Presentation at the

Frankfurt MathFinance Workshop

27 March 2006

A multi-horizon comparison of density forecasts for the S&P 500 using index returns and option prices

Mark B. Shackleton, Stephen J. Taylor and Peng Yu

Department of Accounting and Finance

Lancaster University, England

Overview of this presentation

- Context
- Contributions
- The density forecasts
- Econometric methods
- Data
- The empirical results
- Concluding remarks

Context

- Option markets provide market expectations about the future prices of the underlying asset.
- Option traders can make use of:
  - Historical prices for the underlying asset,
  - Other, forward-looking, information.
- Hence forecasts obtained from option prices may outperform alternatives that are obtained from historical prices.

- There is a considerable literature about volatility forecasting, that compares option-based and historical forecasts.

  Surveys in: Poon & Granger, 2003, Taylor, 2005.

- In contrast, there is almost no literature that compares density forecasts.

  Two recent examples are:

  Anagnou-Basioudis, Bedendo, Hodges & Tompkins, 2005.

  Liu, Shackleton, Taylor & Xu, 2005.

- Option-based forecasts of index volatility are often found to be more accurate than historical forecasts,
    - When daily returns define historical forecasts,
    - And even when intraday returns are used:

      Blair, Poon & Taylor, 2001,

      Martens & Zein, 2004,

      Jiang & Tian, 2005.

- We may conjecture that similar conclusions apply to density forecasts.

Contributions

- Estimate risk-neutral densities for <u>all horizons</u> from option prices.
    - S & P 500 index, each day from 1990 to 2004.
    - Price dynamics of Heston, 1993.
    - Applied previously by Bakshi, Cao & Chen, 1997.
    - Contrasts with a multitude of single-horizon methods:

      Jackwerth & Rubinstein, 1996, Melick & Thomas, 1997,

      Ait-Sahalia & Lo, 1998, Bliss & Panigirtzoglou, 2002, etc

- Evaluate ex ante risk-transformations from risk-neutral to real-world densities.
    - Three transformations, 2 parametric and 1 non-parametric.
    - Seven forecast horizons, from 1 day to 12 weeks.
    - Contrasts with ex post transformations, and horizons that coincide with option expiry dates, in:

      Bakshi, Kapadia & Madan, 2003,

      Bliss & Panigirtzoglou, 2004,

      Anagnou-Basioudis et al, 2005, Liu et al, 2005.

- Compare ARCH and option-based density forecasts.
  - Using either daily or intraday returns to produce historical forecasts.
  - Seven forecast horizons.
  - Previous comparisons only use daily returns and they are restricted to horizons that coincide with option expiry dates.

The density forecasts

- Historical – Obtained from current and past index levels.
- Risk-neutral, Q – Provided by current option prices.
- Real-world, P – Given by a transformation of a Q-density.
- Mixtures – A weighted combination of historical and option-based densities.

Historical densities

One period of time = forecast horizon.

t counts time periods.

Prices F_t . Returns $r_t = \log(F_t / F_{t-1})$.

Four ARCH specifications for the historical density $f(r_{t+1} | I_t)$,

- constant conditional mean μ,
- two specifications for conditional variances h_t,
- conditional density either normal or a t-density.

- Using one-period returns: $r_t = \mu + \varepsilon_t$

$$\frac{h_t}{N_t} = \omega + \frac{(\alpha_1 + \alpha_2 d_{t-1})\varepsilon_{t-1}^2 + \beta_{GJR} h_{t-1}}{N_{t-1}},$$

$$d_{t-1} = 1 \text{ if } \varepsilon_{t-1} < 0, \quad \text{otherwise } d_{t-1} = 0.$$

The term N_t represents the number of trading days during period *t*.

