-based utility function are higher than either power or exponential utility functions. The mean risk aversion estimate is near 9.0 for a one-week horizon and 3.3 when the longer horizon is considered. In fact, by observing Figure 1, the risk aversion for the habit-based case is systematically above the risk aversion of the exponential utility. This might be expected by taking into account that the temporal dependency created by habit allows for high risk aversion with neither high average nor volatile interest rates. Both the power and exponential utility functions do not share these properties. At the same time, in the Campbell-Cochrane model, the surplus consumption ratio is strongly pro-cyclical and this characteristic magnifies the counter-cyclicality of marginal utility relative to the basic power utility function.

Figure 1 clearly reflects the counter-cyclical property of risk aversion under habit formation for the one-week horizon. The highest risk aversion coefficient of 14.47 is obtained for September 2001. Moreover, the correlation coefficient between the annual market index return calculated monthly and the estimate of risk aversion are -0.45 and -0.39 for one week and four weeks respectively.

Figure 2 compares the estimated PDFs for two different expiration days at a four-week horizon. Panel A shows PDFs estimated with option prices of 24/8/2001, i.e. before the terror attacks of 11 September. On that day, all PDFs have a similar shape. Of course, risk-adjusted PDFs appear (slightly) shifted to the right. Panel B shows PDFs estimated with option prices of 21/9/2001, which reflected the impact on market prices of the events of 11 September. Compared with panel A, the probability mass of the tails, and especially on the left tail, is much higher reflecting the higher uncertainty. Interestingly, risk-adjusted PDFs display lower skewness and kurtosis than those of the RND, suggesting that the latter distribution overstate these moments on stress periods. Moreover, habit-adjusted PDF departs significantly from the other two risk-adjusted PDFs. In particular, it displays the lowest skewness and kurtosis and the highest mean. Finally, this figure illustrates that the risk adjustments used in this paper is more subtle than a simple mean shift.

Panels A, B and C of Table 4 present the Berkowitz tests using the power, exponential and habit-based utility functions respectively. The results are strikingly similar for all three specifications. As in the case of risk-neutral densities, for the whole period we cannot reject the hypothesis that risk-adjusted densities are good predictors of future realizations of the underlying index for both horizons. However, contrary to the RNDs case, the key evidence for the four-week horizon is that risk-adjusted densities also provide adequate forecasting ability for both sub-periods. The results are clearly consistent with the need for a risk premium adjustment. On the other hand, what is certainly surprising is that the specific risk-adjustment imposed does not seem to be relevant when trying to improve the forecasting ability of alternative risk-adjusted densities. This is even true for the habit-case where both the level and the behaviour of risk aversion over time are clearly different relative to more traditional utility functions.

In any case, given the similarities in the forecasting power, and noting that the counter-cyclical behaviour of risk aversion is easily conciliated with the counter-cyclical expected market risk premium over sufficiently long investment horizons, we put more confidence into risk aversion estimates around 8-9 than to the levels obtained under either power or exponential preference specifications.

Finally, for the very short horizon of one-week, it is quite clear that risk-adjustments are not the issue when forecasting throughout either risk-neutral or risk-adjusted densities. As mentioned in the introduction, this should be expected. On the one hand, Figure 1 suggests that the volatility of the marginal rate of substitution is relatively more pronounced for one week than for a four-week horizon. Hence the price of risk seems to be more volatile in very short horizons. However, the quantity of risk embedded in one week is probably negligible compared to the quantity of risk for the four-week horizon. The results are identical independently of whether risk-adjustments are taken into account or not suggesting that rejection of the null during the second sub-period is not related to the risk-neutrality assumption.

Table 5 contains the results from the tests designed to analyze the misspecification of the estimated risk-adjusted densities on the tails of the distribution. As before, the results reported in Panels A, B and C of Table 5 are the same independently of the actual preference specification employed. Hence, time-varying risk aversion needed to reflect the business cycle behaviour embedded in the stock market does not seem to be the issue in explaining the forecasting ability of risk-adjusted densities neither in the whole body of the distribution nor in the tails.

