Probabilistic Forecasts of Volatility and its Risk Premia

Worapree Maneesoonthorn, Gael M. Martin*,
Catherine S. Forbes and Simone Grose

Department of Econometrics and Business Statistics, Monash University

February 21, 2012

Abstract

The object of this paper is to produce distributional forecasts of asset price volatility and its associated risk premia using a non-linear state space approach. Option and spot market information on the latent variance process is captured by using dual 'model-free' variance measures to define a bivariate observation equation in the state space model. The premium for variance diffusive risk is defined as linear in the latent variance (in the usual fashion) whilst the premium for variance jump risk is specified as a conditionally deterministic dynamic process, driven by a function of past measurements. The inferential approach adopted is Bayesian, implemented via a Markov chain Monte Carlo algorithm that caters for the multiple sources of non-linearity in the model and for the bivariate measure. The method is applied to spot and option price data on the S&P500 index from 1999 to 2008, with conclusions drawn about investors' required compensation for variance risk during the recent financial turmoil. The accuracy of the probabilistic forecasts of the observable variance measures is demonstrated, and compared with that of forecasts yielded by alternative methods. To illustrate the benefits of the approach, it is used to produce forecasts of prices of derivatives on volatility itself. In addition, the posterior distribution is augmented by information on daily returns to produce value at risk predictions. Linking the variance risk premia to the risk aversion parameter in a representative agent model, probabilistic forecasts of (approximate) relative risk aversion are also produced.

*Corresponding author. Email: gael.martin@monash.edu. This research has been supported by Australian Research Council (ARC) Discovery Grant DP0985234 and ARC Future Fellowship FT0991045. We would like to thank the Editor and two anonymous referees for very detailed, thoughtful and constructive comments on the original draft. We are also grateful to Herman van Dijk and John Geweke, plus various other participants at the: Workshop on Bayesian Econometric Methods (Rimini, 2009); Frontiers in Financial Econometrics Workshop (Princeton, 2009); Econometric Society Australasian Meeting (Canberra, 2009); Bayes on the Beach Workshop (Queensland, 2010); Finance Integrity Research Network Doctoral Tutorial (Melbourne, 2010); and European Seminar on Bayesian Econometrics (Rotterdam, 2010), for very helpful comments on earlier versions of the paper.

1 Introduction

Volatility plays a central role in financial applications, with accurate forecasts of the volatility of future returns being critical for asset pricing, portfolio management and the calculation of risk measures. The aim of this paper is to exploit information from both the spot and option markets to produce full *distributional* forecasts of both this key quantity itself, and any premia that volatility commands in the option market. Using a non-linear state space framework, both the objective volatility of the underlying and the volatility risk premia factored into option prices, are modelled as latent state processes. Applying Bayesian simulation techniques, the predictive distributions of all latent and observable quantities are estimated, and the accuracy thereof compared with that of predictions produced via alternative approaches.

Unlike previous work, in which option prices are linked to state variables via complex pricing formulae, and spot price information enters via daily returns (e.g. Eraker, 2004, Forbes, Martin and Wright, 2007, Johannes, Polson and Stroud, 2009), we use direct non-parametric, or 'model-free' measures of volatility itself, constructed from intraday data, to define a bivariate observation equation in a state space model. An explicit parametric model is then specified for the latent variance, including its two associated risk premia: one that compensates for small and regular movements in the variance (the variance diffusive risk premium), and the other compensating for rare jumps (the variance jump risk premium). The diffusive risk premium is parameterized in the conventional way, as proportional to the latent variance itself. The variance jump risk premium is also allowed to be time varying with a conditionally deterministic process, driven in part by the past 'observed' jump risk premium, used to capture the dynamic behaviour of this component of the model. The avoidance of *price* measures in the state space model means that explicit modelling of the asset price process itself, including the premium for price jump risk, is not necessary. Hence, the volatility forecasts are robust to any reasonable assumptions about the price process.

