Density forecast comparisons for stock prices, obtained from high-frequency returns and daily option prices

Rui Fan [a,*] Lancaster University

Matteo Sandri [a] Lancaster University

Stephen J. Taylor [a] Lancaster University

This version: May 5, 2015

Abstract

We compare density forecasts for the prices of Dow Jones 30 stocks, obtained from 5-minute high-frequency returns and daily option prices. We use the Heston model which incorporates stochastic volatility to extract risk-neutral densities from option prices. From historical high-frequency returns, we use the HAR-RV model to calculate realized variances and lognormal price densities. We use a nonparametric transformation to transform risk-neutral densities into real-world densities and make comparisons based on log-likelihoods. The lognormal Black-Scholes model gives the highest log-likelihoods for all four horizons ranging from one day to one month, both before and after applying transformations. The HAR-RV model and the Heston model give similar log-likelihoods for all four horizons.

Keywords: HAR models, density forecasts, stock options, high-frequency prices, risk-neutral densities, risk-transformations

JEL classification: C14, C22, C53, G13

[a] Department of Accounting and Finance, Lancaster University Management School, LA1 4YX, UK.
[*] Contact author. Tel.: +44 (0)7831 970830. Email: r.fan1@lancaster.ac.uk (R. Fan).

1. Introduction

Density forecasts are of importance to central bankers, risk managers and other decision takers for activities such as policy-making, risk management and derivatives pricing. They can also be used to assess market beliefs about economic and political events when derived from option prices.

Volatility forecasts produce forward-looking information about the volatility of the asset price in the future, while density forecasts are more sophisticated, as they provide information about the whole distribution of the asset's future price. Since option prices reflect both historical and forward-looking information, volatility forecasters might rationally prefer implied volatilities from option prices to realized variance calculated from historical time series. We anticipate a similar preference could apply to density forecasts. There is a considerable literature comparing volatility forecasts obtained from option prices with volatility forecasts obtained from the history of asset prices. Blair et al. (2001), Jiang and Tian (2005), Giot and Laurent (2007) and Busch et al. (2011) state that option forecasts are more informative and accurate than historical forecasts of index volatility even when the historical information set includes high-frequency returns.[1] Few studies, however, make similar comparisons for density forecasts. Liu et al. (2007), Shackleton et al. (2010) and Yun (2014) provide comparisons for UK and US stock indices, while Hog and Tsiaras (2010) focus on crude oil prices, and Ivanova and Gutierrez (2014) look at interest rates. These studies show option-based density forecasts outperform historical forecasts for a one-month horizon. There are no known previous results for individual

[1] Further comparisons are in Poon and Granger (2003), Martens and Zein (2004) and Taylor et al. (2010).

stocks, so our contribution is to provide the first comparison for density forecasts obtained from option prices and historical intraday returns for individual stocks.

Many methods have been proposed to obtain risk-neutral densities from option prices. Parametric methods include a lognormal mixture (Ritchey, 1990; Jondeau and Rockinger, 2000), a generalized beta distribution (Anagnou-Basioudis et al., 2005; Liu et al., 2007), and a lognormal-polynomial (Madan and Milne, 1994; Jondeau and Rockinger, 2000). Other approaches include discrete probabilities (Jackwerth and Rubinstein, 1996), a nonparametric kernel regression (Ait-Sahalia and Lo, 1998; Bates, 2000), and densities obtained from implied volatility splines (Bliss and Panigirtzoglou, 2002). All these methods, however, only provide densities for horizons which match option expiry dates. We instead fit a stochastic process, to obtain densities for all horizons.

As the implied volatility smile effect indicates that risk-neutral densities are not lognormal and volatility is not constant, some studies use a stochastic process to model volatility. Heston (1993) assumes the volatility follows a mean-reverting square-root process and gives a closed form solution for option prices. We use Heston's model in our study as its parameters can be calibrated from daily option records and it also has a tractable density formula based on inverting characteristic functions. Extensions of the Heston (1993) model are in Bates (1996) who also incorporates jumps, and in Duffie et al. (2000), Eraker (2004), Eraker et al. (2003) and Pan (2002) who include a jump process in both price and volatility components. However, we do not evaluate a jump component because Bakshi et al. (2003) and Shackleton et al. (2010) both find that adding jumps does not improve their empirical results much. Furthermore, our nonparametric transformations can systematically improve mis-specified risk-neutral densities.

We compare density forecasts derived from option prices using the Heston (1993) model and forecasts obtained from historical time series using the Corsi (2009) Heterogeneous Autoregressive model of Realized Variance (HAR-RV). However, the risk-neutral density is a suboptimal forecast of the future distribution of the asset price as there is no risk premium in the risk neutral world, while in reality investors are risk-averse. Hence we need to use economic models and/or econometric methods to transform risk-neutral densities into real-world [2] densities. Pricing kernel transformations include power and/or exponential utility functions (Bakshi et al., 2003; Bliss and Panigirtzoglou, 2004; Liu et al., 2007), and the hyperbolic absolute risk aversion (HARA) function (Kang and Kim, 2006). Liu et al. (2007) use both utility and statistical calibration transformations, and they show that statistical calibration gives a higher log-likelihood than a utility transformation. Shackleton et al. (2010) compare parametric and nonparametric transformations, obtaining good results for the latter. Hence we also transform the risk-neutral densities into real-world densities using a nonparametric transformation.

Early studies including Bakshi et al. (2003), Bliss and Panigirtzoglou (2004) and Anagnou-Basioudis et al. (2005) use the full dataset to make risk-transformations. The real-world densities obtained are then ex post because each forecast is made using some information from later asset prices. However it is best to apply ex ante transformations as in Shackleton et al. (2010). Thus we only use past and present asset

[2] Similar to Liu et al. (2007) and Shackleton et al. (2010), we use “real-world” rather than other alternative adjectives, such as “risk-adjusted”, “statistical”, “empirical”, “physical”, “true”, “subjective” and “objective”, etc., which are all used in the literature to indicate that the price distributions incorporate risk preferences.

and option prices to construct real-world densities. We investigate seventeen stocks from Dow Jones 30 Index for four horizons ranging from one day to one month for the period from 2003 to 2012.

This paper is structured as follows. Section 2 describes the density forecasting methods, namely the Heston (1993) model for densities inferred from option prices, the Corsi (2009) HAR-RV model for density forecasts obtained from historical high-frequency stock prices and the nonparametric transformation of Shackleton et al. (2010). It also includes the econometric methods used to obtain ex-ante parameters and evaluate density forecasts. Section 3 describes the Dow Jones 30 stock and option prices data employed in the study. Section 4 focuses on the empirical analysis. Section 5summarizes the findings and concludes.

2. Methodology

2.1 Option pricing with stochastic volatility

We want to extract the risk-neutral density for the underlying asset from option prices, and a realistic process for an individual stock must incorporate a stochastic volatility component, whose increments are correlated with the price increments. We need to calculate an enormous number of theoretical option prices, so fast calculations are essential. The stochastic volatility process of Heston (1993) meets all our requirements as it has closed-form densities and option prices.

The risk-neutral price dynamics for the stock price S, which incorporate the stochastic variance V, is defined as below

$$\frac{dS}{S} = (r - q)dt + \sqrt{V}dW_1 \tag{1}$$

where r is the risk-free interest rate, q is the dividend yield, and W_1 is a Wiener process. For the variance, we have the familiar square-root process of Cox, Ingersoll and Ross (1985) written as

$$dV = \kappa(\theta - V)dt + \sigma\sqrt{V}dW_2 \tag{2}$$

We let ρ denote the correlation between the two Wiener processes W_1 and W_2, while θ is the level towards which the stochastic variance V reverts, and κ denotes the rate of reversion of V towards θ. The volatility of volatility parameter σ controls the kurtosis of the returns. More complicated affine jump-diffusion processes which have closed-form solutions are described by Duffie et al. (2000). We do not consider these, noting that Shackleton et al. (2010) obtained no benefits from including price jump in their study.

Similar to the Black-Scholes formula, at time 0 the Heston call price formula is derived by assuming

$$C(S_0, V_0, 0) = S_0P_1 - KP(0,T)P_2. \tag{3}$$

The first term S_0 is the current value of the spot price, while the second term $KP(0,T)$ is the present value of the strike price K. Each P_j in equation (3) is a conditional probability that the call option expires in-the-money. The term P_2 is derived from the characteristic function of S_T under the risk-neutral measure Q, while P_1 is derived from the characteristic function of S_T under a related measure Q^* for different drift rates.

Probabilities are obtained from the conditional characteristic function of $log(S_T)$, which is denoted by $g(\Phi)$ and defined for all real numbers Φ, with $i = \sqrt{-1}$, as

$$g(\Phi) = E\left(e^{i\Phi log(S_T)} \middle| S_0, V_0\right). \quad (4)$$

This is a complex-valued function. Heston (1993) solves the PDEs to get the characteristic function solution

$$g(\Phi) = e^{\,C+DV_0+i\Phi\, log(S_0)} \quad (5)$$

When the asset pays continuous dividends, so q>0. S_0 is replaced by S_0e^{-qT} in (3) and (5). For options on futures, q=r. Each desired probability can be obtained by inverting the characteristic function, which is given as

$$P(S_T \geq K \mid S_0, V_0) = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left[\frac{e^{-i\Phi\, log(K)}\, g(\Phi)}{i\Phi}\right] d\Phi \quad (6)$$

where $Re[.]$ is the real part of a complex number (Kendall et al. 1987). This integral can be evaluated rapidly and accurately by numerical methods. It also provides the conditional cumulative distribution function of S_T, therefore

$$F(y) = P(S_T \leq y \mid S_0, V_0) = 1 - P(log(S_T) \geq log(y) \mid S_0, V_0).$$

From routine calculations, the conditional risk-neutral density for positive values of y is hence

$$f(y) = \frac{dF}{dy} = \frac{1}{\pi y}\int_0^\infty Re\left[e^{-i\Phi\, log(y)} g(\Phi)\right] d\Phi. \quad (7)$$

2.2 High-frequency HAR methods

The HAR-RV model of Corsi (2009) is a simple AR-type model for the realized volatility which combines different volatility components calculated over different time horizons. The HAR-RV model states that the multiperiod realized variance is the average of the corresponding one-period measures denoted as

$$RV_{t,t+h} = h^{-1}[RV_{t+1} + RV_{t+2} + \cdots + RV_{t+h}] \tag{8}$$

where h=1, 2, …, by definition $RV_{t,t+h} \equiv RV_{t+h}$ and we use h=5 and h=22 to represent the weekly and monthly realized volatility. Here the time period for predictions is from t to t+h, both counting trading days. In contrast, our options notation is a time period from 0 to T, both measured in years.

The HAR-RV model of Corsi (2009) is stated as a regression of the next RV on today's RV and the average RVs over the latest week and month:

$$RV_{t+1} = \beta_0 + \beta_D RV_t + \beta_W RV_{t-5,t} + \beta_M RV_{t-22,t} + \varepsilon_{t+1}.$$

To make predictions for the next h-day period, the regression specification is simply:

$$RV_{t,t+h} = \beta_{0,h} + \beta_{D,h} RV_t + \beta_{W,h} RV_{t-5,t} + \beta_{M,h} RV_{t-22,t} + \varepsilon_{t,t+h}. \tag{9}$$

Some volatility forecast models also employ standard deviations as opposed to variances. Andersen et al. (2007) present the standard deviation form of HAR-RV model as

$$(RV_{t,t+h})^{1/2} = \beta_{0,h} + \beta_{D,h} RV_t^{1/2} + \beta_{W,h} (RV_{t-5,t})^{1/2} + \beta_{M,h} (RV_{t-22,t})^{1/2} + \varepsilon_{t,t+h} \tag{10}$$

Given the logarithmic daily realized volatilities are approximately unconditionally normally distributed, Andersen et al. (2007) also predict the realized variance in logarithmic form as

$$log\left(RV_{t,t+h}\right) = \beta_{0,h} + \beta_{D,h}\, log(RV_t) + \beta_{W,h}\, log\left(RV_{t-5,t}\right) + \beta_{M,h}\, log\left(RV_{t-22,t}\right) + \varepsilon_{t,t+h} \quad (11)$$

We also use the logarithmic form of realized variance in our study. However, Pong et al. (2004) state that we cannot simply take the exponential of a forecast of logarithmic volatility to get a forecast of the variance, as the forecasts obtained will be biased. We thus follow Granger and Newbold (1976) to get the volatility forecast. In their notation,

$$x_{n+h} = f_{n+h} + e_{n+h}^{(x)} \quad (12)$$

where $e_{n+h}^{(x)}$ is the h-step forecast error of x_{n+h} and f_{n+h} is the optimal forecast of x_{n+h} made at time n. Using $I_n = \{x_{n-j}, j \geq 0\}$, we define $S^2(h)$ to be the variance of the h-step forecast error of x_{n+h}:

$$S^2(h) = var\left(e_{n+h}^{(x)}\right). \quad (13)$$

The optimal forecast of $exp(x_{n+h})$ using I_n is then given by

$$g_{n+h}^{(x)} = \exp\left(f_{n+h} + \frac{1}{2}S^2(h)\right) \quad (14)$$

assuming $\{x_n\}$ is a Gaussian process. This is a standard assumption for $log(RV_t)$.