- Using intra-period returns:

$Intra_t$ is the total of some set of squared intra-period returns for period *t*.

$$\frac{h_t}{N_t} = \omega + \frac{\gamma Intra_{t-1} + \beta_{Intra} h_{t-1}}{N_{t-1}}.$$

Figure 7

One-day ahead density forecasts obtained from ARCH models on July 30th, 2001

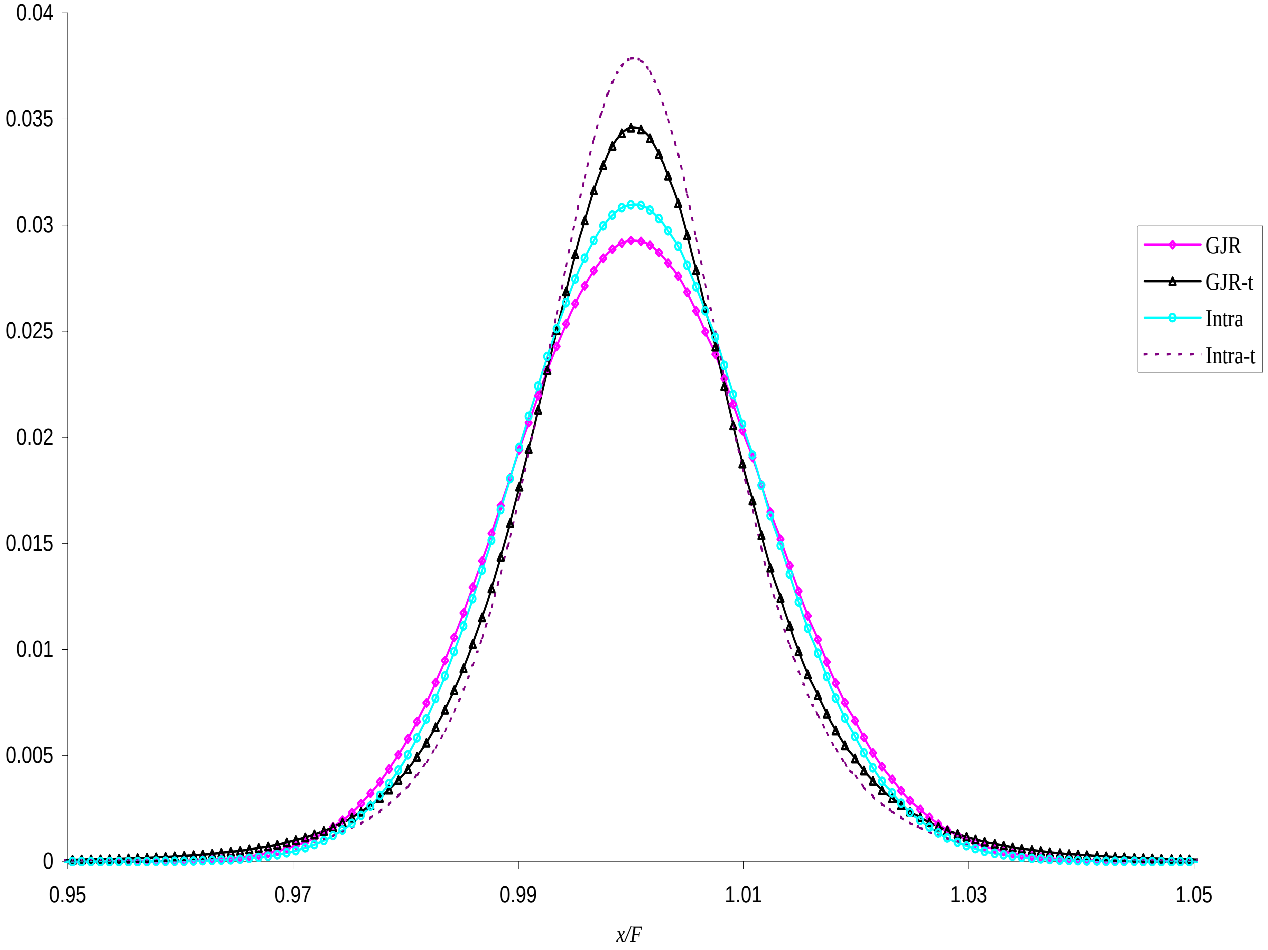

Risk-neutral densities

- Time t is now measured in years.
- The clock starts at time zero.
- Futures prices $F(t)$ are assumed to have Q-dynamics:

$$dF/F = \sqrt{V}\,dW_1,$$

$$dV = \kappa(\theta - V)dt + \xi\sqrt{V}\,dW_2,$$

$$\rho\,dt = dW_1 dW_2\,.$$

- Special case: $\kappa = \xi = 0$, so V is constant. Not in the paper yet.

- The characteristic function is analytic and defined by:

$$\tilde{g}(\psi) = E^{Q}[\, \exp(i\psi F(T)) \,|\, F(0), V(0) \,].$$

- Numerical integration provides the risk-neutral density:

$$g_{Q,T}(x) = \frac{1}{\pi x} \int_0^{\infty} \mathrm{Re}[\exp(-i\psi \log(x))\tilde{g}(\psi)]\, d\psi.$$

- ….. and also the probabilities $P_1(K)$ and $P_2(K)$ appearing in the formula

  for European, call option prices:

$$c(K) = e^{-rT}[F_0 P_1(K) - K P_2(K)],$$

where K is the strike price.

Real-world densities (drift-correction method)

Not included in the paper yet.

- Include a drift term $\lambda_1 FVdt$ in the diffusion for the price F.
- And a term $\lambda_2 Vdt$ in the diffusion for the variance V.
- The real-world, P-dynamics are then:

$$dF/F = \lambda_1 Vdt + \sqrt{V}\,dW_1,$$

$$dV = [\lambda_2 V + \kappa(\theta - V)]dt + \xi\sqrt{V}\,dW_2,$$

$$\rho\, dt = dW_1 dW_2 .$$

- To implement, must estimate λ_1 and λ_2 from historical data.

Real-world densities (calibration methods)

- At time 0, suppose we consider a forecast horizon *T*.
- We have these risk-neutral functions for the futures price F_T:
  - ➢ The density $g_{Q,T}(x)$ and
  - ➢ The cumulative distribution function (c.d.f.) $G_{Q,T}(x)$.
- Define a random variable by $U_T = G_{Q,T}(F_T)$. Suppose it has:
  - ➢ Real-world density $c_T(u)$ and
  - ➢ Real-world c.d.f. $C_T(u)$.

- Then the real-world functions for the futures price F_T are its:
  - Density $g_{P,T}(x) = g_{Q,T}(x)c_T(u)$, with $u = G_{Q,T}(x)$,
  - c.d.f. $G_{P,T}(x) = C_T(G_{Q,T}(x))$.

- These formulae can be implemented if we are willing to assume that the calibration function $C_T(u)$ is time-invariant and hence can be estimated.

- Some calibration functions can be linked to a utility function that is employed in a representative agent model (Liu et al, 2005).

- The c.d.f. of the Beta distribution defines our preferred, parametric calibration function:

$$C_T(u) = \int_0^u v^{j-1}(1-v)^{k-1} dv \,/\, B(j,k).$$

  - The calibration parameters, *j* and *k*, depend on the horizon *T*.
  - $B(j,k) = \Gamma(j)\Gamma(k) / \Gamma(j+k)$.

- The real-world density can be evaluated rapidly using numerical methods. It is:

$$g_{P,T}(x) = \frac{G_{Q,T}(x)^{j-1}(1-G_{Q,T}(x))^{k-1}}{B(j,k)} g_{Q,T}(x).$$

- A non-parametric calibration function can be estimated by a standard, kernel estimator.

- Figure 8 illustrates the two calibration transformations from a *Q*-density to a *P*-density.
    - The *Q*-densities have variances that are higher than those of credible *P*-densities.
    - Consequently, the transformations reduce the variance.