The focus on probabilistic forecasting of volatility is in marked contrast to the focus of related work that uses both spot- and option-based volatility measures in inferential settings, whereby either *point* forecasts of volatility are the key output (e.g. Blair, Poon and Taylor, 2001, Martens and Zein, 2004, Pong, Shackleton, Taylor and Xu, 2004, Koopman, Jungbacker and Hol, 2005, Martin, Reidy and Wright, 2009, Busch, Christensen and Nielsen, 2011), or forecasting volatility *per se* is not the primary goal of the analysis (e.g. Bollerslev, Gibson and Zhou, 2011, Garcia, Lewis, Pastorello and Renault, 2011, Wu, 2011). Moreover, the emphasis on probabilistic forecasting - whilst, of course, inherent to a Bayesian approach - is also consistent with more general developments in the recent forecasting literature, in which distributional forecasts are increasingly viewed as the principal object of interest, encapsulating as they do all relevant information about the likely future values of the variable of interest. (See Corradi and Swanson, 2006, Gneiting, Balabdaoui and Raftery, 2007, Gneiting, 2008, Geweke and Amisano, 2010, and McCabe, Martin and Harris, 2011, for relevant expositions.) The modelling and prediction of time-varying volatility premia is similar in spirit to the analyses in Brandt and Wang (2003), Bollerslev, Gibson and Zhou (2011) and Todorov (2010), amongst others, but with the probabilistic representation of risk premia predictions being novel.

The Bayesian predictive distributions are produced via a Markov chain Monte Carlo (MCMC) algorithm that caters for multiple sources of non-linearity in the state space model and that allows for multi-move sampling of the latent variances. The conditionally deterministic specification for the variance jump risk premium is computationally convenient, with the posterior distribution of the risk premium at any time point - including future time points - able to be estimated from the MCMC draws of the parameters and latent variables to which the premium is functionally related.

The method is applied to empirical spot and option price data for the S&P500 index over the 1999 to 2008 period, with the spot-based measure constructed from high frequency observations on the index, and the option-implied volatility measure based on the VIX volatility index published by the Chicago Board of Option Exchange (CBOE). One-day-ahead distributional forecasts are produced for all latent and observable quantities of interest, for an evaluation period including both the lead up to the recent global financial crisis and the peak of the crisis at the end of 2008. Most notably, the extraction of forecasts for the variance risk premia enables a picture to be constructed of the extent to which investors' expectations of future risks - and, correspondingly, their demand for compensation - is affected by extreme daily movements in the market. The accuracy of probabilistic forecasts of the (observable) measures of variance is compared with that of forecasts produced by univariate observation-driven time series models for these quantities, as well as forecasts generated by state space models nested in our general structure.

Illustrations of how the distributional forecasts can be used in financial applications are provided. Firstly, we demonstrate how the predictions of the latent variance may be used to produce probabilistic forecasts of the prices of futures written on volatility itself. Secondly, the model is augmented by observations on daily returns to produce probabilistic forecasts of returns, from which one-day-ahead value at risk (VaR) predictions are extracted. Finally, coupled with a particular form of representative agent model, forecasts of the variance risk premia are transformed into probabilistic forecasts of the (approximate) relative risk aversion of the representative investor.

The remainder of the paper is organized as follows. Section 2 describes the measurement of latent volatility using high frequency spot and option price data. Section 3 then outlines a model for latent volatility, including the specification of the dynamic risk premia, completing the state space model. A description of the MCMC algorithm used to sample the latent variables and static parameters, and used subsequently to produce the desired forecast distributions, is provided in Section 4, with all details relevant to the sampling of the latent variances provided in Appendix A. The results of an extensive empirical investigation of intraday spot and option price data for the S&P500 index from July 1999 to December 2008 are reported in Section 5, with supporting material and derivations included in the remaining appendices. Some conclusions are given in Section 6.

2 The Measurement of Latent Volatility

Denoting by p_t the (natural) logarithm of the spot asset price P_t at time t, we assume a continuous time jump diffusion process,

$$dp_t = \mu_t dt + \sqrt{V_t} dB_t^p + dJ_t^p \ , \ t \geq 0, \tag{1}$$

where μ_t is a function of locally finite variation, V_t is a strictly positive variance process, B_t^p is standard Brownian motion, $dJ_t^p = Z_t^p dN_t^p$, and J_t^p is a random jump process that allows for occasional jumps in p_t of random size Z_t^p, with $P(dN_t^p = 1) = \delta dt$ and $P(dN_t^p = 0) = (1 - \delta)\, dt$. The *quadratic variation* for the return over the horizon $t-1$ to t (call this day t), $r_t = p_t - p_{t-1}$, is then given by

$$Q\mathcal{V}_{t-1,t} = \int_{t-1}^{t} V_s ds + \sum_{t-1<s\leq t}^{N_t^p} (Z_s^p)^2 \, . \tag{2}$$