2.3 Lognormal densities, from the Black-Scholes model and HAR-RV forecasts

In the Black-Scholes model, we assume the prices follow geometric Brownian motion

$$dS/S = \mu dt + \sigma dW \quad (15)$$

where μ is the expected return per annum, and is equal to the risk free rate plus the asset's risk premium and minus the dividend yield.

Since the distribution of stock price S_T is then lognormal, the distribution of $log(S_T)$ is normal:

$$log(S_T) \sim N(log(S_0) + \mu T - \frac{1}{2}\sigma^2 T, \sigma^2 T)$$

Under the risk-neutral or the *Q*-distribution, the risk-neutrality assumption requires a drift rate *r-q* instead of μ, and hence we have

$$\log(S_T) \sim N(\log(S_0) + (r-q)T - \frac{1}{2}\sigma^2 T, \sigma^2 T)$$

and $$E^Q[S_T] = S_0 e^{(r-q)T} = F \tag{16}$$

where F is the no-arbitrage, futures price for time T.

The risk-neutral density of S_T then depends on three parameters (F, σ, T) and is given by the lognormal density

$$\psi(x|F,\sigma,T) = \frac{1}{x\sigma\sqrt{2\pi T}} e^{-\frac{1}{2}\left(\frac{log(x)-\left[log(F)-\frac{1}{2}\sigma^2 T\right]}{\sigma\sqrt{T}}\right)^2}. \tag{17}$$

Similarly, a risk-neutral, lognormal density from the HAR-RV model can be given by replacing $\sigma\sqrt{T}$ by a term $\widehat{RV}_T$ to give:

$$\psi\left(x\middle|F,\widehat{RV}_{t,t+h}\right) = \frac{1}{x\sqrt{2\pi\widehat{RV}_{t,t+h}}} e^{-\frac{1}{2}\left(\frac{log(x)-\left[log(F)-\frac{1}{2}\widehat{RV}_{t,t+h}\right]}{\sqrt{\widehat{RV}_{t,t+h}}}\right)^2} \tag{18}$$

The quantity $\widehat{RV}_{t,t+h}$ is calculated from (5.18) and (5.21) with the horizon h (measured in trading days) and maturity T (measured in years).

2.4 Nonparametric transformations

The risk-neutral, Q-densities are not satisfactory specifications of the real-world densities. One reason is that Q-variance obtained from option prices is usually higher than the real-world variance, because there is a negative volatility risk premium (Carr and Wu, 2009). Consequently there are fewer observations than predicted in the tails of the Q-densities. A second reason is that the equity risk premium is, by definition, absent from all the risk-neutral densities. Hence it is necessary to use some technique to transform risk-neutral densities into real-world densities.

We consider the nonparametric calibration method in this study. Nonparametric calibration functions are re-estimated for each period t. At time t (which counts trading days), the nonparametric transformation for a selected horizon h is determined by a set of t-h+1 cumulative, risk-neutral probabilities

$$u_{s+h} = F_{Q,s,T}(p_{s+h}|\Theta_s), \qquad 0 \leq s \leq t-h, \tag{19}$$

with T (years) matching h (trading days), s a time before t-h+1, $F_{Q,\,s,\,T}$ the cumulative distribution function of the price p_{s+h}, and with Θ_s a vector of density parameters. We assume the observations u_{s+1} are i.i.d. and their c.d.f. is given by the calibration function $C_T(u)$.

The values of the variables u for the Heston model are given by (6). The variables u for the HAR-RV model can be derived in the following way. For the risk-neutral dynamics,

$$log(p_{s+h}) \sim N\left(log\left(F_{s,s+h}\right) - \frac{1}{2}\widehat{RV}_{s,s+h}, \widehat{RV}_{s,s+h}\right)$$

with $F_{s,\,s+h}$ the futures price at time s for a transaction at time s+h and with $\widehat{RV}_{s,s+h}$

the forecast of RV for the period from time s to $s+h$ inclusive. From the outcome $\hat{p}_{s+h}$ we calculate

$$u_{s+h} = F_{Q,s,T}(\hat{p}_{s+h}|\Theta_s)$$

$$= \Phi\left(\frac{log(\tilde{p}_{s+h}) - (log(F_{s,s+h}) - \frac{1}{2}\widehat{RV}_{s,s+h})}{\sqrt{\widehat{RV}_{s,s+h}}}\right). \quad (20)$$

The values of the variables u for the Black-Scholes model are given in a similar way[3]

$$u_{s+h} = \Phi\left(\frac{log(\tilde{p}_{s+h}) - (log(\tilde{F}_{s,s+h}) - \frac{1}{2}\sigma^2 T)}{\sigma\sqrt{T}}\right) \quad (21)$$

We use $\varphi()$ and $\Phi()$ to represent the density and the c.d.f. of the standard normal distribution. We then transform the observations u_i, whose domain is from 0 to 1, to new variables $y_i=\Phi^{-1}(u_i)$, and then fit a nonparametric kernel c.d.f. to the set $\{y_1, y_2, \ldots, y_{t-h+1}\}$. We use a normal kernel with bandwidth B to obtain the kernel density and c.d.f.:

$$\hat{h}_T(y) = \frac{1}{(t-h+1)B}\sum_{i=1}^{t-h+1}\varphi\left(\frac{y-y_i}{B}\right),$$

$$\hat{H}_T(y) = \frac{1}{t-h+1}\sum_{i=1}^{t-h+1}\Phi\left(\frac{y-y_i}{B}\right). \quad (22)$$

The bandwidth B decreases as t increases. We apply the standard formula of Silverman (1986), where $B=0.9\sigma_y/t^{0.2}$ and σ_y is the standard deviation of the terms y_i.

The empirical calibration function is then

[3] When calculating densities and u variables, we use forward prices evaluated on the day of forecast, and calculate for four horizons, i.e. one day, one week, two weeks and one month.

$$\hat{C}_T(u) = \hat{H}_T\left(\Phi^{-1}(u)\right) \tag{23}$$

which is calculated at time *t*. At the same time, we let $f_{Q,T}(x)$ and $F_{Q,T}(x)$ denote the risk-neutral density and the cumulative distribution function of the random variable p_T. We define $u_T = F_{Q,T}(p_T)$. We follow Bunn (1984) and denote the calibration function $C_T(u)$, which is the real-world c.d.f. of the random variable u_T. Now we consider the real world c.d.f. of p_T, with *Pr* referring to the real world probabilities. The c.d.f. is

$$Pr(p_T \leq x) = Pr\left(F_{Q,T}(p_T) \leq F_{Q,T}(x)\right) = Pr\left(u_T \leq F_{Q,T}(x)\right) = C_T\left(F_{Q,T}(x)\right) \tag{24}$$

Consequently replacing $C_T(.)$ by $\hat{C}_T(.)$, the predictive real-world c.d.f. of p_T is

$$F_{P,T}(x) = \hat{C}_T\left(F_{Q,T}(x)\right) \tag{25}$$

The real-world density is

$$\begin{aligned} f_{P,T}(x) &= \frac{d}{dx}\hat{H}_T\left(\Phi^{-1}\left(F_{Q,T}(x)\right)\right) = \frac{d}{dx}\hat{H}_T(y) = \frac{dy}{dx}\frac{d\hat{H}_T(y)}{dy} \\ &= \frac{du}{dx}\frac{dy}{du}\hat{h}_T(y) = \frac{f_{Q,T}(x)\hat{h}_T(y)}{\varphi(y)}. \end{aligned} \tag{26}$$

Also the nonparametric calibration density is

$$\hat{c}_T(u) = \frac{d}{du}\hat{C}_T(u) = \frac{d}{du}\hat{H}_T(y) = \frac{d\hat{H}_T(y)}{dy}\frac{dy}{du} = \frac{\hat{h}_T(y)}{\varphi(y)}. \tag{27}$$

2.5 Parameter estimation

The densities are all evaluated out-of-sample and thus the parameter values are obtained ex ante, i.e. the values at time t are estimated based on the information available at time t. For the HAR variances we estimate all parameters from regressions over five-year windows. For Black-Scholes lognormal densities, we use the nearest-the-money, nearest-to-expiry option implied volatility.

For the Heston model, we estimate the risk-neutral parameters of the asset price dynamics every day. On each day, we estimate the initial variance V_t, the rate of reversion κ_t, the unconditional expectation of stochastic variance θ_t, the volatility of volatility σ_t, and the correlation ρ_t between the two Wiener processes. Assume there are N_t European, call[4] option contracts traded on day t, denoted by $i=1, \ldots, N_t$, and the market prices are $c_{t,i}$, for strike prices $K_{t,i}$, and expiry times $T_{t,i}$. We also assume $p_{t,i}$ is the futures price for the asset, calculated for a synthetic futures contract which expires in $T_{t,i}$ years. Then we calibrate the five risk-neutral Heston parameters by minimizing the total squared errors in

$$\sum_{i=1}^{N_t}\left(c_{t,i} - c\left(p_{t,i}, K_{t,i}, T_{t,i}, V_t, \kappa_t, \theta_t, \sigma_t, \rho_t\right)\right)^2 \tag{28}$$

with $c(.)$ the solution for the European call option price from the Heston model given in (3).[5]

2.6 Econometric methods

2.6.1 Maximum log-likelihood

[4] We use put-call parity to obtain the equivalent European call prices from the put prices, and then apply them to (5.5), this is also discussed in section 4.1.2.

[5] Christoffersen and Jacobs (2004) conclude that it is a “good general-purpose loss function in option valuation applications”. Christoffersen et al. (2006) also employed it in the study of S&P 500 dynamics.

There are several ways to evaluate density forecasts, and we will use the standard log-likelihood criterion previously employed by Bao et. al (2007), Liu et. al (2007) and Shackleton et al. (2010). For a given horizon h, assuming method m gives densities $f_{m,t}(x)$ at times $i, .., j$ for the asset price at times $i+h, ..., j+h$. Our goal is to find the method which maximizes the out-of-sample log-likelihood of observed asset prices, and this log-likelihood for method m is given by

$$L_m = \sum_{t=i}^{j} log\left(f_{m,t}(p_{t+h})\right) \tag{29}$$

To compare two methods we apply a version of the log-likelihood ratio test in Amisano and Giacomini (2007). The null hypothesis states that two different density forecasting methods m and n have equal expected log-likelihood. The test is based on the log-likelihood differences

$$d_t = log\left(f_{m,t}(p_{t+h})\right) - log\left(f_{n,t}(p_{t+h})\right), \qquad i \leq t \leq j. \tag{30}$$

Amisano and Giacomini (2007) follow Diebold and Mariano (1995) and add the assumption that the differences are uncorrelated. Hence the AG test statistic is

$$t_{i,j} = \frac{\bar{d}}{s_d/\sqrt{(j-i+1)}} = \frac{L_m - L_n}{s_d\sqrt{(j-i+1)}} \tag{31}$$

This statistic follows a standard normal distribution, where $\bar{d}$ is the mean and s_d is the standard deviation of the terms d_t.

When h>1 the forecasts overlap and it is plausible to expect some autocorrelation in the differences. A Newey-West adjustment should then be made when estimating the variance of $\bar{d}$. Assuming the terms d_t are stationary,

$$var(\bar{d}) = var\left(\frac{d_1 + d_2 + \cdots + d_n}{n}\right)$$

$$= \frac{1}{n^2}[nvar(d_1) + 2(n-1)cov(d_1, d_2) + \cdots + 2cov(d_1, d_n)]$$

$$= \frac{var(d_1)}{n}\left[1 + 2\left(\frac{n-1}{n}\right)\rho_1 + 2\left(\frac{n-2}{n}\right)\rho_2 + \cdots + 2\left(\frac{1}{n}\right)\rho_{n-1}\right]$$

where the autocorrelations are $\rho_\tau = cor(d_t, d_{t+\tau})$. The typical estimate of the variance of $\bar{d}$ is

$$\frac{s_d^2}{n}[1 + 2\omega_1\hat{\rho}_1 + \cdots + 2\omega_k\hat{\rho}_k]$$

and a standard set of weights for k estimated autocorrelations is $\omega_\tau = \frac{k+1-\tau}{k+1}$, $1 \leq \tau \leq k$.

2.6.2 Diagnostic tests

Appropriate diagnostic tests use properties of time series derived from density forecasts. Rosenblatt (1952) introduces the probability integral transform (PIT), and states that the PIT values are i.i.d. uniform for known densities. Diebold et al. (1998) initiated the idea of using PIT values to evaluate density forecasts. Following this and Shackleton et al. (2010), we also employ a series of observed cumulative probabilities to check the accuracy of the forecasts. For a given method *m* the PIT probabilities are given by

$$u_{t+1} = \int_0^{p_{t+1}} f_{m,t}(x)dx, \tag{32}$$

for prices p_{t+1} matched with densities $f_{m,t}(x)$.