Figure 8

One-day ahead density forecasts obtained from options on July 30th, 2001

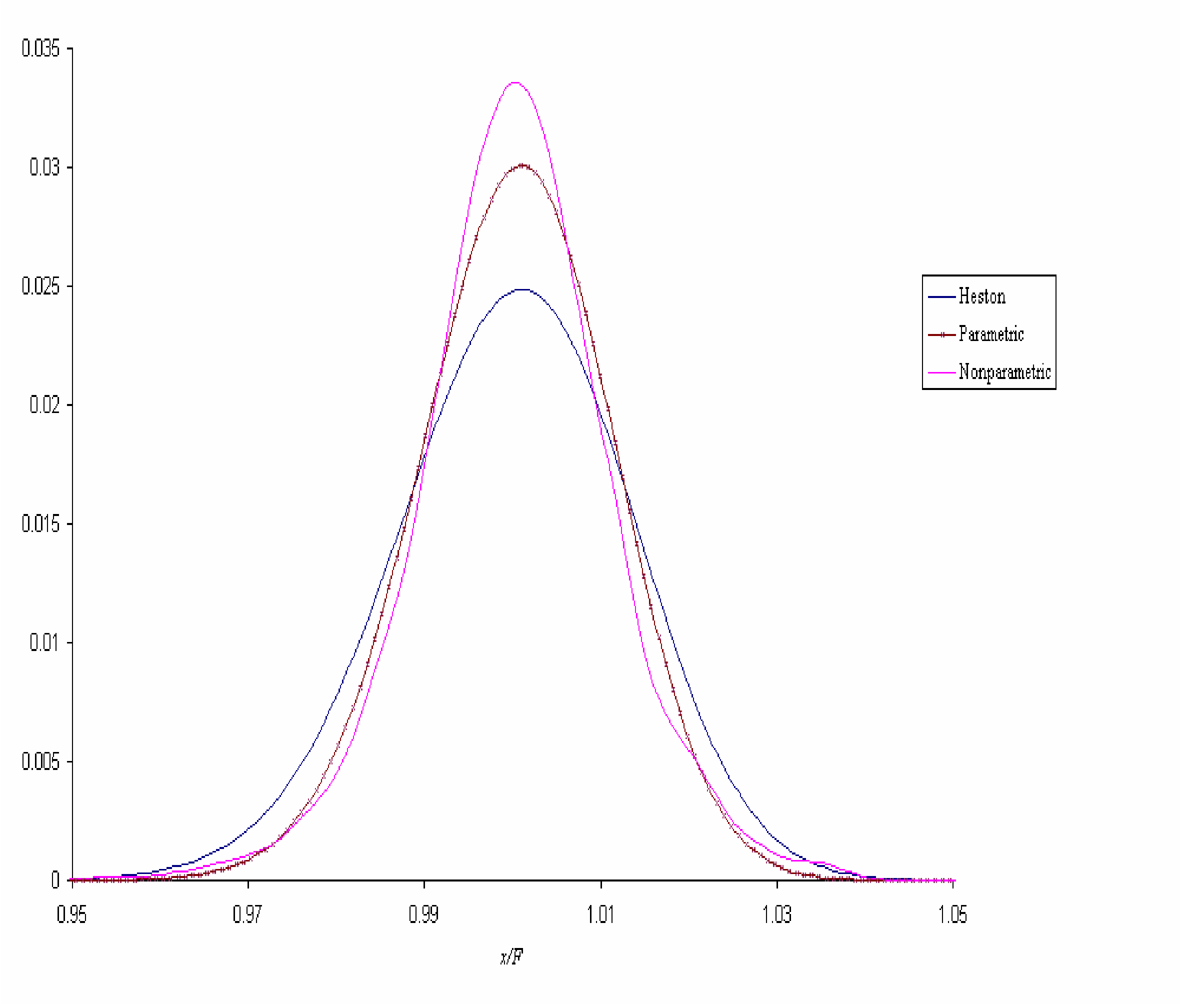

Mixture densities

The mixture density is:

$$g_{mix,T}(x) = \alpha\, g_{P,T}(x) + (1-\alpha) g_{ARCH,T}(x), \qquad 0 \le \alpha \le 1.$$

- If the “best” value of α is between 0 and 1, then incremental information is provided by both the ARCH and the option-based densities.