That is, $Q\mathcal{V}_{t-1,t}$ is equal to the sum of the *integrated variation* of the continuous sample path component of p_t,

$$\mathcal{V}_{t-1,t} = \int_{t-1}^{t} V_s ds, \tag{3}$$

and the sum of the $N_t^p - N_{t-1}^p$ squared price jumps that occur on day t. Denoting $r_{t_i} = p_{t_i} - p_{t_{i-1}}$ as the *ith* transaction (logarithmic) return on day t, it is now standard knowledge (Barndorff-Nielsen and Shephard, 2002, 2004, Andersen, Bollerslev, Diebold and Labys 2003) that

$$RV_t = \sum_{t-1<t_i\leq t}^{M} r_{t_i}^2 \xrightarrow{p} Q\mathcal{V}_{t-1,t} \tag{4}$$

and

$$BV_t = \frac{\pi}{2} \sum_{t-1<t_i\leq t}^{M-1} |r_{t_i}| \left|r_{t_{i-1}}\right| \xrightarrow{p} \mathcal{V}_{t-1,t}, \tag{5}$$

where RV_t is referred to as *realized variance*, BV_t as *bipower variation*, and M is equal to the number of intraday returns on day t used in the construction of these measures.[1]

The quantities in (4) and (5) are thus observable measures of the relevant latent variation (or, in a typical abuse of terminology, *volatility*) quantities in a standard continuous time setting for an asset price. Given that both measures are constructed without reference to specific parametric models for the components of (1), they are often referred to as non-parametric, or 'model-free', measures of volatility (see, for example, Andersen, Bollerslev and Diebold, 2010). Increasingly, observation-driven time series models fitted directly to the measures are used to produce volatility predictions; some recent examples being Andersen, Bollerslev and Diebold (2007), Aït-Sahalia and Mancini (2008), Martin, Reidy and Wright (2009), Martens, van Dijk and de Pooter (2009) and Liu and Maheu (2009). Alternatively, when allied with explicit parametric models for $Q\mathcal{V}_{t-1,t}$ (or $\mathcal{V}_{t-1,t}$), the measures are used as observable quantities in a state space framework, with filtering techniques employed to extract estimates of the latent quantities; see Barndorff-Nielsen and Shephard (2002), Creal (2008) and Jacquier and Miller (2010) for illustrations.

We adopt the state space approach in this paper, with BV_t used as a measure of the unobserved integrated variance in (3). An explicit parametric model is specified only for the process driving the continuous sample path variation in p_t, namely the latent point-in-time variance process V_t, $t \geq 0$, with no parametric model adopted at this point for the latent price jump process, J_t^p. With the value of M used in the calculation of BV_t in (5) assuming a finite value in practice, we view BV_t as a *noisy* measure of $\mathcal{V}_{t-1,t}$. Using a rectangular approximation to the integral that defines $\mathcal{V}_{t-1,t}$, we then specify the first measurement equation as

$$BV_t = V_t + u_{BV_t};\ u_{BV_t} \sim N(0, \sigma_{BV}^2 V_t^2). \tag{6}$$

The adoption of a Gaussian measurement error in (6) (conditional on the volatility path) is motivated, in part, by computational convenience. However, related asymptotic theory regarding BV_t provides some informal justification for both this modelling assumption and the use of a state-dependent measurement error variance, with Barndorff-Nielsen and Shephard (2006) demonstrating that under the assumption of no price jumps, as $M \to \infty$, BV_t converges to a normal distribution with a variance that is a function of integrated quarticity, $\mathrm{I}Q_{t-1,t} = \int_{t-1}^{t} V_s^2 ds$. The approximation of the two integral quantities, $\mathcal{V}_{t-1,t}$ and $\mathrm{I}Q_{t-1,t}$, by corresponding time t quantities in (6) is made for computational ease, with the 'augmenta-

[1]Implicit in (4) and (5) is the assumption that microstructure noise effects are absent. The formal incorporation of microstructure noise in the assumed process for intraday returns has led to modifications of RV_t and BV_t that are consistent estimators of $Q\mathcal{V}_{t-1,t}$ and $\mathcal{V}_{t-1,t}$ respectively, in the presence of such noise. See *Journal of Econometrics, 2011, Volume 160(1): Special Issue on Realized Volatility,* for several recent contributions to this literature.

tion' methods of Elerian, Chib and Shephard (2001), Eraker (2001) and Jones (2003) being feasible techniques for managing a finer discretization of the integrals if so desired.