However, risk-adjusted densities perform better than RNDs at the four-week horizon. In particular, for the complete sample period, the difference between the actual frequency observed on the right tail and the probability mass assigned by our estimated densities is not statistically significant. Again, by recognizing the existence of a risk premium we improve the forecasting ability of our estimated densities. However, for the bull market of the first sample period, the true distribution assigns more probability to high returns than the estimated risk-adjusted densities. Not even a risk-adjustment which incorporates the actual business cycle embedded in the stock market is able to adequately incorporate the seemingly optimistic view of investors during the first sub-period.

Finally, the results in the tails of the distribution for a one-week horizon are again quite similar independently of using RNDs or the risk-adjusted counterparts, although the right tail of the distribution is estimated even worst.

7. Conclusions

Option prices provide information about how investors assess the likelihood of alternative outcomes for future market prices of underlying assets. The main objective of this paper is to analyse the value of information contained in prices of options on the IBEX 35 index at the Spanish Stock Exchange Market. The forward looking information is extracted using implied risk-neutral density functions estimated by a mixture of two-lognormals. Moreover, three alternative risk-adjustments are also considered. On the one hand, the classic power and exponential utility functions are analyzed. On the other, a habit-based specification that allows for a counter-cyclical variation of risk aversion is also discussed.

Our results show that between 1996 and 2004, we cannot reject the hypothesis that the RNDs provide accurate predictions of the distributions of future realisations of the IBEX 35 index at both one-week and four-week horizons. Interestingly, when the whole period is divided into two sub-periods, we find that RNDs are not able to consistently predict the excellent behaviour of the stock market from October 1996 to March 2000 at the four-week horizon. In this period, option prices assign a low risk-neutral probability to large rises compared with realisations. On the other hand, RNDs are good predictors of realizations for the period between April 2000 and December 2004 at the same horizon. This suggests that the overall ability of RNDs as a forecasting device is just a consequence of two distinct sub-periods compensating each other.

These results tend to confirm the necessity of risk premium adjustments at four-week horizons with a (probably) countercyclical risk aversion parameter, which seems to be especially relevant for bull markets. When forecasting through risk-adjusted densities the performance of the first sub-period is statistically improved. We cannot reject the hypothesis that risk-adjusted densities provide adequate predictions of the distributions of future realisations of the IBEX 35 index at a four-week horizon. What is more important, and contrary to the RNDs, the results are consistent throughout sub-periods. However, all risk-adjusted densities generate similar forecasting statistics. Then, at least for a horizon of four-weeks, the actual risk adjustment does not seem to be the issue. It is enough to recognize that the stock and option markets price risk. Of course, future research should be concentrated on the importance of risk-adjustments at longer investment horizons.

At the one-week horizon risk adjustments do not improve the forecasting ability of RNDs, suggesting that at very short horizons the assumption of risk-neutrality is reasonable.

REFERENCES

Aït-Sahalia, Y. and A. Lo (2000). "Nonparametric risk management and implied risk aversion", Journal of Econometrics 94, 9-51.

Alonso, F., R. Blanco and G. Rubio (2005). "Testing the forecasting performance of IBEX 35 option-implied risk-neutral densities", Working Paper Series, Bank of Spain, no. 0505.

Anagnou, I., M. Bedendo, S. Hodges and R. Tompkins (2003). "The relation between implied and realized probability density functions", Working Paper, Financial Options Research Centre, University of Warwick.

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

Black, F. (1976). "The pricing of commodity contracts", Journal of Financial Economics 3, 167-179.

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

Bliss, R. and N. Panigirtzoglou (2004). "Option-implied risk aversion estimates", Journal of Finance, 59, 407-446.

Breeden, D. and R. Litzenberger (1978). "Prices of state contingent claims implicit in option prices", Journal of Business, 51, 621-652.