Econometric methods

Parameter estimation

- Always ex ante. Parameters are re-estimated for each forecast.
- Each day, (1) the volatility for the lognormal *Q*-densities is simply the at-the-money, nearest-to-expiry implied volatility.
- (2) least squares provides the five parameters of the Heston *Q*-dynamics:
  - by matching option market prices with theoretical prices,
  - using all option expiry dates and all strikes.

- Each day, the maximum likelihood method provides estimates of:
  - ARCH parameters,
  - The parametric risk-transformation parameters,
  - The mixture parameter α,
  - Separately for each forecast horizon.

- The ML method is applied to all available futures prices; one price is observed at the <u>same</u> time-interval as the forecast horizon.
- Thus the density forecasts do not overlap.

Criteria for evaluating density forecasts

- For a selected, constant forecast horizon….
- Method *m* provides a series of density forecasts $g_{m,t}(x)$,
- Made at times $V \le t \le W$,
- Evaluated at times from $V+1$ to $W+1$ inclusive,
- The out-of-sample, log-likelihood is used to rank methods:

$$L_m = \sum_{t=V}^{W} \log(g_{m,t}(F_{t+1})).$$

Diagnostic tests are applied to observed cumulative probabilities *u*.

- For a general method *m* these are defined by

$$u_{t+1} = \int_{0}^{F_{t+1}} g_{m,t}(x)dx, \quad V \leq t \leq W.$$

- Are observations from $U[0,1]$, for correctly specified densities…
- … and are then i.i.d. for non-overlapping densities.
- We evaluate:
  - The Kolmogorov-Smirnov test,
  - The LR3 test of Berkowitz (2001).

Data

- Forecasts are made for the prices of S&P 500 index futures:
  - Mar, Jun, Sep, Dec contracts.
- Futures prices are used from 1982 to 2004:
  - Daily settlement prices,
  - Intraday prices, recorded every 5 minutes.

- Prices of American options on futures are used from 1990 to 2004:
    - Daily settlement prices.
    - All out-of-the-money call and put options.
    - European prices calculated by subtracting an estimate of the early exercise premium; the premium is often negligible.
- The daily panels of option prices contain, on average:
    - Prices for 3.1 distinct expiry dates,
    - 37 prices for each expiry date.
- Figure 1 shows typical implied volatilities.