We then evaluate if the values of u are compatible with i.i.d. observations from the uniform distribution. We can employ the Kolmogorov and Smirnov test. The KS test checks the maximum difference between the empirical and theoretical cumulative functions. For forecasts made at times $i \leq t \leq j$, the sample c.d.f. of $\{u_{i+1}, \ldots, u_{j+1}\}$, evaluated at u, is the proportion of values less than or equal to u, i.e.

$$\tilde{C}(u) = \frac{1}{j-i+1} \sum_{t=i+1}^{j+1} S(u - u_t) \tag{33}$$

with $S(x)$=1 if $x \geq 0$, and $S(x)$=0 if $x < 0$. The test statistic is given by

$$KS = \sup_{0 \leq u \leq 1} \left| \tilde{C}(u) - u \right|. \tag{34}$$

The KS test is widely applied because it is easy to implement. However, one needs to be cautious when interpreting the test results, as the KS test checks for uniformity under the i.i.d. assumption rather than tests i.i.d. and uniformity jointly.

Some researchers doubt the power of the KS test when evaluating density forecasts. Berkowitz (2001) invented the BK test, which states that if the PIT is i.i.d. uniform, then the normal inverse cumulative function of the PIT is i.i.d. normal. The advantage of the BK test is that it can test independence and uniformity jointly. The BK test has been applied in Clements and Smith (2000), Clements (2004), Guidolin and Timmermann (2005) and Shackleton et al. (2010),.

The BK method transforms the observations u_i to new variables $y_i = \Phi^{-1}(u_i)$, with $\Phi()$ the c.d.f. of the standard normal distribution. The null hypothesis of the test is that the values of y are i.i.d. and follow a standard normal distribution, against the alternative hypothesis that y is a stationary, Gaussian, AR(1) process with no restrictions on the mean, variance and autoregressive parameters. Let

$$y_t - \mu = \rho(y_{t-1} - \mu) + \varepsilon_t. \tag{35}$$

Then the null hypothesis is that $\mu = 0$, $\rho = 0$, and $var(\varepsilon_t) = 1$. The log-likelihood for T observations from (35) is

$$-\frac{T}{2}log(2\pi) - \frac{1}{2}log[\sigma^2/(1-\rho^2)] - \frac{(y_1 - \mu/(1-\rho))^2}{2\sigma^2/(1-\rho^2)} - \frac{T-1}{2}log(\sigma^2)$$

$$-\sum_{t=2}^{T}\left(\frac{(y_t - \mu - \rho y_{t-1})^2}{2\sigma^2}\right) \tag{36}$$

Here σ^2 is the variance of ε_t. The log-likelihood is written as a function of the unknown parameters of the model, $L(\mu, \sigma^2, \rho)$. The log-likelihood ratio test (LR3) is

$$LR_3 = -2(L_0 - L_1) = -2\big(L(0,1,0) - L(\hat{\mu}, \hat{\sigma}^2, \hat{\rho})\big). \tag{37}$$

Here hats denote maximum-likelihood values, L_0 and L_1 are the maximum log-likelihoods for the null and alternative hypotheses, and the test statistic has an asymptotic x_3^2 distribution. One disadvantage of the BK test is that models cannot be easily compared if they are all accepted or rejected. The AG test, which we discussed before, compares the log-likelihoods between models and solves this problem.

3. Data

3.1 Option data

We investigate the Dow Jones Industrial Average (DJIA) 30 Index stocks[6] for 10 years from 1st January 2003 to 31st December 2012. The option data are obtained from Ivy DB OptionMetrics, which includes price information for all U.S. listed equity options, based on daily closing quotes at the CBOE.

Although the components of the Dow Jones 30 index have changed many times in its history, we simply use the constituent stocks at the end of our sample period (24th September 2012 is the date of the last change). A detailed list of the component stocks of the DJIA is shown in Table 1.

The OptionMetrcis database also includes information about end-of-day security prices and zero-coupon interest rate curves. The security price file provides the closing price for each security on each day from CRSP.

3.2 Option prices

In terms of filtering option price records, we follow the criteria of Carr and Wu (2003, 2009 and 2010) and Huang and Wu (2004). We delete an option record when the bid price is zero or negative. We also delete an option record when the bid price is greater than the ask price. As do Carr and Wu (2009), we eliminate all the options which have maturity equal to or more than one year. Following Carr and Wu (2003), Huang and Wu (2004), Shackleton et al. (2010) and Taylor et al. (2010), we delete all data for options with maturity equal to or less than seven calendar or five business days.

[6] To date, complete results are available for seventeen stocks. Further results will be included later.

All the equity options are American. OptionMetrics provides implied volatilities, calculated from binomial trees which incorporate dividends and permit early exercise. We use equivalent European option prices defined by assuming the European and American implied volatilities are equal. This method assumes the early exercise premium can be obtained from constant volatility pricing models. The assumption is particularly reasonable for out-of-the-money options which have small early exercise premia.

European call and put prices for the same strike and maturity theoretically contain the same information. Either the call option or the put option will be out-of-the-money (OTM), or under rare circumstances both are at-the-money (ATM). Options are ATM when the strike price equals the stock price (S=K), calls are OTM when $S<K$ and puts are OTM when $S>K$; they are nearest-the-money if $|S - K|$ is nearer zero than for all other contemporaneous strikes. We choose to use the information given by the prices of OTM and ATM options only, because in-the-money (ITM) options are less liquid and have higher early exercise premia. We use put-call parity to obtain equivalent European call prices from the European OTM put prices.

3.3 Interest rates

We follow Taylor et al. (2010) to get the interest rate corresponding to each option's expiry by linear interpolation of the two closest zero-coupon rates supplied by Ivy DB.

3.4 IBM example

We use IBM to illustrate our data and results. A total of 109,111 option prices are investigated in our sample period for IBM stock. The average number of option prices used per day is 44, consisting of 19 OTM calls and 25 OTM puts. Table 2 presents the quantity, moneyness and maturity of the option contracts used in this paper.

3.5 Futures prices

We calculate synthetic futures prices, which have the same expiry dates as the options, as the future value of the current spot price minus the present value of all the dividends expected during the life of the futures contract until the option expiry time T, i.e.

$$F = e^{rT}\big(S - PV(dividends)\big) \tag{38}$$

3.6 High-frequency stock prices

We use the transaction prices of DJIA 30 Index stocks for ten years during the period between 1st January 1998 and 31st December 2012. The data are obtained from pricedata.com. The prices provided are the last prices in one-minute intervals. After an inspection of the high-frequency data, we find a number of problematic days which do not have complete trading records. We set the price equal to that for the previous minute when there is a missing record, and we delete a day when there are more than 40 consecutive missing prices. The days deleted are usually close to holidays such as New Year's Day, Easter, Independence Day, Thanksgiving Day and Christmas.

Between 2003 and 2012, 17 days are deleted because of missing high-frequency prices and these days usually only have prices for half a day. There are also 8 days with unsatisfactory option price data. All 25 days are deleted from the high-frequency and option files leaving a sample of 2488 days for each firm for the ten-year period ending on 31st December 2012.

The stocks are traded for six-and-a-half-hours, from 9:30 EST to 16:00 EST. We calculate realized variances from 5-minute returns because Bandi and Russell (2006) state that the 5-minute frequency provides a satisfactory trade-off between maximizing the accuracy of volatility estimates and minimizing the bias from microstructure effects. As usual, returns are changes in log prices. We have 77 5-minute intraday returns for each day after deleting the data in the first five minutes to avoid any opening effects. The realized variance for day *t* is the sum of the squares of the 5-minute returns $r_{t,i}$:

$$RV_t = \sum_{i=1}^{77} r_{t,i}^2. \tag{39}$$

However, this calculation of realized variance is downward biased as a measurement of close-to-close volatility over a 24-hour period. This is because we only include the information during the trading period when we calculate the realized variance for a day, so the variation overnight (from close-to-open) is excluded. We thus need to scale the realized variance up. We multiply forecasts from the HAR-RV model by a scaling factor. The denominator of the scaling factor is the sum of the squares of the 5-minute returns representing the open market period, while the numerator of the scaling factor is the sum of the squares of the daily returns representing open and closed market periods. We use a rolling window for the scaling factor, hence if we forecast the realized variance on day *t*, then we use the information about returns up to and including day *t* to calculate

$$\widehat{RV}_{t,t+h}\left(\frac{\sum_{i=1}^{t} r_t^2}{\sum_{i=1}^{t}\sum_{j=1}^{77} r_{t,j}^2}\right).$$

This quantity replaces $\widehat{RV}_{t,t+h}$ in (18) when the high-frequency, lognormal densities are evaluated.

4. Empirical results

4.1 Heston risk-neutral parameters

Table 3 shows the summary statistics for risk-neutral parameters calibrated for IBM and across all stocks for each day in our sample period. The risk-neutral parameters minimize the mean squared error (MSE) of option prices on each day.

For IBM, our median estimate of the stochastic variance θ is 0.3457, equivalent to an annualized volatility level of 58.80%. The mean estimate of the rate of reversion κ is 1.6861, for which the half-life parameter of the variance process is then about 5 months. The median estimate of the volatility of volatility parameter σ which controls the kurtosis of returns is 0.8617. Also the median estimate of the correlation ρ is -0.6652, consistent with estimates in the literature.

4.2 Examples of density forecasts

The one-day ahead Heston, lognormal and HAR densities for IBM calculated on January 2nd 2003 are shown in Figure 1. The Heston density is negatively skewed while the lognormal density is slightly positively skewed. The HAR density is seen to have less variance than the Heston and the lognormal densities. The one-month ahead Heston, lognormal and HAR densities for IBM calculated on January 2nd 2003 and shown in Figure 2 display similar properties.

4.3 Examples of cumulative probabilities and nonparametric transformations

The one-day ahead risk-neutral densities give the cumulative distribution functions $F_{Q,t}(x)$ for the next stock price p_{t+1}, and the observed risk-neutral probabilities $u_{t+1}=F_{Q,t}(p_{t+1})$ are not consistent with uniform probabilities, as expected. The sample cumulative probabilities $\tilde{C}(u)$ are calculated using (33), and the deviations between the sample c.d.f. and a uniform c.d.f., namely $\tilde{C}(u) - u$, are plotted in Figure 3 for IBM, for one-day-ahead forecasts obtained from the Heston model. We can observe from the figure that there are few observations u close to either zero or one; only 7.3% of the variables u are below 0.1 and only 5.1% of them are above 0.9. The KS test statistic is the maximum value of $|\tilde{C}(u) - u|$, which is equal to 7.1%, hence the null hypothesis of a uniform distribution is rejected at the 0.01% significance level. The shape of the curve may be explained by the fact that the historical volatility is lower than the risk-neutral volatility, hence the risk-neutral probabilities of large price changes exceed the real-world probabilities. The corresponding plot for IBM for one-day-ahead forecasts obtained from Black-Scholes lognormal densities is shown in Figure 4 and is similar.

The nonparametric transformation of the probabilities u_{t+1} used in the calculation of the real-world density is calculated from (27). The calibration densities $\hat{c}(u)$, for one-day ahead Heston and Black-Scholes lognormal forecasts are shown in Figure 5 and 6; these densities use the values of u for all 10 years from 2003 to 2012. The purpose of the calibration is to create real-world densities which have uniformly distributed observed probabilities u_{t+1}. The differences $\tilde{C}(u) - u$ after applying the nonparametric calibration method for one-day ahead forecasts from Heston and Black-Scholes lognormal densities are shown in Figure 3 and 4. The differences are much nearer zero compared to the risk-neutral densities. Comparable figures and results are obtained for longer horizon density forecasts.

4.4 Log-likelihood comparison

Table 4 gives the log-likelihoods for IBM and another sixteen stocks from 2003 to 2012 under six measures. The density forecasts are made overlapping for four horizons, namely one day, one week (5 trading days), two weeks (10) and one month (22).[7] The log-likelihood of the HAR model is defined as the benchmark level, the log-likelihoods of the other five density forecasting methods exceeding the benchmark are summarized in the table.

For IBM stock, the lognormal Black-Scholes model gives the highest log-likelihoods for all four horizons ranging from one day to one month, for both risk-neutral and transformed real-world density. The HAR model and the Heston model give similar

[7] For a horizon h, we set $T=h/252$ to calculate option implied densities.

forecasts for all four horizons both before and after applying transformations. The log-likelihoods for nonparametric transformation are always higher than those under risk-neutral measure for all methods and horizons, and the differences range from 66.3 to 192.8.