Bondarenko, O. (2003). "Estimation of risk-neutral densities using positive convolution approximation", Journal of Financial Econometrics 116, 85-112.

Cambpell, J. and J. Cochrane (1999). "By force of habit: A consumption-based explanation of aggregate stock market behavior", Journal of Political Economy 107, 205-251.

Chen, X. and S. Ludvigson (2004). “Land of addicts? An empirical investigation of habit-based asset pricing models”, Working Paper, Department of Economics, New York University.

Craig, B., E. Glatzer, J. Keller and M. Scheicher (2003). “The forecasting performance of German stock option densities”, Discussion Paper 17, Studies of the Economic Research Centre, Deutsche Bundesbank.

Diebold, F., T. Gunther and A. Tay (1998). “Evaluating density forecasts, with applications to financial risk management”, International Economic Review 39, 863-883.

Ghysels, E., P. Santa-Clara and R. Valkanov (2004). “There is a risk-return tradeoff after all”, forthcoming in the Journal of Financial Economics.

Jackwerth, J. (2000). “Recovering risk aversion from option prices and realized returns”, Review of Financial Studies 13, 433-451.

Melick, W. and C. Thomas (1997). “Recovering an asset’s implied PDF from option prices: an application to crude oil during the Gulf crisis”, Journal of Financial and Quantitative Analysis 32, 91-115.

Seillier-Moiseiwisch, F. and P. Dawid (1993). “On testing the validity of sequential probability forecasts”, Journal of the American Statistical Association 88, 355-359.

Table 1
Berkowitz Tests for Risk-Neutral Densities Estimated with a Mixture of Two Lognormals
October 1996-December 2004

The reported LR value is the Berkowitz likelihood ratio test for i.i.d. normality of the inverse-normal transformed inverse probability transforms of the realizations as given by $LR = -2[L(0,1,0)-L(\mu,\sigma,\rho)]$ which is distributed as a $\chi^2(3)$. The *LR(i)* statistic is the Berkowitz likelihood ratio test for independence. Rejection of the test for independence suggests that rejection of the density as a good forecast may be due to serial correlation rather than poor forecasting performance.

Forecast Horizon	Oct. 1996-Dec. 2004		Oct. 1996-Mar. 2000		Apr. 2000-Dec. 2004	
	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)
1 week	4.59 (0.20)	0.91 (0.34)	3.53 (0.32)	0.34 (0.56)	12.76 (0.01)	0.30 (0.59)
4 weeks	3.92 (0.27)	1.31 (0.25)	9.45 (0.02)	0.83 (0.36)	0.41 (0.94)	0.02 (0.89)

Table 2
Brier´s Score Tail Tests for Risk-Neutral Densities Estimated with a Mixture of Two Lognormals
October 1996-December 2004

Tests of misspecification for tails of estimated densities. For the right tail and the left tail, the frequency with which actual observations fall in those areas and the probability mass assigned by the mixture of lognormals and splines are reported. The values of the ASN test statistic based on the Brier´s score are also reported. The statistic is asymptotically distributed as a standard normal distribution.

% Tails	Oct. 1996-Dec. 2004			Oct. 1996-Mar. 2000			Apr. 2000-Dec. 2004		
	Freq.	Prob. Forecast[1]	ASN	Freq.	Prob. Forecast	ASN	Freq.	Prob. Forecast	ASN
1 week: Right Tail 5%	0.05	0.08	-0.94	0.10	0.08	0.45	0.02	0.07	-1.66
1 week: Left Tail 5%	0.05	0.09	-1.42	0.05	0.09	-1.01	0.05	0.08	-1.00
4 weeks: Right Tail 10%	0.12	0.07	2.06	0.22	0.08	3.61	0.05	0.06	-0.54
4 weeks: Left Tail 10%	0.05	0.09	-1.29	0.05	0.10	-1.16	0.05	0.08	-0.68

Table 3
Estimated Parameters from Alternative Utility Specifications and Estimates of Risk Aversion