Figure 1

Implied volatilities for out-of-the-money options on November 7, 1997

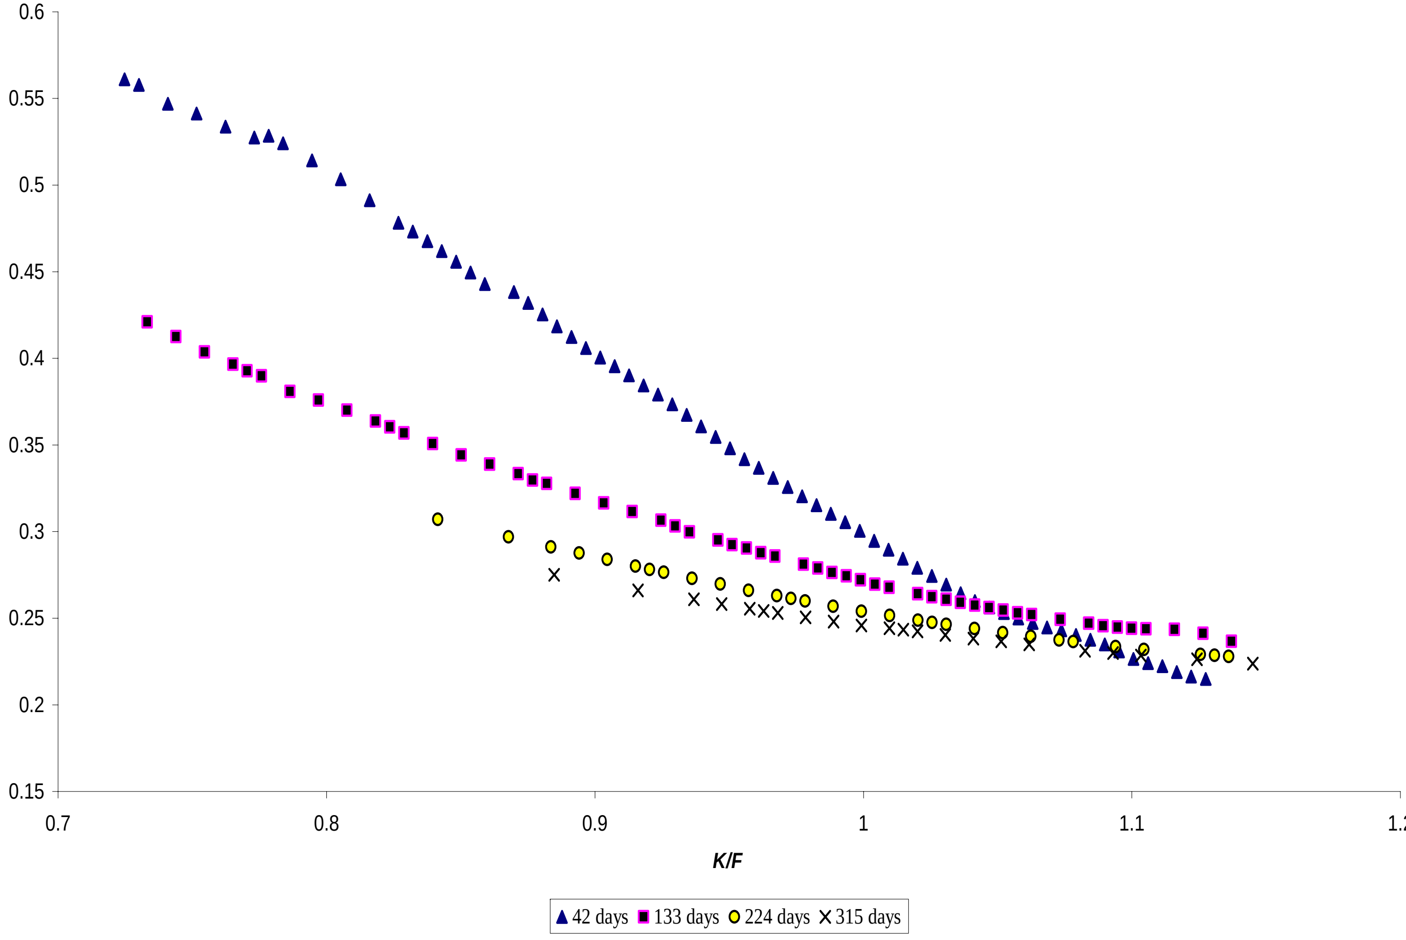

The empirical results

Density forecasts are evaluated:

- From 1991 to 2004.
- For 7 horizons:
  - One trading day,
  - 1, 2, 4, 6, 8 and 12 weeks.
- For non-overlapping forecasts.
- With all parameters estimated <u>ex ante</u>.

Historical densities

- There are no surprises.
- The usual likelihood comparisons show:
  - Intraday returns are more informative than daily returns.
  - A fat-tailed conditional density is required.
  - … for all horizons.
- Consequently, we concentrate on the “Intra-t” specification.