4.5 Diagnostic tests

The KS statistic tests if the densities are correctly specified under the i.i.d. assumption. Table 5 summarizes the *p*-values for the KS test for six density forecasting methods for four horizons for IBM. Since the null hypothesis is rejected at α significance level when $p<\alpha$, All the risk-neutral measure *p*-values reject the null hypothesis at 5% significance level, which might be due to the mis-specified risk-neutral densities that are conditionally normal. All nonparametric transformations have satisfactory *p*-values greater than 50%.

The BK test LR3 statistic tests the null hypothesis that the variables $y_i=\Phi^{-1}(u_i)$ are i.i.d. and follow a standard normal distribution, against the alternative hypothesis of a stationary, Gaussian, AR(1) process with no restrictions on the mean, variance and autoregressive parameters. Table 6 presents the BK test LR3 statistic, and the estimates of the variance and AR parameters for six density forecasting methods and four horizons for IBM.

The MLEs of the autoregressive parameters are between -0.01 and 0.01 for one-day horizon, hence there is no significant evidence of time-series dependence. However, the MLEs for one-week horizon range between -0.04 and -0.08, thus five of them reject the null hypothesis that the autoregressive parameter is 0 at 5% significance level. The longer two-weeks and one-month horizons also provide no evidence of dependent observations. The MLEs of the variance parameter are near one for correctly specified densities. The low estimates for one-day lognormal and Heston forecasts under Q measure might be explained by the fact that the risk-neutral standard deviations on average are higher than the historical standard deviations.

The LR3 test statistic is significant at 5% level when it exceeds 7.81. Table 6 indicates that the null hypothesis is rejected for all risk-neutral forecasts and one-week forecasts. The null hypothesis is accepted for all real-world forecasts for one day, two-weeks and one-month horizons. The significant values of LR3 test statistic might be attributed to the negative estimates of the AR parameter for one-week horizon, or the mis-specified risk-neutral density which is conditionally normal.

For one day horizon, two of the AG test statistics are insignificant at 5% level when the best method, nonparametric lognormal, is compared with the five alternatives, the AG test statistics equal -0.37 and 1.27 for tests against nonparametric HAR and nonparametric Heston methods. The AG test has similar test values and the same conclusion when the Newey-West adjustment is made to the estimated variance. The insignificant values become -0.33 and 0.99 when twenty autocorrelations are considered. The AG test results show that the best method for one week horizon is significantly better than two of the remaining five methods at the 5% level, and the best method is statistically better than one method at the 5% level for two weeks horizon, while the best method is statistically indifferent to the other methods at the longest, one month horizon, when the Newey-West adjustment is employed.

5. Conclusions

We compare density forecasts for the prices of Dow Jones 30 stocks, obtained from 5-minute high-frequency returns and daily option prices by using Heston, lognormal Black-Scholes, lognormal HAR-RV and transformed densities. Our comparison criterion is the log-likelihood of observed stock prices. For the sixty-eight combinations from seventeen stocks for four horizons, the lognormal Black-Scholes model gives the highest log-likelihoods for fifty-eight combinations. The HAR-RV model and the Heston model have similar forecast accuracy for different horizons, either before or after applying a transformation which enhances the densities.

Jiang and Tian (2005) suggest that daily option prices are more informative than daily and intraday index returns when forecasting the volatility of the S&P 500 index over horizons from one to six months. Shackleton et al. (2010) similarly imply that option prices are more informative when based on mid-term forecast horizons due to the forward-looking nature of option prices. They only use option prices for the contracts with maturities of more than one week, hence the short horizons of one day and one week density forecasts are extrapolations which are not backed by active trading. They state that the historical density is best for the one day horizon as we can forecast the volatility for tomorrow accurately by calculating the realized variance from recent high-frequency returns.

Most density research only focuses on either risk-neutral densities or ex post real-world density forecasts for horizons matching option expiry dates, while we generate ex ante real-world densities for different forecast horizons. We use a nonparametric transformation to transform the risk-neutral density into real-world density. The log-likelihoods for the nonparametric transformation are always higher than those under the risk-neutral measure for all methods and horizons. The nonparametric transformation also gives better diagnostic test results. Hence central banks, risk managers and other decision takers should not merely look at risk-neutral densities, but should also obtain more accurate predictions by using risk transformations applied to risk-neutral densities. The relatively unsatisfactory performance of the Heston model for individual firms might be attributed to the illiquidity of the OTM options. Compared to the index, the individual firm stocks options have fewer strikes that are traded.

Appendix 1. Assumptions about prices, dividends and options

Stock prices jump when dividends are assigned. We apply the Heston dynamics to futures prices which do not jump. We also need to assume all synthetic futures prices have the same dynamics. We assume futures and options contracts expire at time T_1, and there is a dividend at time τ_1 between time 0 and time T_1. The second expiry time for futures and options is T_2 and there is another dividend at τ_2 between time T_1 and T_2. We can use the same dynamics for all futures from simple dividend assumptions; this is easy for continuous dividends but harder for discrete dividends. We denote the futures price at t for delivery at T to be $F_{t,T}$. Our discussion below refers to dividend constants $c_1, c_2, \ldots$, which do not need to be calculated.

We assume, at time t before time τ_i, that the expected dividends are

$$
\begin{aligned}
E_t[dividend\ at\ \tau_i] &= c_1 e^{r(\tau_1 - t)} S_t && \tau = 1, t < \tau_1 \\
&= c_2(1 - c_1) e^{r(\tau_2 - t)} S_t && \tau = 2, t < \tau_2 \\
&= c_3(1 - c_1)(1 - c_2) e^{r(\tau_3 - t)} S_t && \tau = 3, t < \tau_3
\end{aligned}
$$

etc. We assume futures prices are set by no-arbitrage conditions, so

$$F_{t,T} = e^{r(T-t)}[S_t - PV(expected\ dividends\ from\ t\ to\ T)].$$

Then for the first contract

$$
\begin{aligned}
F_{t,T_1} &= e^{r(T_1 - t)}\left[S_t - e^{-r(\tau_1 - t)} c_1 e^{r(\tau_1 - t)} S_t\right] \\
&= (1 - c_1) e^{r(T_1 - t)} S_t && 0 \le t < \tau_1, \\
&= e^{r(T_1 - t)} S_t && \tau_1 \le t \le T_1.
\end{aligned}
$$

Then we have

$$ln\left(F_{t,T_1}/S_t\right) = ln(1-c_1) + r(T_1 - t) \qquad 0 \le t < \tau_1,$$

$$= r(T_1 - t) \qquad \tau_1 \le t \le T_1.$$

Thus

$$d\left(lnF_{t,T_1}\right) = d(lnS_t) - rdt \qquad t \ne \tau_1$$

Also S_t jumps down by $c_1 S_{\tau_1}$ at time $t=\tau_1$, but F_{t,T_1} does not jump at $t=\tau_1$.

Similarly, for the second contract

$$F_{t,T_2} = e^{r(T_2-t)}\left[S_t - e^{-r(\tau_1-t)}c_1 e^{r(\tau_1-t)}S_t - e^{-r(\tau_2-t)}c_2(1-c_1)e^{r(\tau_2-t)}S_t\right]$$

$$= e^{r(T_2-t)}(1-c_1)(1-c_2)S_t \qquad 0 \le t < \tau_1,$$

$$= e^{r(T_2-t)}(1-c_2)S_t \qquad \tau_1 \le t < \tau_2,$$

$$= e^{r(T_2-t)}S_t \qquad \tau_2 \le t \le T_2.$$

Hence we have

$$d\left(lnF_{t,T_2}\right) = d(lnS_t) - rdt \qquad t \ne \tau_1, \tau_2,$$

$$= d\left(lnF_{t,T_1}\right) \qquad 0 \le t \le T_1.$$

And we also have

$$\frac{F_{t,T_2}}{F_{t,T_1}} = e^{r(T_2-T_1)}(1-c_2) \qquad 0 \le t \le T_1.$$

We estimate the Heston parameters for the prices of European options which expire at $T_1, T_2, \ldots, T_N$, and strike prices are available as $K_{i,j}$, with $1 \le i \le N$ and $1 \le j \le n_i$. At time 0 we have Black-Scholes implied volatilities $\sigma_{i,j}$, these give market prices from the standard formula for options on futures,

$$c_{i,j} = c_B\left(F_{0,T_i}, T_i, K_{i,j}, r, \sigma_{i,j}\right).$$

Here we have

$$F_{0,T_i} = e^{rT_i}[S_0 - PV(expected\ dividends\ from\ zero\ to\ T_i]$$

and S_0 is the spot price.

Our target is to estimate the Heston parameters θ as:

$$\hat{\theta} = arg\min_{\theta} \sum_i \sum_j \left[c_{ij} - c_{Heston}\left(F_{0,T_i}, T_i, K_{ij}, r, \theta\right)\right]^2$$

At time 0 and for any future time τ, we can obtain the density of $S_\tau = F_{\tau,\tau}$ by evaluating the Heston-density with initial price $F_{0,\tau}$ and parameters $\hat{\theta}$.

Appendix 2. Tables and Figures

Table 1

List of Dow Jones 30 constituent stocks as at 24th September 2012.

Number	Company	Exchange	Symbol	Industry	Date added
1	3M	NYSE	MMM	Conglomerate	1976/8/9
2	Alcoa	NYSE	AA	Aluminum	1959/6/1
3	American Express	NYSE	AXP	Consumer finance	1982/8/30
4	AT&T	NYSE	T	Telecommunication	1999/11/1
5	Bank of America	NYSE	BAC	Banking	2008/2/19
6	Boeing	NYSE	BA	Aerospace and defense	1987/3/12
7	Caterpillar	NYSE	CAT	Construction and mining equipment	1991/5/6
8	Chevron Corporation	NYSE	CVX	Oil & gas	2008/2/19
9	Cisco Systems	NASDAQ	CSCO	Computer networking	2009/6/8
10	Coca Cola	NYSE	KO	Beverages	1987/3/12
11	DuPont	NYSE	DD	Chemical industry	1935/11/20
12	ExxonMobil	NYSE	XOM	Oil & gas	1928/10/1
13	General Electric	NYSE	GE	Conglomerate	1907/11/7
14	Hewlett-Packard	NYSE	HPQ	Computers & technology	1997/3/17
15	The Home Depot	NYSE	HD	Home improvement retailer	1999/11/1
16	Intel	NASDAQ	INTC	Semiconductors	1999/11/1
17	IBM	NYSE	IBM	Computers & technology	1979/6/29
18	Johnson & Johnson	NYSE	JNJ	Pharmaceuticals	1997/3/17
19	JPMorgan Chase	NYSE	JPM	Banking	1991/5/6
20	McDonald's	NYSE	MCD	Fast Food	1985/10/30
21	Merck	NYSE	MRK	Pharmaceuticals	1979/6/29
22	Microsoft	NASDAQ	MSFT	Software	1999/11/1
23	Pfizer	NYSE	PFE	Pharmaceuticals	2004/4/8
24	Procter & Gamble	NYSE	PG	Consumer goods	1932/5/26
25	Travelers	NYSE	TRV	Insurance	2009/6/8
26	UnitedHealth Group	NYSE	UNH	Managed health care	2012/9/24
27	United Technologies Corporation	NYSE	UTX	Conglomerate	1939/3/14
28	Verizon	NYSE	VZ	Telecommunication	2004/4/8
29	Wal-Mart	NYSE	WMT	Retail	1997/3/17
30	Walt Disney	NYSE	DIS	Broadcasting and entertainment	1991/5/6

Table 2

Summary statistics for IBM option data. The information about out-of-the-money (OTM) and at-the-money (ATM) options on IBM stock from 2003 to 2012.

	Total	Average per day	Maximum per day	Minimum per day
Calls	47709	19	46	6
Puts	61402	25	74	5
Total	109111	44	115	12

Moneyness/maturity	S/K	<1 month	Between 1 and 6 months	>6 months	Subtotal
Deep OTM put	>1.05	6462	30100	13596	50158
		(5.92%)	(27.59%)	(12.46%)	(45.97%)
OTM put	1.01-1.05	2040	5123	1839	9002
		(1.87%)	(4.70%)	(1.69%)	(8.25%)
At/near the money	0.99-1.01	1049	2641	973	4663
		(0.96%)	(2.42%)	(0.89%)	(4.27%)
OTM call	0.95-0.99	2278	5733	2330	10341
		(2.09%)	(5.25%)	(2.14%)	(9.48%)
Deep OTM call	<0.95	3168	20393	11386	34947
		(2.90%)	(18.69%)	(10.44%)	(32.03%)
Subtotal		14997	63990	30124	109111
		(13.74%)	(58.65%)	(27.61%)	(100.00%)

Table 3

Summary statistics for risk-neutral calibrated parameters for IBM and across all stocks. Estimates are summarized for the risk-neutral dynamics (2). The parameters are estimated each day from 2003 to 2012, from the OTM and ATM options, through minimizing the MSE of the fitted option prices. We apply the constraint $\kappa \leq 36$.