Panel A: Estimated Parameters from Alternative Utility Specifications				
	Power	Exponential	Habit	
Horizon	γ	Γ	γ	δ
1 week	3.54	0.000421	1.55	0.61
4 weeks	1.70	0.000197	1.35	0.27

Panel B: Risk Aversion Estimates							
	Power	Exponential			Habit		
Horizon		Range[1/]	Mean	Median	Range	Mean	Median
1 week	3.54	1.84-5.23	3.50	3.47	7.08-14.47	8.91	8.72
4 weeks	1.70	0.88-2.45	1.64	1.62	3.17-3-63	3.31	3.30

1/ Risk aversion estimates over the whole sample

Table 4
Berkowitz Tests for Power-, Exponential- and Habit-Utility-Adjusted Densities Estimated with a Mixture of Two Lognormals
October 1996-December 2004

The reported LR value is the Berkowitz likelihood ratio test for i.i.d. normality of the inverse-normal transformed inverse probability transforms of the realizations as given by $LR = -2[L(0,1,0)-L(\mu,\sigma,\rho)]$ which is distributed as a $\chi^2(3)$. The *LR(i)* statistic is the Berkowitz likelihood ratio test for independence. Rejection of the test for independence suggests that rejection of the density as a good forecast may be due to serial correlation rather than poor forecasting performance.

Panel A: Power Utility						
Forecast Horizon	Oct. 1996-Dec. 2004		Oct. 1996-Mar. 2000		Apr. 2000-Dec. 2004	
	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)
1 week	2.99 (0.39)	1.00 (0.32)	1.78 (0.62)	0.35 (0.55)	12.49 (0.01)	0.41 (0.59)
4 weeks	2.43 (0.49)	1.74 (0.19)	6.06 (0.11)	1.05 (0.31)	1.83 (0.61)	0.01 (0.94)

Panel B: Exponential Utility						
Forecast Horizon	Oct. 1996-Dec. 2004		Oct. 1996-Mar. 2000		Apr. 2000-Dec. 2004	
	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)
1 week	3.00 (0.39)	1.19 (0.28)	1.99 (0.58)	0.44 (0.51)	12.42 (0.01)	0.52 (0.47)
4 weeks	2.51 (0.47)	1.76 (0.18)	6.21 (0.10)	1.13 (0.29)	1.67 (0.64)	0.00 (0.95)

Panel C: Habit-Based Utility						
Forecast Horizon	Oct. 1996-Dec. 2004		Oct. 1996-Mar. 2000		Apr. 2000-Dec. 2004	
	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)	LR (p-value)	LR(i) (p-value)
1 week	2.93 (0.40)	1.04 (0.31)	1.81 (0.61)	0.35 (0.55)	12.23 (0.01)	0.46 (0.50)
4 weeks	2.49 (0.48)	1.76 (0.18)	6.17 (0.10)	1.05 (0.30)	1.79 (0.62)	0.01 (0.93)

Table 5
Brier´s Score Tail Tests for Power-, Exponential- and Habit-Utility-Adjusted Densities Estimated with a Mixture of Two Lognormals
October 1996-December 2004

Tests of misspecification for tails of estimated densities. For the right tail and the left tail, the frequency with which actual observations fall in those areas and the probability mass assigned by the mixture of lognormals and splines are reported. The values of the ASN test statistic based on the Brier´s score are also reported. The statistic is asymptotically distributed as a standard normal distribution.