Risk-neutral densities

The median values of the Heston parameters are:

- $\theta = 0.0452$; the annualized volatility then reverts towards 21.3%.
- $\kappa = 4.15$; the "half-life" of the variance process is then 2 months.
- $\xi = 0.79$; this "vol. of vol." term controls the kurtosis of returns and it is high compared with estimates obtained from returns.
- $\rho = -0.66$; which is a consequence of the skewed implied volatility curves.

The empirical Q-densities are not satisfactory P-densities.

- The Q-variance is higher than the real-world variance.
- So we see too few observations in the tails of the Q-densities.
- The “risk transformations” are therefore necessary.

- Define a time series by $u_{t+1} = G_{Q,t,1}(F_{t+1})$.
- Their sample c.d.f. is denoted by $\tilde{C}(u)$.
- The difference between sample and uniform probabilities, $\tilde{C}(u) - u$, is shown on Figure 4 for the Heston Q-densities.

Figure 4

The function $\widetilde{C}(u)-u$ for one-day forecasts obtained from Heston's model and risk-transformations

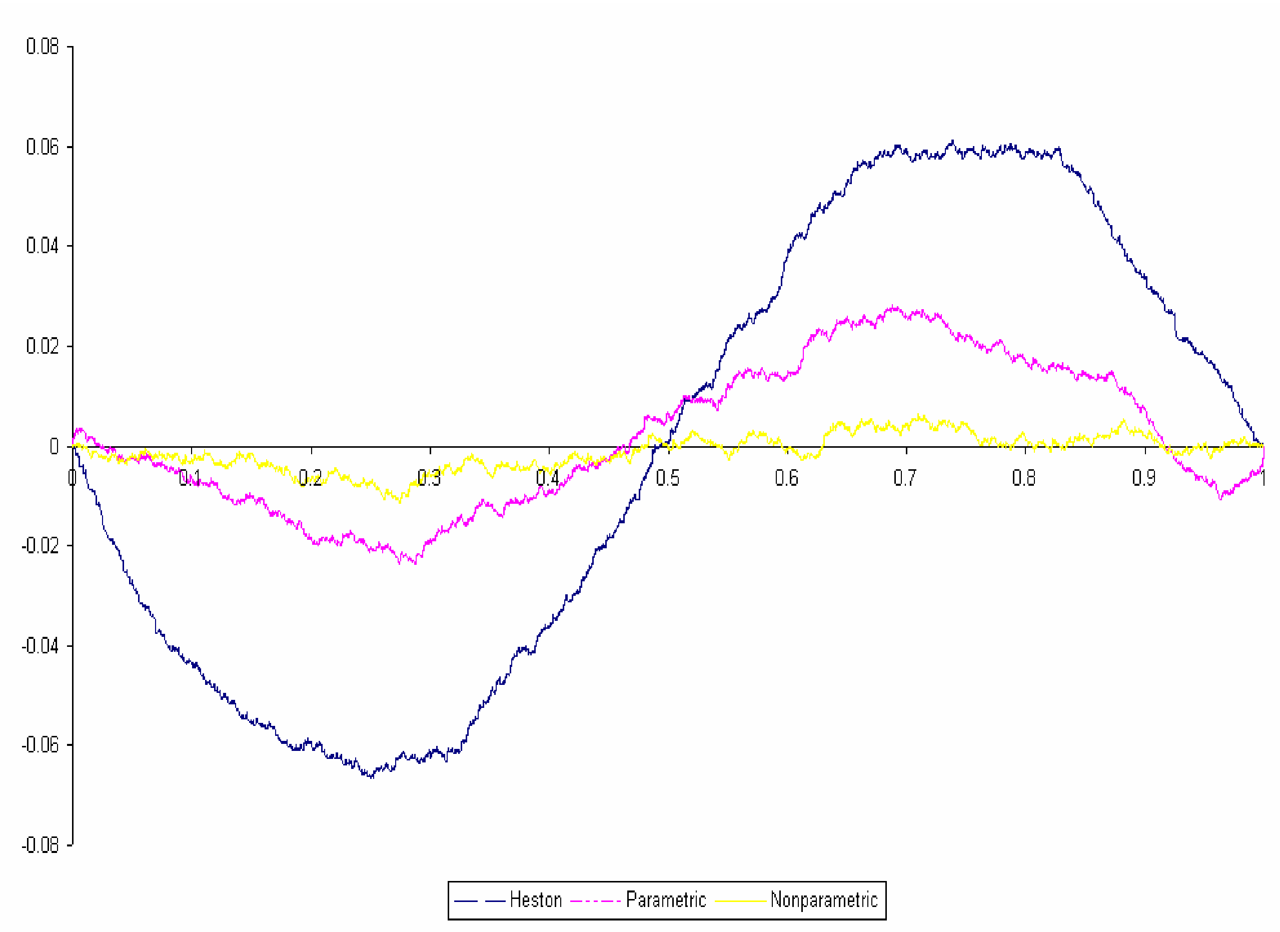

Likelihood comparisons

- Table 4 shows out-of-sample, log-likelihoods, above the GJR level.
- We rank methods using these numbers, and focus on:
  - The historical, Intra-t densities.
  - The lognormal and Heston *Q*-densities.
  - The parametric calibration transformation: *P1*-densities.
  - The non-parametric calibration transformation: *P2*-densities.
  - The drift-correction transformation: *P3*-densities.
- Our conclusions depend upon the forecast horizon.

	Intra-t	Lognormal			Heston			
Horizon		*Q*	*P1*	*P2*	*Q*	*P1*	*P2*	*P3*
1 day	123	27	73	101	-2	104	127	94
1 week	34	17	33	36	18	41	35	41
2 w.	19	14	28	26	15	27	22	26
4 w.	12	12	13	16	16	20	20	23
6 w.	16	16	17	17	20	16	19	20
8 w.	5	5	9	9	7	8	7	7