	κ	θ	σ	ρ	v_0
IBM					
Mean	1.6861	0.5042	1.2038	-0.6723	0.0653
Median	0.1661	0.3457	0.8617	-0.6652	0.0444
Standard deviation	3.6779	0.4201	2.1596	0.1051	0.0726
Averages across all firms					
Mean	3.0401	0.4037	1.9675	-0.6331	0.1081
Median	1.1136	0.2308	1.0267	-0.6305	0.0692
Standard deviation	5.2434	0.3594	5.6694	0.1462	0.1206

Table 4

Log-likelihoods for overlapping forecast. The numbers shown are the log-likelihoods of the HAR untransformed density forecasts and the log-likelihoods of the other forecasts in excess of the HAR benchmark values. The letter *Q* defines untransformed and risk-neutral densities, while the letter *P* denotes nonparametric transformation of the *Q* densities defined by (26). The numbers in bold in each row refer to the best method with the highest log-likelihood for the selected forecast horizon.

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
IBM							
1 day	2487	-4312.5	124.1	33.0	**128.5**	-9.3	113.2
1 week	2483	-6419.1	157.3	100.1	**217.4**	100.9	167.2
2 weeks	2478	-7222.1	189.3	78.1	**270.9**	76.1	176.2
1 month	2466	-8232.5	179.9	77.2	**257.6**	65.7	151.1
Alcoa							
1 day	2487	-1616.5	81.3	68.2	**117.1**	6.9	77.7
1 week	2483	-3687.9	60.9	77.8	**107.6**	7.0	82.9
2 weeks	2478	-4575.0	131.0	108.8	**161.8**	-5.1	111.1
1 month	2466	-5693.6	305.0	202.5	**332.0**	-17.7	239.7
Boeing							
1 day	2487	-3706.5	168.5	178.7	**208.5**	119.5	174.6
1 week	2483	-5644.3	110.8	115.3	**150.9**	44.4	134.8
2 weeks	2478	-6439.4	100.4	72.3	**119.8**	-57.5	109.2
1 month	2466	-7387.4	**158.9**	57.3	147.5	-186.9	113.4
Cisco							
1 day	2487	-1235.1	260.2	185.3	**269.8**	129.3	242.8
1 week	2483	-3218.6	161.0	223.4	**266.9**	92.9	226.9
2 weeks	2478	-3966.3	109.0	130.3	**189.3**	-29.9	115.3
1 month	2466	-4904.4	81.3	68.1	**133.3**	-181.3	50.0

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Disney							
1 day	2487	-1787.1	163.6	166.3	**231.4**	98.9	200.7
1 week	2483	-3569.7	69.3	85.0	**131.8**	-38.9	93.0
2 weeks	2478	-4368.9	105.3	76.2	**191.5**	-117.6	118.9
1 month	2466	-5342.1	169.2	43.7	**237.0**	-255.5	204.5
General Electric							
1 day	2487	-2636.3	**185.2**	-150.7	-8.6	-702.0	-97.3
1 week	2483	-3330.3	208.8	347.9	**385.7**	135.2	267.3
2 weeks	2478	-3910.0	62.0	75.6	**109.2**	-166.0	-2.3
1 month	2466	-5220.7	160.3	396.3	**453.1**	36.2	335.6
Home Depot							
1 day	2487	-2009.2	78.3	40	**98.5**	-222.6	-59.9
1 week	2483	-4014.8	54.2	77.6	**110.8**	-261.3	-151.3
2 weeks	2478	-4815.7	72.8	46.9	**117.4**	-238.3	-162.7
1 month	2466	-5821.9	92.6	26.2	**136.7**	-321.8	-221.1
Hewlett Packard							
1 day	2487	-2395.3	356.2	238.3	**401.3**	257.6	386.3
1 week	2483	-4299.4	255.6	193.4	**316.8**	248.5	311.5
2 weeks	2478	-5035.9	200	127.8	**245.1**	180	232.9
1 month	2466	-6095.3	280.6	136.7	**332.6**	244.1	302.2
Intel							
1 day	2487	-2395.3	**85.5**	7.6	77.2	-1.9	71
1 week	2483	-4299.4	75.4	65.9	**104.8**	52.1	91.6
2 weeks	2478	-5035.9	**83.6**	26.7	80.3	16	71.8
1 month	2466	-6101	**86.8**	-34.2	51.3	5.5	38.7

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
Johnson & Johnson							
1 day	2487	-2395.3	**171.9**	62.6	163.7	-58.3	106.1
1 week	2483	-4299.4	146.5	59.3	**174.8**	-113.7	46.2
2 weeks	2478	-5035.9	105.4	-4.5	**145.9**	-208.4	-31.4
1 month	2466	-6101.0	104.6	-37.5	**112.7**	-286.4	-137.6
JP Morgan Chase							
1 day	2487	-2395.3	**101.7**	15.1	95.6	5.5	87.6
1 week	2483	-4299.4	62.5	28.4	**93.7**	-2.3	61.3
2 weeks	2478	-5035.9	53.1	9.7	**106.6**	-50.4	39.8
1 month	2466	-6101.0	43.5	-33.3	**79.5**	-112.9	-4.6
McDonald's							
1 day	2487	-2395.3	135.9	92.7	**167.1**	57.2	148.5
1 week	2483	-4299.4	207	429.6	**516.1**	384.1	453.7
2 weeks	2478	-5035.9	78.7	-31.7	**85.8**	-72.6	19.7
1 month	2466	-6101	**153.8**	-42.3	121.8	-65.4	58.9
Merck							
1 day	2487	-2395.3	790.4	235.6	**850.8**	570.2	751.0
1 week	2483	-4299.4	553.8	137.1	**609.6**	353.0	492.6
2 weeks	2478	-5035.9	582.3	102.2	**648.8**	459.7	586.7
1 month	2466	-6101.0	431.9	-20.5	**464.6**	266.2	404.9
Pfizer							
1 day	2487	-2395.3	197.6	91.5	**222.3**	1.5	180.2
1 week	2483	-4299.4	82.5	57.9	**113.1**	3.5	71.0
2 weeks	2478	-5035.9	64.9	30.3	**96.2**	-33.7	48.3
1 month	2466	-6101.0	49.1	12.9	**101.3**	-123.7	13.7

Forecast horizon	No. of obs.	HAR		Lognormal		Heston	
		Q	*P*	*Q*	*P*	*Q*	*P*
AT&T							
1 day	2487	-2395.3	**85.1**	74.0	121.9	-597.4	-284.7
1 week	2483	-4299.4	57.0	59.6	**105.7**	-754.4	-494.7
2 weeks	2478	-5035.9	79.0	86.5	**139.0**	-703.0	-435.1
1 month	2466	-6101.0	116.6	109.4	**171.9**	-786.2	-491.3
Walmart							
1 day	2487	-2395.3	183.2	127.4	**213.7**	83.6	183.5
1 week	2483	-4299.4	69.7	56.2	**102.0**	-25.0	50.7
2 weeks	2478	-5035.9	40.5	5.4	**71.1**	-98.0	-7.9
1 month	2466	-6101.0	29.3	-36.8	**48.9**	-140.6	-31.5
American Express							
1 day	2487	-3046.9	271.7	154.4	267.2	53.1	**305**
1 week	2483	-4829.3	117.9	89	**153.4**	-75.1	12.1
2 weeks	2478	-5608.6	94.9	52.3	**144.5**	-40.5	10.5
1 month	2466	-6456.2	92.7	33.7	**145.5**	-91.9	9.9

Table 5

KS test results for IBM overlapping forecast. The numbers are the *p*-values of the KS test for the null hypothesis that the terms u_t are uniformly distributed. The letter Q defines risk-neutral densities, while the letter P denotes nonparametric transformation of the real-world densities defined by (26). * indicates that the *p*-values are greater than 50%. The null hypothesis is rejected at α significance level when $p<\alpha$.

Forecast horizon	No. of obs.	HAR (%)		Lognormal (%)		Heston (%)	
		Q	P	Q	P	Q	P
1 day	2487	0.42	*	0.00	*	0.00	*
1 week	2483	0.01	*	0.00	*	0.00	*
2 weeks	2478	0.00	*	0.00	*	0.00	*
1 month	2466	0.00	*	0.00	*	0.00	*

Table 6

BK test results for IBM overlapping forecast. The null hypothesis that the variables $y_i=\Phi^{-1}(u_i)$ are i.i.d. and follow a standard normal distribution is tested against the alternative hypothesis of a stationary, Gaussian, AR(1) process with no restrictions on the mean, variance and autoregressive parameters. The numbers are the LR3 test statistic, and the estimates of the variance and AR parameters. * suggests that the null hypothesis is rejected at 5% significance level when LR3>7.81.

Forecast horizon		HAR		Lognormal		Heston	
		Q	P	Q	P	Q	P
1 day	AR	-0.01	-0.01	0.01	0.00	0.01	0.00
	Variance	1.17	0.97	0.79	0.97	0.78	0.97
	LR3	42.19*	1.74	74.23*	1.36	75.42*	1.47
1 week	AR	-0.04	-0.07	-0.06	-0.08	-0.05	-0.06
	Variance	1.18	0.96	0.86	0.96	0.84	0.96
	LR3	50.06*	15.07*	44.06*	19.08*	44.69*	11.08*
2 weeks	AR	0.01	0.00	0.01	0.00	0.01	0.01
	Variance	1.11	0.96	0.82	0.96	0.81	0.96
	LR3	30.61*	2.42	67.22*	2.51	56.32*	2.54
1 month	AR	0.01	-0.02	0.01	-0.02	-0.02	-0.01
	Variance	1.12	0.96	0.86	0.96	0.90	0.96
	LR3	44.77*	3.42	62.91*	4.12	23.80*	2.64

Figure 1. Heston, lognormal and HAR one-day ahead density forecasts for IBM on January 2nd 2003.