Panel A: Power Utility									
% Tails	Oct. 1996-Dec. 2004			Oct. 1996-Mar. 2000			Apr. 2000-Dec. 2004		
	Freq.	Prob. Forecast[1]	ASN	Freq.	Prob. Forecast	ASN	Freq.	Prob. Forecast	ASN
1 week: Right Tail 5%	0.05	0.10	-1.48	0.10	0.10	-0.02	0.02	0.09	-1.97
1 week: Left Tail 5%	0.05	0.06	-0.68	0.05	0.07	-0.56	0.05	0.06	-0.41
4 weeks: Right Tail 10%	0.12	0.09	1.26	0.22	0.10	2.93	0.05	0.08	-0.98
4 weeks: Left Tail 10%	0.05	0.07	-0.61	0.05	0.08	-0.69	0.05	0.06	-0.19

Panel B: Exponential Utility									
% Tails	Oct. 1996-Dec. 2004			Oct. 1996-Mar. 2000			Apr. 2000-Dec. 2004		
	Freq.	Prob. Forecast[1]	ASN	Freq.	Prob. Forecast	ASN	Freq.	Prob. Forecast	ASN
1 week: Right Tail 5%	0.05	0.10	-1.47	0.10	0.10	-0.03	0.02	0.09	-1.95
1 week: Left Tail 5%	0.05	0.07	-0.72	0.05	0.07	-0.53	0.05	0.06	-0.48
4 weeks: Right Tail 10%	0.12	0.09	1.26	0.22	0.10	2.91	0.05	0.08	-0.96
4 weeks: Left Tail 10%	0.05	0.07	-0.66	0.05	0.08	-0.70	0.05	0.06	-0.26

Panel C: Habit-Based Utility									
% Tails	Oct. 1996-Dec. 2004			Oct. 1996-Mar. 2000			Apr. 2000-Dec. 2004		
	Freq.	Prob. Forecast[1]	ASN	Freq.	Prob. Forecast	ASN	Freq.	Prob. Forecast	ASN
1 week: Right Tail 5%	0.05	0.10	-1.47	0.10	0.10	-0.02	0.02	0.10	-1.97
1 week: Left Tail 5%	0.05	0.06	-0.64	0.05	0.07	-0.55	0.05	0.06	-0.36
4 weeks: Right Tail 10%	0.12	0.09	1.28	0.22	0.10	2.95	0.05	0.08	-0.97
4 weeks: Left Tail 10%	0.05	0.07	-0.61	0.05	0.08	-0.69	0.05	0.06	-0.19

1/ The probability forecast is obtained as the average of probabilities from the estimated densities

Figure 1
Risk Aversion Estimated from Power, Exponential and Habit-Based Utilities
October 1996-December 2004