With a view to identifying not only the objective volatility itself, but also the premia that it commands in the option market, we supplement BV_t with a variance measure constructed from the prices of options written on the underlying asset. Specifically, in the context of the general jump diffusion model in (1), Jiang and Tian (2005) and Carr and Wu (2009) (and as based on the earlier work of Britten-Jones and Neuberger, 2000) demonstrate that the risk-neutral conditional expectation, $E^*(.)$, of quadratic variation,

$$E^*(Q\mathcal{V}_{t,t+\tau}|\mathcal{F}_t) = E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t) + E^*\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right], \tag{7}$$

is implied by a continuum (over strike K) of option prices with maturity $\tau > 0$, $C(t+\tau, K)$,

$$MF_t^{QV} = 2\int_0^\infty \frac{C(t+\tau,K) - C(t,K)}{K^2} dK, \tag{8}$$

where $\mathcal{F}_t = \sigma\{V_s; s \leq t\}$ is the sigma-algebra generated by the point-in-time variance process.[2] The notation MF_t^{QV} is used here to denote the measure as a non-parametric, or 'model free' representation of the risk-neutral expectation of quadratic variation over the maturity period, with the expression 'model-free' being consistent with common parlance. Given an approximation of MF_t^{QV}, constructed from a finite set of observed option prices on day t, the option-based measure can be used to extract information on the risk-neutral conditional expectation of $Q\mathcal{V}_{t,t+\tau}$. When used jointly with a spot-price based measure, and allied with a model for the risk premia factored into the conditional expectation, the two measures can be used to conduct inference about *both* the objective process *and* all relevant risk premia. (See Eraker, 2008, and Duan and Yeh, 2010, for other uses of the model-free volatility measure in inferential settings.)

With the aim of avoiding the explicit specification of the price jump process in the state space model, we adjust the option-based measure as follows. Define $\widehat{MF}_t^{QV}$ as a measure of $E^*(Q\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)$ in (7), produced by using a finite set of observed option prices on day t to approximate MF_t^{QV} in (8).[3] Also define $\widehat{E}^*\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right]$ as an estimate of the risk neutral expectation of the jump variation, and

$$MF_t = \widehat{MF}_t^{QV} - \widehat{E}^*\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right] \tag{9}$$

[2] Carr and Wu (2009) demonstrate that $E^*(Q\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)$ is not given *exactly* by MF_t^{QV} in the presence of price jumps, but that the approximation error is insignificant for practical purposes. In any case, any such error will be absorbed by the measurement error, under our approach.

[3] As detailed in Section 5.1, $\widehat{MF}_t^{QV}$ is represented by the (appropriately scaled) VIX of the CBOE, for the empirical analysis of the S&P500 stock index data conducted therein.

as the resultant estimate of $E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)$. In the spirit of Bollerslev, Gibson and Zhou (2011), we assume that the risk neutral expectation of jump variation is some constant proportion ($c \geq 1$) of the associated objective expectation, invoking the *a priori* assumption that the premium factored into option prices for price jump risk is positive. Robustness to different values of c is then documented. We thus specify

$$\widehat{E}^*\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right] = c\widehat{E}\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right] \tag{10}$$

for some constant c, with $\widehat{E}\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right]$ measured as follows. As in Tauchen and Zhou (2011), we identify the presence of a significant price jump on any day t by the realization of the statistic,

$$Z_{Jt} = \frac{RJ_t}{\sqrt{(2.61-2)\, M^{-1} \max\left(1, \frac{TP_t}{BV_t^2}\right)}}, \tag{11}$$

where $RJ_t = \frac{RV_t - BV_t}{RV_t}$, and where tri-power quarticity, TP_t, serves as a consistent estimator of the integrated quarticity, $\mathrm{I}Q_{t-1,t}$. Realized price jump variation on day t is then extracted by means of

$$JV_t = I\left(Z_{Jt} > Z_\alpha\right) \times \max\left(0, RV_t - BV_t\right), \tag{12}$$

with a time series of values for JV_t produced accordingly.[4] As is consistent with the bulk of the empirical literature and with the theoretical assumption of independent jumps over time, JV_t (as based on the empirical data analyzed in Section 5) exhibits little autocorrelation. Hence, a short-memory autoregressive model of order one is used to produce point predictions, conditional on data up to time t, of $JV_{t,t+1}$, $JV_{t,t+2}$, ..., $JV_{t,t+\tau}$, with $\widehat{E}\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right]$ given as the aggregate of these τ predictions.