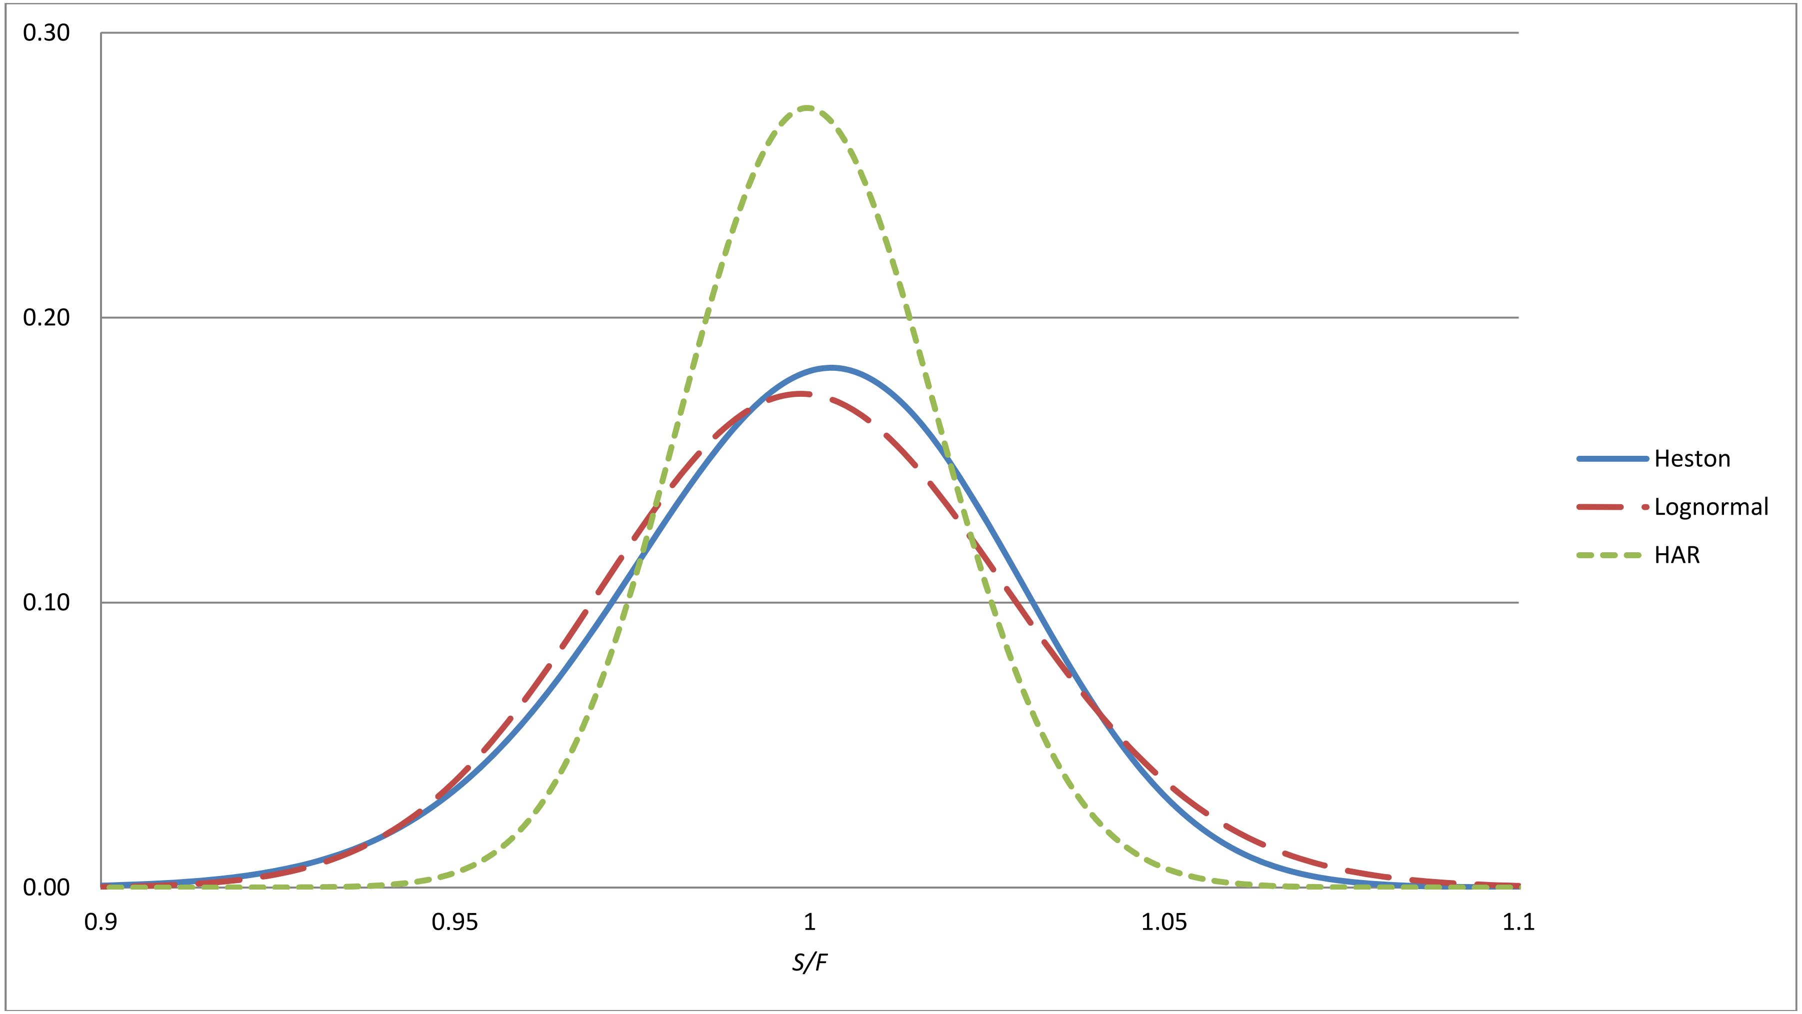

Figure 2. Heston, lognormal and HAR one-month ahead density forecasts for IBM on January 2nd 2003.

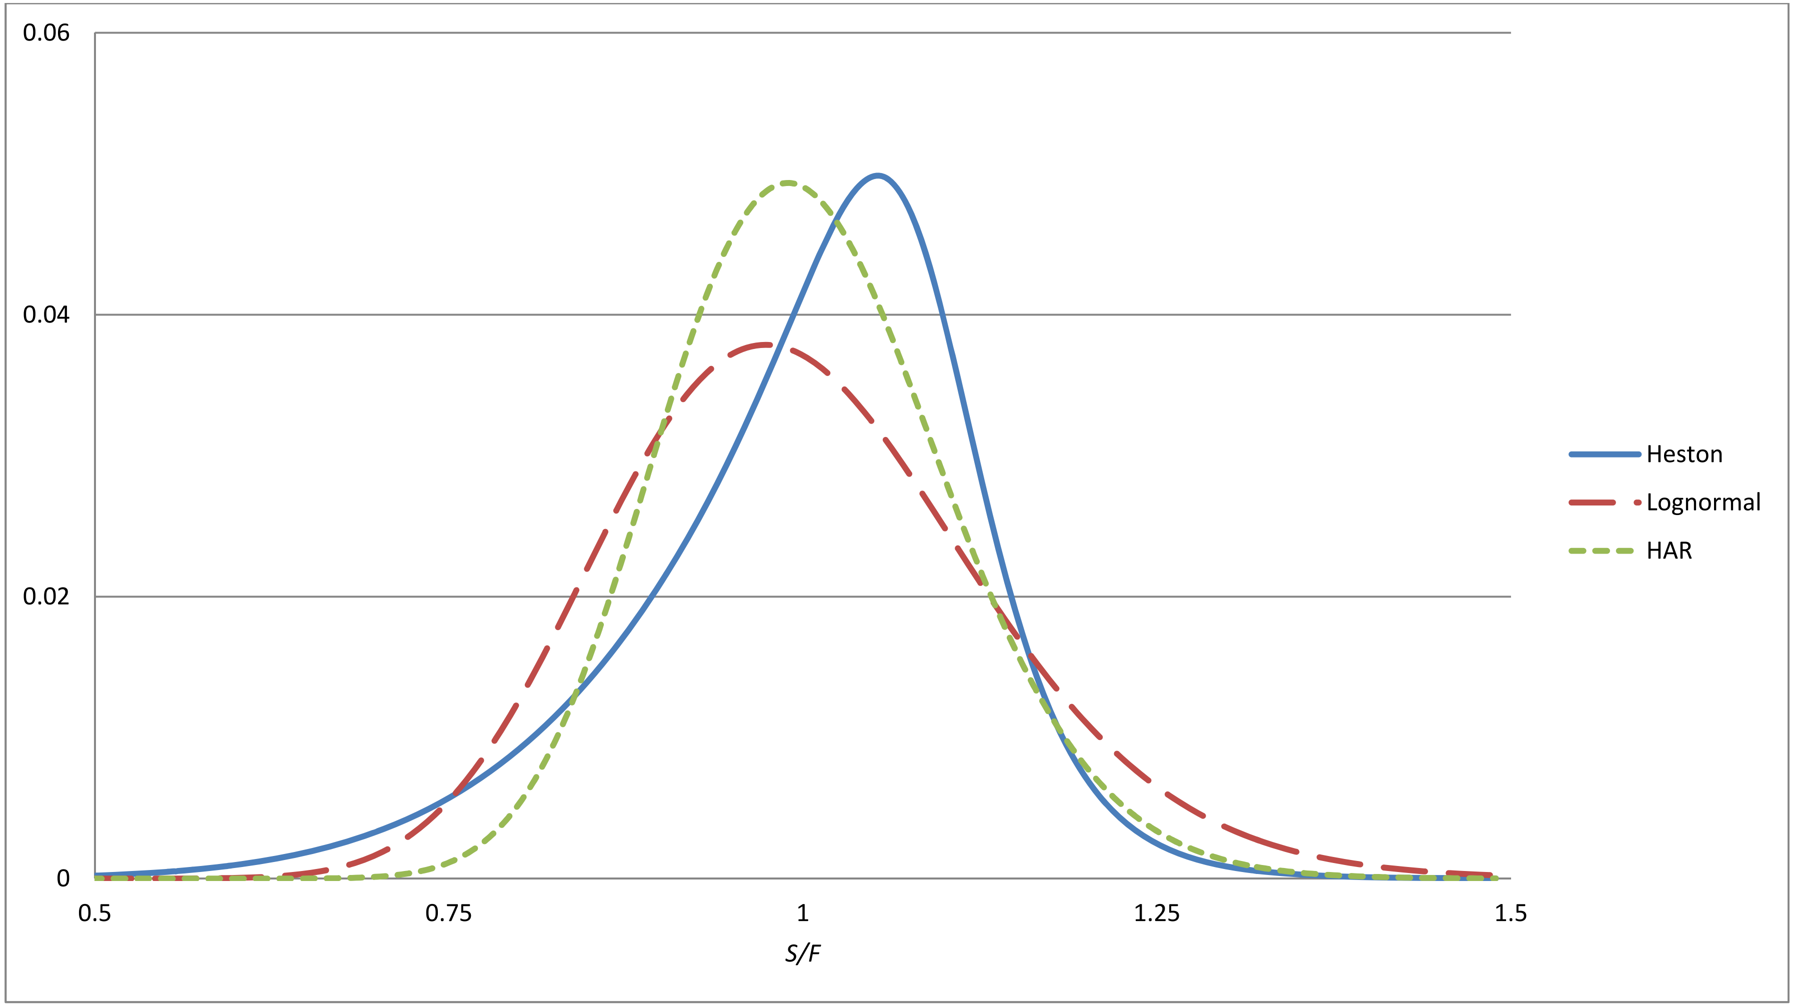

Figure 3. Function $\tilde{C}(u) - u$ for one-day ahead forecasts from the Heston model and a nonparametric transformation for IBM.

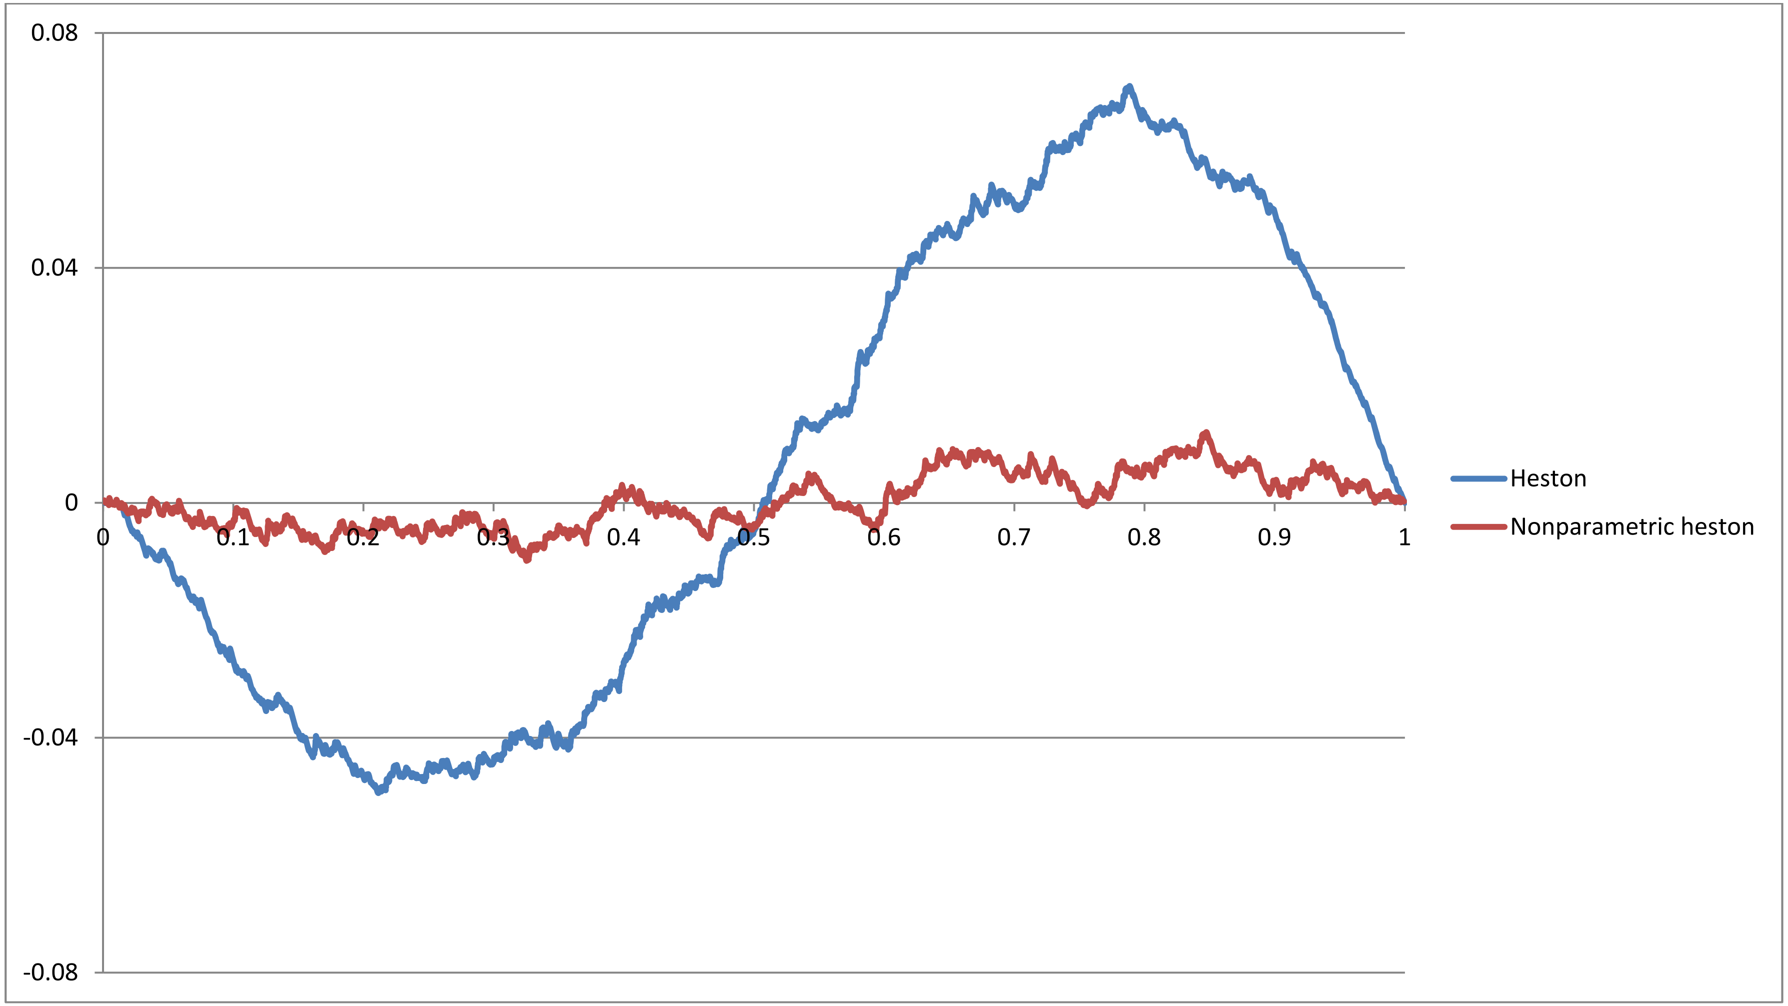

Figure 4. Function $\tilde{C}(u) - u$ for one-day ahead forecasts from the Black-Scholes model and a nonparametric transformation for IBM.