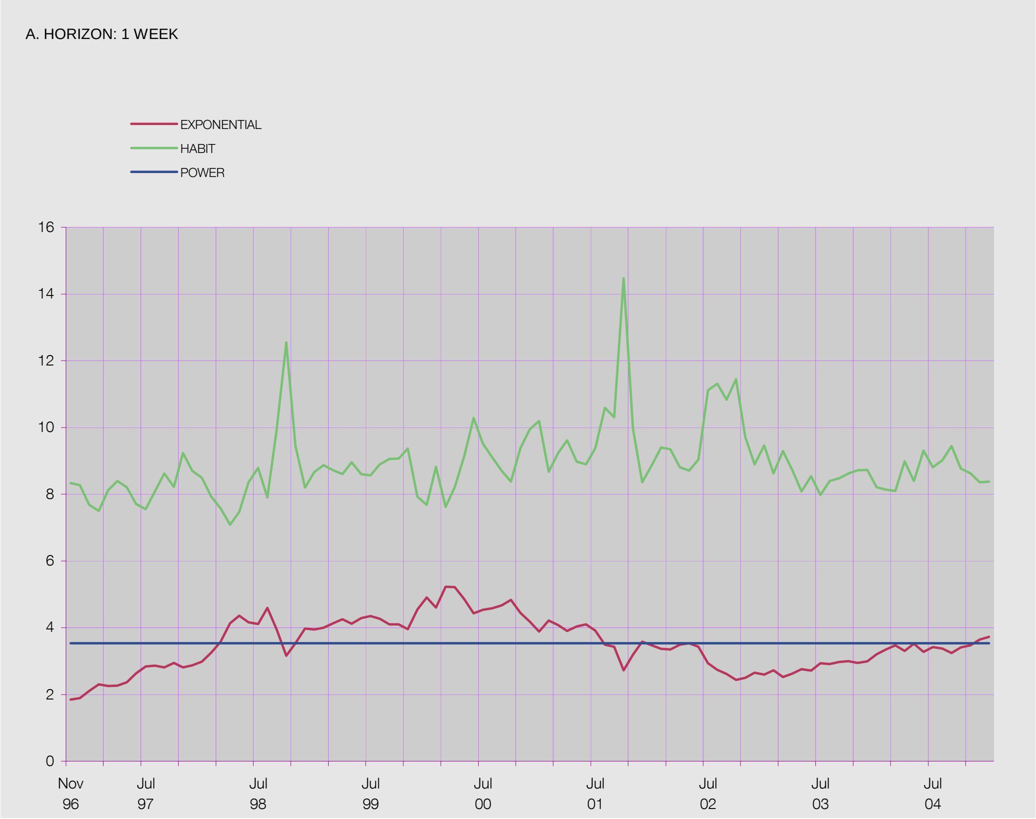

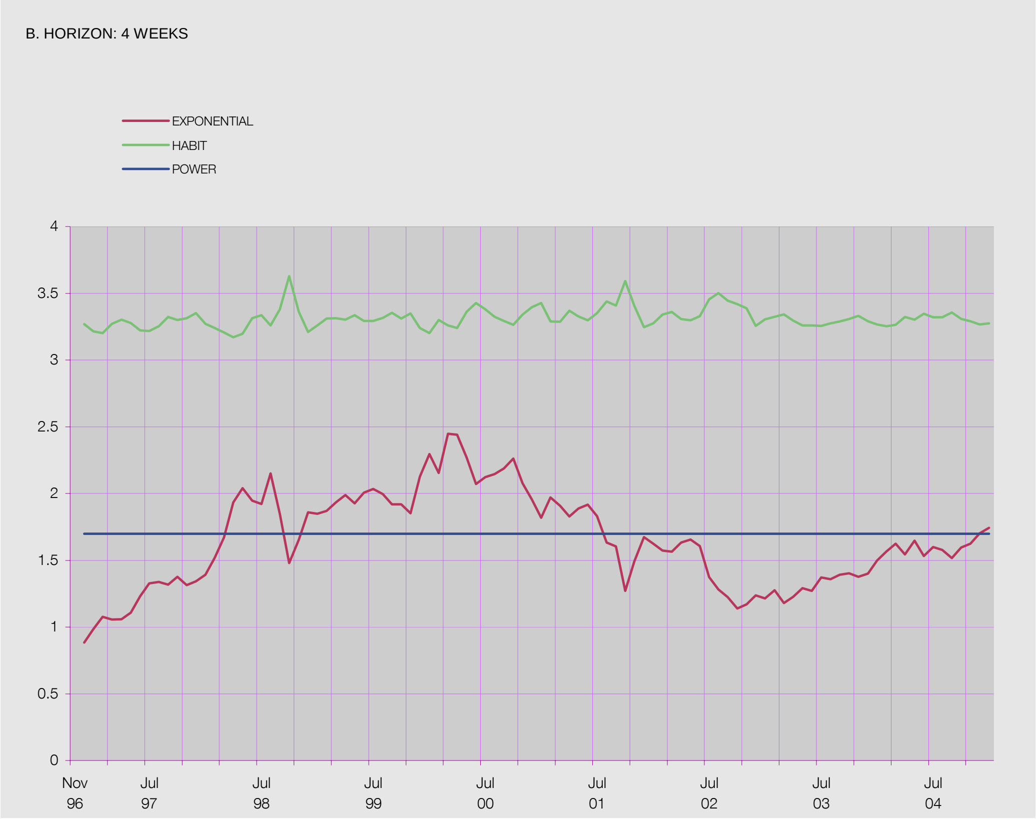

Figure 2
Estimated PDFs for Selected Days. 4-week Horizon