The second measurement equation is then specified as

$$MF_t = E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t) + u_{MF_t}; \quad u_{MF_t} \sim N(0, \sigma_{MF}^2 V_t^2), \tag{13}$$

where u_{MF_t} captures both the error associated with the discretization and truncation of the integral in (8) and the estimation error in $\widehat{E}^*\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}} (Z_s^p)^2 \middle| \mathcal{F}_t\right]$. As with (6), a (conditionally) Gaussian measurement error is adopted in (13) primarily for convenience, with a state-dependent measurement error variance motivated by the following additional empirical

[4] Tri-power quartricity is computed as $TP_t = M\mu_{4/3}^{-3} \sum_{t_i \in [t-1,t]}^{M} |r_{t_i}|^{4/3} |r_{t_{i-1}}|^{4/3} |r_{t_{i-2}}|^{4/3}$, where $\mu_{4/3} = 2^{2/3}\Gamma\left(\frac{7}{6}\right)\Gamma\left(\frac{1}{2}\right)^{-1}$ and $TP_t \to \mathrm{I}Q_{t-1,t}$ as $M \to \infty$. Under the null hypothesis of no price jumps, Z_{Jt} is asymptotically $N(0,1)$ as $M \to \infty$. Thus, in testing whether a price jump is present on a particular day, the $\alpha-$level critical value from a standard normal distribution (Z_α) applies.

considerations. Firstly, measurement error in $\widehat{MF}_t^{QV}$ is induced (in part) by the truncation of the volatility smile used in its calculation (see Jiang and Tian, 2005). Given the occurrence of more extreme smiles in periods of high levels of volatility, the extent of this measurement error is likely to be dependent on V_t. Secondly, the state dependent measurement error may also capture any possible dependence of the intensity of variance jumps on the level of V_t (e.g. Eraker, 2008), a dependence that (as will be seen in the following section) we do not explicitly model.

3 A Parametric Model for Latent Volatility and its Risk Premia

3.1 Constant risk premium parameters

In addition to the specification of the measurement (or observation) equations in (6) and (13), which link the spot- and option-based volatility measures to the latent volatility process, the production of volatility forecasts within a state space framework requires a dynamic model for the latent objective volatility process itself and its risk premia, with the latter associated with an assumed risk-neutralized volatility process. We begin with an outline of the parametric continuous time models of objective and risk-neutral volatility that give rise to the latent state process, highlighting the inferential issues associated with the unknown quantities of interest, including the risk premium parameters - with the latter assumed constant for the purpose of this initial exposition. The resulting state space model is then given towards the end of the section. In Section 3.2 we allow for a dynamic jump risk premium parameter, with the full state space model augmented by this additional (conditionally deterministic) latent variable.

The objective latent variance V_t in (1) is assumed to evolve according to the following standard jump diffusion model,

$$dV_t = \kappa[\theta - V_t]dt + \sigma_v \sqrt{V_t} dB_t^v + dJ_t^v, \tag{14}$$

where $dJ_t^v = Z_t^v dN_t^v$, $Z_t^v \sim Exp(\mu_v)$, $P(dN_t^v = 1) = \delta_J dt$ and $P(dN_t^v = 0) = (1 - \delta_J)\, dt$. At this point in the development of the model, we require no particular assumptions to be adopted for the relationship between the Brownian and jump increments in (1) and (14), including the equality (or otherwise) of the jump intensities, δ and δ_J. Hence, the framework accommodates a range of models adopted in the literature for stochastic volatility (e.g. Duffie, Pan and Singleton 2000, Eraker, Johannes and Polson, 2003, Eraker, 2004, Broadie, Chernov and Johannes, 2007).