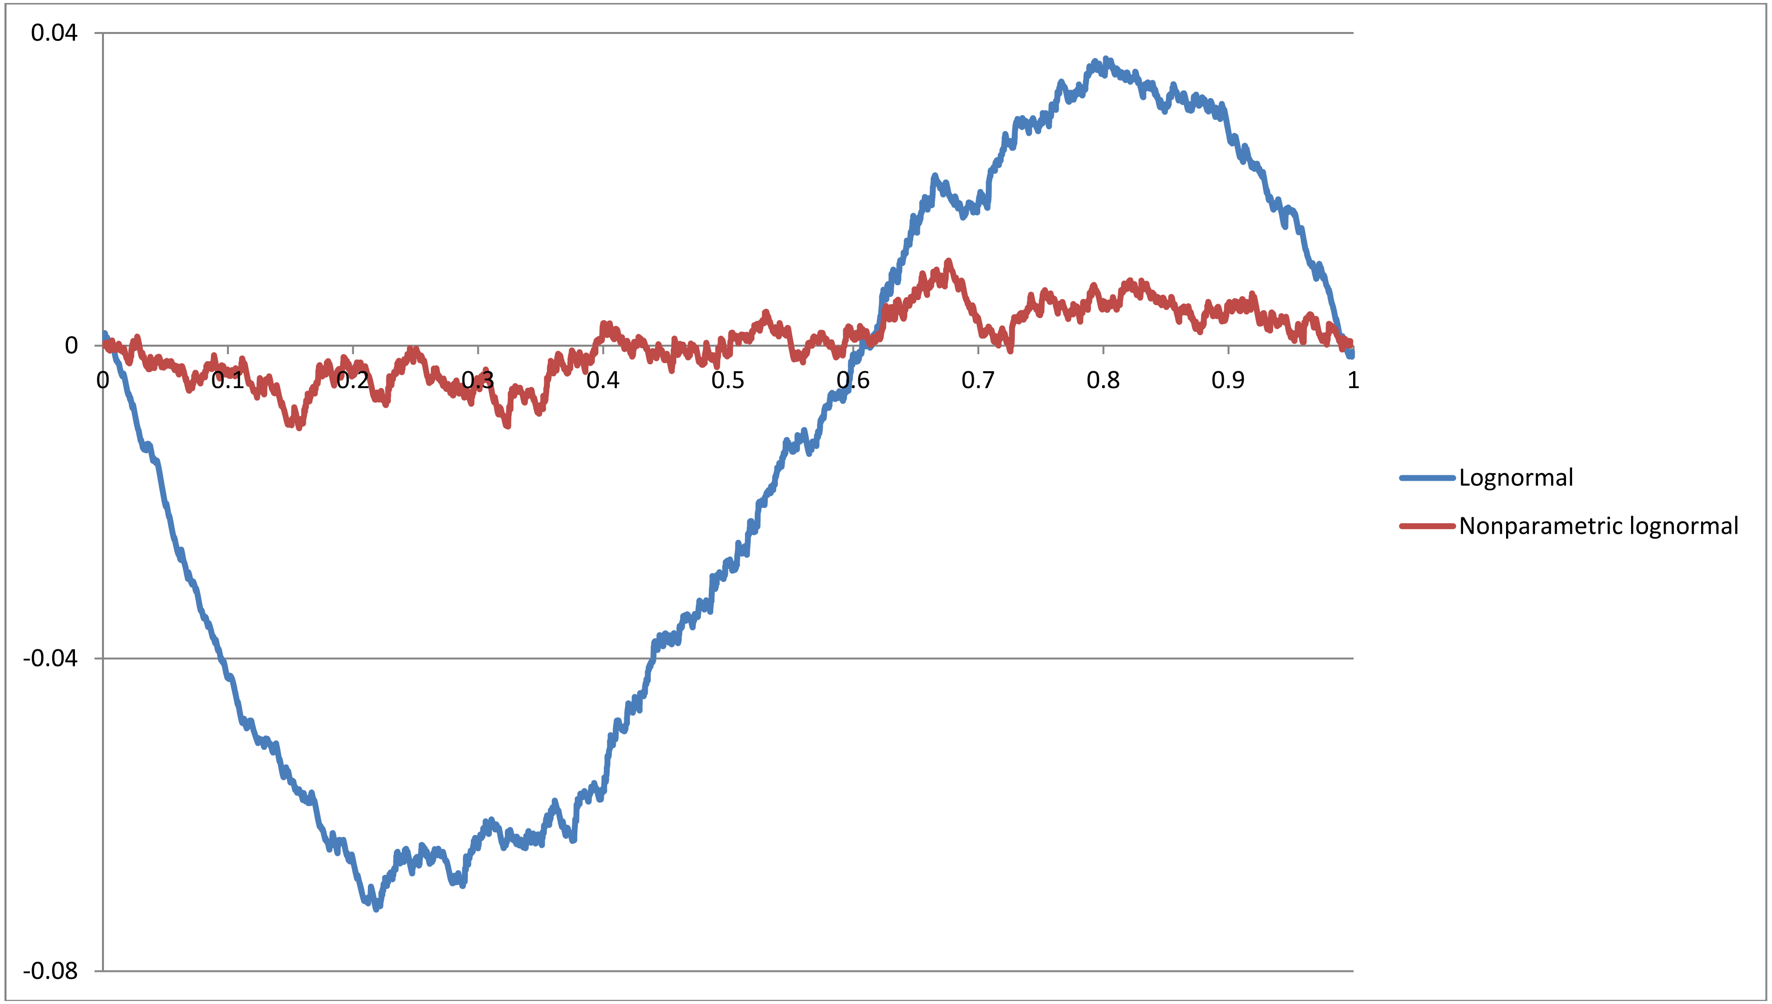

Figure 5. Nonparametric calibration densities $\hat{c}(u)$ from one-day ahead Heston forecasts for IBM.

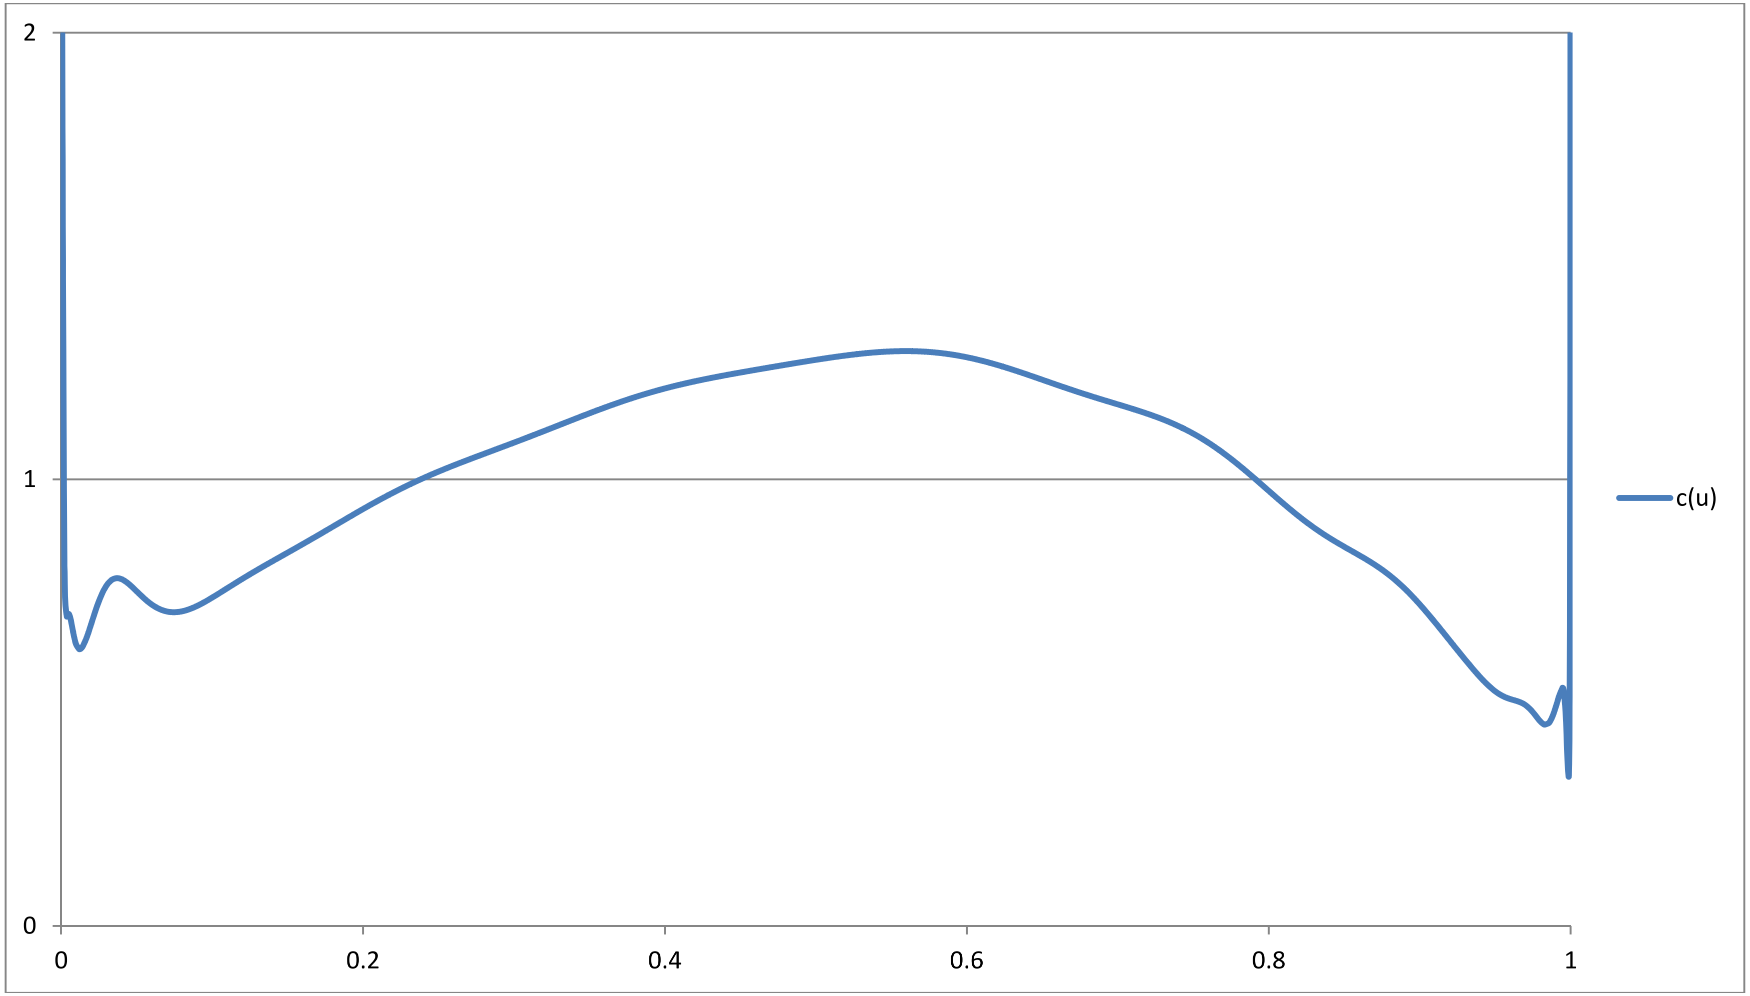

Figure 6. Nonparametric calibration densities $\hat{c}(u)$ from one-day ahead Black-Scholes lognormal forecasts for IBM.