12 w.	8	6	7	7	10	8	9	8

One-day horizon: methods ranked by log-likelihood

Heston, *P2*	127	
Intra-t	123	…. Historical data scores highly
Heston, *P1*	104	
Lognormal, *P2*	101	
Heston, *P3*	94	
Lognormal, *P1*	73	
Lognormal, *Q*	27	…. *P*-densities are much better
Heston, *Q*	-2	than *Q*-densities

- *P2* is the best transformation.

Horizons of 1, 2 and 4 weeks

- The 5 P-density methods outrank the Intra-t densities (14 cases out of 15).
- The 5 P-density methods have similar log-likelihoods.
- Each P-density method outranks the corresponding Q-method.

Horizons of 6, 8 and 12 weeks

- All 8 methods provide similar log-likelihoods.
- Investigate overlapping densities to make more incisive comparisons.

Mixture densities

Consider: α * option-based density + $(1-\alpha)$ * historical density.

Selected results, for the *P1* transform of the Heston densities:

- One-day horizon:
    - Average α is 0.42.
    - Reject $\alpha = 0$ and $\alpha = 1$ at very low significance levels.
- One-week horizon: Average 57%. Reject special cases at the 5% level.
- Longer horizons: the benefits from mixing are less obvious.

Diagnostic test failures

The *P2*-densities fail the least KS and LR3 tests. At the 5% level:

	Intra-t	Lognormal			Heston			
Horizon		*Q*	*P1*	*P2*	*Q*	*P1*	*P2*	*P3*
1 day		K L	K		K L	K		K
1 week	K L	K L	K L	L	K L		L	
2 weeks		K L			K L			
4 weeks	K	K L	K					

Concluding remarks

- Informative, real-world, densities can be calculated from option prices for a general forecast horizon.
- The Heston *P2*-densities outperform historical densities:
  - when out-of-sample, likelihoods are compared,
  - when diagnostic test results are compared,
  - …. even when intraday returns are used.
- Mixture densities provide the best results for short horizons.