Equilibrium arguments can be used to produce the following risk-neutral process for the variance,

$$dV_t = \kappa^*[\theta^* - V_t]dt + \sigma_v\sqrt{V_t}dB_t^{*v} + dJ_t^{*v}, \tag{15}$$

which, in conjunction with an appropriate risk-neutralized version of the price process in (1), is used to price options on the asset. In common with Eraker (2004) and Broadie, Chernov and Johannes (2007), we assume $dJ_t^{*v} = Z_t^{*v}dN_t^{*v}$ and $Z_t^{*v} \sim Exp(\mu_v^*)$, and in common with Eraker, we impose $\delta_J^* = \delta_J$. That is, we allow only the jump size to change across probability measures, with jump intensity held constant, and assume that investors factor into their option pricing the following premium for variance jump risk,

$$\lambda_J = \mu_v^* - \mu_v. \tag{16}$$

Also implicit in the move from (14) to (15) is the transformation

$$\kappa^*[\theta^* - V_t] = \kappa[\theta - V_t] - \lambda_D V_t,$$

where

$$\kappa^* = \kappa + \lambda_D;\ \theta^* = \frac{\kappa\theta}{\kappa^*} \tag{17}$$

and λ_D is a scalar parameter. The term $\lambda_D V_t$ represents the premium associated with variance diffusive risk, with the value of λ_D determining the magnitude (and sign) of the premium factored into option prices for the risk associated with small and regular changes in the non-traded state variable, V_t.

Empirical estimates of λ_D and λ_J reported in the literature are (not surprisingly) model dependent. In purely diffusive stochastic volatility models (in which no jumps or jump premia are parameterized - in either the price or the variance), estimates of λ_D are typically negative (see, for example, Guo, 1998, Eraker, 2004, Forbes, Martin and Wright, 2007, Bollerslev, Gibson and Zhou, 2011), which implies slower reversion ($\kappa^* < \kappa$) to a higher mean level ($\theta^* > \theta$) under the risk neutral distribution. However, as outlined in some detail by Broadie, Chernov and Johannes (2007), empirical conclusions regarding the significance and sign of λ_D are less clear cut once jumps (and associated premia) are included in the model specification. Overall, current empirical evidence points to λ_J being significantly greater than zero, with the significance of λ_D tending to be reduced accordingly. A positive value for λ_J leads to the same *qualitative* result as a negative value of λ_D. That is, either numerical outcome leads to a higher long run expectation of V_t under the risk-neutral measure than under the objective measure, given that the two long-run means are given respectively by $E^*(V_t) = [\kappa\theta + (\mu_v + \lambda_J)\delta_J] / [\kappa + \lambda_D]$ and $E(V_t) = [\kappa\theta + \mu_v\delta_J] / \kappa$. This implies, in turn, that (call) options are priced higher under the risk-neutral process, on average, than if they had been priced under the objective measure. That is, these signs for the risk premium parameters ($\lambda_D < 0$ and $\lambda_J > 0$ respectively) imply that investors are willing to pay a premium for options, as a hedge against movements in the spot price that result from either the diffusive or jump components of the random variance (or both).

As is clear from (16) and (17), observed option prices can be used to identify the parameters of the objective process, and the risk premia, λ_J and λ_D, only if additional information on the objective parameters and/or λ_J and λ_D, is introduced. Earlier analyses (based on a variety of versions of the model presented here) solved this identification problem: by jointly estimating the objective and risk-neutral processes using option and spot *price* data (e.g. Chernov and Ghysels, 2000, Pan, 2002, Polson and Stroud, 2003, Eraker, 2004, Forbes, Martin and Wright, 2007); by using option price data only to estimate risk-neutral parameters, and extracting estimates of the risk premium parameters via separate return-based estimates of the objective parameters (e.g. Guo, 1998, Broadie, Chernov and Johannes, 2007); or by either imposing theoretical restrictions on the risk premia (Bates, 2000) or using externally estimated values (Johannes, Polson and Stroud, 2009). Most importantly, in all of these studies, the link between observed market option prices and the objective variance in (14) is quite complex, with volatility imputed via the numerical approximation of a theoretical option price formula. The link between the spot prices and the latent variance also occurs 'indirectly' via the variance specification of a measurement equation for the (typically) daily return.