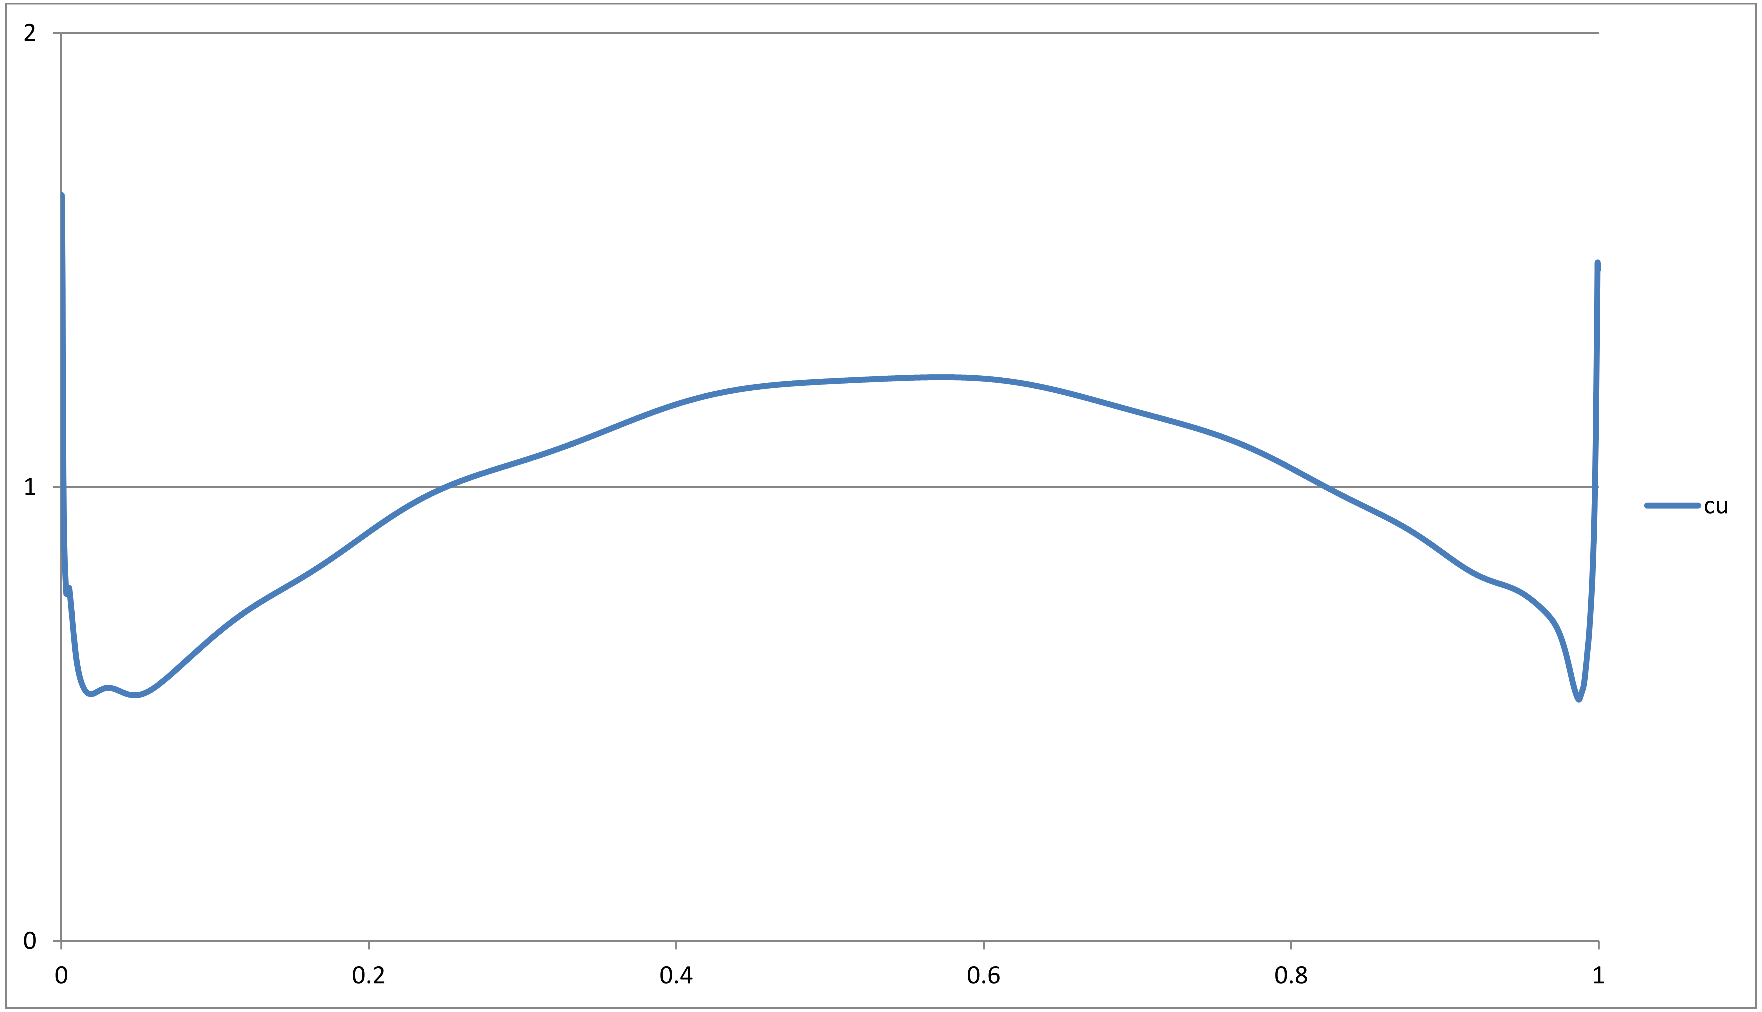

References

Ait-Sahalia, Y. and A.W. Lo, 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53, 499-547.

Amisano, G. and R. Giacomini, 2007. Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business and Economic Statistics* 25, 177-190.

Anagnou-Basioudis, I., M. Bedendo, S.D. Hodges and R. Tompkins, 2005. Forecasting accuracy of implied and GARCH-based probability density functions. *Review of Futures Markets* 11, 41-66.

Andersen, T., T. Bollerslev and F.X. Diebold, 2007. Roughing it up: including jump components in measuring, modeling and forecasting asset return volatility. *Review of Economics and Statistics* 89, 701-720.

Bakshi, G., N. Kapadia and D.B. Madan, 2003. Stock return characteristics, skew laws and the differential pricing of individual equity options. *Review of Financial Studies* 16, 101-143.

Bandi F.M. and J.R. Russell, 2006. Separating microstructure noise from volatility. *Journal of Financial Economics* 79, 655-692.

Bao, Y., T.-H. Lee and B. Satoglu, 2007. Comparing density forecast models. *Journal of Forecasting* 26, 203-225.

Bates, D., 1996. Jumps and stochastic volatility: exchange rate processes implicit in Deutsche mark options. *Review of Financial Studies* 9, 69-107.

Bates, D., 2000. Post-’87 crash fears in S&P 500 futures options. *Journal of Econometrics* 94, 181-238.

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

Blair, B.J., S.-H. Poon and S.J. Taylor, 2001. Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high frequency index returns. *Journal of Econometrics* 105, 5-26.

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.

Bunn, D.W., 1984. Applied Decision Analysis. McGraw-Hill.

Busch, T., B.J. Christensen and M.Ø. Nielsen, 2011. The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets. *Journal of Econometrics* 160, 48–57.

Carr, P. and L. Wu, 2003. Finite moment log stable process and option pricing. *Journal of Finance* 58, 753-777.

Carr, P. and L. Wu, 2009. Variance risk premiums. *Review of Financial Studies* 22, 1311-1341.

Carr, P. and L. Wu, 2010. Stock options and credit default swaps: a joint framework for valuation and estimation. *Journal of Financial Econometrics* 8, 409-449.

Christoffersen, P. and Jacobs, K., 2004. The importance of the loss function in option valuation. *Journal of Financial Economics* 72, 291-318.

Christoffersen, P., Jacobs, K. and Mimouni, K., 2006. Models for S&P 500 dynamics: evidence from realized volatility, daily returns and option prices. McGill University.

Clements, M.P., 2004. Evaluating the Bank of England density forecasts of inflation. *Economic Journal* 114, 855-877.

Clements M.P. and J. Smith, 2000. Evaluating density forecasts of linear and non-linear models: applications to output growth and unemployment. *Journal of Forecasting* 19, 255-276.

Cox, J.C., J.E. Ingersoll and S.A. Ross, 1985. A theory of the term structure of interest rates. *Econometrica* 53, 385–407.

Corsi, F., 2009. A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7, 174-196.

Diebold, F.X. and R.S. Mariano, 1995. Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253-263.

Diebold, F.X., A. Gunther and A.S. Tay, 1998. Evaluating density forecasts with applications to financial risk management. *International Economic Review* 39, 863-883.

Duffie, D., J. Pan, and K.J. Singleton, 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343-1476.

Eraker, B., 2004. Do equity prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59, 1367-1403.

Eraker, B., M.S. Johannes, and N.G. Polson, 2003. The impact of jumps in returns and volatility. *Journal of Finance* 53, 1269-1300.

Giot, P. and S. Laurent, 2007. The information content of implied volatility in light of the jump/continuous decomposition of realized volatility. *Journal of Futures Markets* 27, 337-359.

Guidolin, M. and A. Timmermann, 2005. Economic implication of bull and bear regimes in UK stock and bond returns. *Economic Journal* 115, 111-143.

Granger, C.W.J., P. Newbold, 1976. Forecasting economic time series. Academic Press, New York.

Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6,327– 343.

Hog E. and L. Tsiaras, 2010. Density forecasts of crude-oil prices using option implied and arch-type models. *Journal of Futures Markets* 31, 727-754.

Huang J., L. Wu, 2004. Specification analysis of option pricing models based on time-changed levy processes. *Journal of Finance* 59, 1405-1439.

Ivanova, V. and J.M. Gutierrez, 2014. Interest rate forecasts, state price densities and risk premium from Euribor options. *Journal of Banking and Finance* 48, 210-223.

Jackwerth, J.C. and Rubinstein M., 1996. Recovering probability distributions from option prices. *Journal of Finance* 51, 1611-1631.

Jiang G.J. and Y.S. Tian, 2005. The model-free implied volatility and its information content. *Review of Financial Studies* 13, 433-451.

Jondeau E. and M. Rockinger, 2000. Reading the smile: the message conveyed by methods which infer risk neutral densities. *Journal of International Money and Finance* 19, 885-915.

Kang B.J. and T.S. Kim, 2006. Option-implied risk preferences: an extension to wider classes of utility functions. *Journal of Financial Markets* 9, 180-198.

Kendall, M.G., A. Stuart and J.K. Ord, 1983. The advanced theory of statistics, vol.3, 4th edition. London: Charles Griffin.

Liu X., M.B. Shackleton, S.J. Taylor and X. Xu, 2007. Closed-form transformations from risk-neutral to real-world densities. *Journal of Banking and Finance* 31, 1501-1520.

Martens M. and J. Zein, 2004. Predicting financial volatility: High-frequency time-series forecast vis-à-vis implied volatility. *Journal of Futures Markets* 24, 1005-1028.

Pan, J., 2002. The jump risk premium implicit in options: evidence from an integrated time-series study. *Journal of Financial Economics* 63, 3-50.

Pong, S.H. and C. Granger, 2003. Forecasting volatility in financial markets: a review. *Journal of Economic Literature* 41, 478-539.

Pong, S.H., M.B. Shackleton, S.J. Taylor and X. Xu, 2004. Forecasting currency volatility: a comparison of implied volatilities and ar(fi)ma models. *Journal of Banking and Finance* 28, 2541-2563.

Ritchey, R.J., 1990. Call option valuation for discrete normal mixtures. *Journal of Financial Research* 13, 285-296.

Rosenblatt, M., 1952. Remarks on a multivariate transformation. *Annuals of Mathematical Statistics* 23, 470-472.

Shackleton, M.B., S.J. Taylor and P. Yu, 2010. A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices. *Journal of Banking and Finance* 34, 2678-2693.

Silverman, B.W., 1986. Density estimation for statistical and data analysis. Chapman Hall.

Taylor, S.J., 2005. Asset price dynamics, volatility and prediction, Princeton University Press.

Taylor, S.J., P.K. Yadav and Y. Zhang, 2010. The information content of implied volatilities and model-free volatility expectations: evidence from options written on individual stocks. *Journal of Banking and Finance* 34, 871-881.

Yun J., 2014. Out-of-sample density forecasts with affine jump diffusion models. *Journal of Banking and Finance* 47, 74-87.