In contrast, we use the high frequency variance measures available from spot and option markets outlined in Section 2 as the basis for our measurement equations in (6) and (13). Given (15), the first term in the R.H.S. of (13) is now explicitly given by

$$E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t) = a^*_\tau V_t + b^*_\tau + E^*\left(\int_t^{t+\tau} Z^{*v}_s dN^{*v}_s \middle| \mathcal{F}_t\right), \tag{18}$$

where

$$a^*_\tau = \frac{1}{\kappa^*}(1 - e^{-\tau\kappa^*}) \quad \text{and} \quad b^*_\tau = \tau\theta^* - \frac{\theta^*}{\kappa^*}(1 - e^{-\tau\kappa^*}), \tag{19}$$

and where θ^* and κ^* are defined in (17). Invoking at this point the assumptions of independence of dJ^{*v}_t over time, and contemporaneous independence between Z^{*v}_t and dN^{*v}_t, it follows that

$$E^*\left(\int_t^{t+\tau} Z^{*v}_s dN^{*v}_s \middle| \mathcal{F}_t\right) = E^*\left(\int_t^{t+\tau} Z^{*v}_s dN^{*v}_s\right) = \int_t^{t+\tau} \mu^*_v \delta_J ds = \tau\left[\mu_v + \lambda_J\right]\delta_J, \tag{20}$$

where the last equality follows from (16). Using an Euler discretization of (14) over the daily time interval, and with ΔN^v_t also linked to a daily interval, the state space model is thus given by

$$
\begin{aligned}
BV_t &= V_t + \sigma_{BV} V_t \xi_{1t} && (21) \\
MF_t &= a_\tau^* V_t + b_\tau^* + \tau \left[\mu_v + \lambda_J\right] \delta_J + \sigma_{MF} V_t \xi_{2t} && (22) \\
V_t &= \kappa\theta\Delta t + (1 - \kappa\Delta t)\, V_{t-\Delta t} + \sigma_v \sqrt{\Delta t}\sqrt{V_{t-\Delta t}}\xi_{3t} + Z_t^v \Delta N_t^v, && (23)
\end{aligned}
$$

with independent sequences,

$$
\begin{aligned}
\xi_t &= (\xi_{1t}, \xi_{2t}, \xi_{3t})' \overset{iid}{\sim} N(0_3, I_3) && (24) \\
Z_t^v &\overset{iid}{\sim} Exp(\mu_v) && (25) \\
\Delta N_t^v &\overset{iid}{\sim} Bernoulli(\delta_J \Delta t), && (26)
\end{aligned}
$$

for all $t = 1, 2, ..., T$. Setting $\Delta t = 1$ in (23), the state equation for V_t describes the evolution of the point-in-time (annualized) variance from one day to the next, with BV_t being a noisy measure of that quantity. The parameter θ is estimated as an annualized quantity, matching the annualized magnitude of V_t, whilst the parameter κ is treated as a daily quantity, measuring the rate of mean reversion in the annualized V_t per day. In accordance with this treatment of κ, $\tau = 22$ (trading) days and MF_t in (22) is an aggregated (annualized) variance over the trading month.

3.2 A dynamic model for the jump risk premium

In the context of our state space framework, it is straightforward to extend the model to cater for dynamic behaviour in the risk premium parameters, λ_D and λ_J, and produce probabilistic forecasts thereof via the Bayesian approach. Our motivation for investigating such an extension is essentially empirical, with evidence presented in Section 5 justifying the use of a dynamic model for λ_J (specifically) for the S&P500 index data analysed therein. An extension of this kind is also consistent with recent explorations of time varying volatility premia in the literature, most notably in Bollerslev, Tauchen and Zhou (2009), Bollerslev, Sizova and Tauchen (2012), Todorov (2010) and Bollerslev, Gibson and Zhou (2011). This literature, in turn, draws some motivation from related work in which time-variation in investor risk aversion, explored in a variety of settings, is deemed to be plausible; see, for example, Brandt and Wang (2003), Gordon and St-Amour (2004) and Bekaert, Engstrom and Xing (2009).

With our focus being on modelling the dynamic behaviour of daily volatility, and on short-term (one-day-ahead) forecasting, it is beholden upon us to propose a dynamic specification for the risk premia that is able to be implemented on a daily basis, without the need to access values of plausible predictors. Based on the assumption that significant *daily* variation in

λ_J, most notably in response to recent *observed* jumps in volatility, is more plausible, from a behavioural point of view, than daily variation in λ_D, we hypothesize that λ_J in (22) follows a dynamic process, whilst λ_D is held constant. (See also Todorov, 2010).[5] Assuming that short memory dynamics drive λ_J, (an assumption that is also supported empirically in Section 5 and which is consistent with some of the arguments presented in the literature cited in the paragraph above) we specify the following conditionally deterministic model for the value of λ_J at time t,

$$\lambda_{Jt} = \lambda_{J0}\left(1-\alpha_1-\alpha_2\right)+\alpha_1\lambda_{Jt-1}+\alpha_2 l_{Jt-1}, \tag{27}$$

where l_{Jt-1} denotes an 'observed' value of λ_{Jt} at time $t-1$. In specifying l_{Jt-1}, we exploit recent theoretical developments in Carr and Wu (2009) (amongst others) to link the difference between the two observed measures of variance over the option maturity period to λ_{Jt}. In brief, using the concept of a variance swap, Carr and Wu quantify the variance risk premium as the difference between the objective and risk-neutral expectations of quadratic variation over the life of a $\tau-$maturity option. With the focus of our paper being on integrated variance specifically, the conditional variance risk premium defined over the maturity period τ is given by the following linear function of the point-in-time latent variance V_t,

$$\begin{aligned} E(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)-\left(MF_t^{QV}-E^*\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}}(Z_s^p)^2\middle|\mathcal{F}_t\right]\right) &= E(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)-E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)\\ &= a_\tau V_t+b_\tau-\left(a_\tau^* V_t+b_\tau^*+\tau\lambda_J\delta_J\right), \end{aligned} \tag{28}$$

where

$$a_\tau=\frac{1}{\kappa}(1-e^{-\tau\kappa}) \quad \text{and} \quad b_\tau=\tau\theta-\frac{\theta}{\kappa}(1-e^{-\tau\kappa}), \tag{29}$$

and a_τ^* and b_τ^* are as given in (19). The conditional variance risk premium in (28) is also seen to be a linear function of λ_J, the premium for variance jump risk, and via (17) is a non-linear function of λ_D, the diffusive risk premium parameter.[6]

---

[5]Preliminary analysis also indicated that maintaining a constant value for λ_D was justified for the data set in question. Note, however, that even with a constant value for λ_D, the diffusive risk *premium* over dt, $\lambda_D V_t dt$, is still a dynamic process, via the assumed linear relationship with V_t.

[6]Note that the *unconditional* variance risk premium is given by:

$$\underbrace{\frac{\theta\lambda_D}{\kappa^*}\left[\tau-\frac{1}{\kappa^*}\left(1-e^{-\tau\kappa^*}\right)\right]}_{\text{Negative if }\lambda_D<0}-\underbrace{\frac{\delta_J}{\kappa}\left[\tau\kappa\lambda_J+\left(a_\tau^*-a_\tau\right)\mu_v\right]}_{\text{Positive if }\lambda_J>0\text{ and }\lambda_D<0\ (\Rightarrow a_\tau^*>a_\tau)}.$$

Hence, the unconditional mean of the aggregate variance risk premium (over τ) defined in this way is negative if the individual risk premium parameters have the anticipated signs. This result corresponds correctly to the spot price-based quantity, $\mathcal{V}_{t,t+\tau}$, being less, on average, than the option-price-based quantity, $MF_t^{QV}-E^*\left[\sum_{t<s\leq t+\tau}^{N_{t+\tau}}(Z_s^p)^2\middle|\mathcal{F}_t\right]$.

Motivated by (28), we set

$$\widehat{E}(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t) - MF_t = a_\tau V_t + b_\tau - [a_\tau^* V_t + b_\tau^* + \tau\lambda_{Jt}\delta_J],$$

with MF_t as given in (9), and solve for an 'observed' value of λ_{Jt} as

$$l_{Jt} = \frac{[(a_\tau V_t + b_\tau) - (a_\tau^* V_t + b_\tau^*)] - \left[\widehat{E}\left(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t\right) - MF_t\right]}{\tau\delta_J}, \tag{30}$$

at each point t, within the estimation algorithm. As a model-based estimate of the objective conditional expectation, $E(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)$, is needed to evaluate the right-hand-side of (30) at each t, we use the following linear function of BV_t,

$$\widehat{E}(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t) = \frac{a_\tau}{a_1} BV_t + \left((b_\tau + \tau\mu_v\delta_J) - \frac{a_\tau\left(b_1 + \mu_v\delta_J\right)}{a_1}\right), \tag{31}$$

where a_1 and b_1 are simply those functions a_τ and b_τ in (29), evaluated at $\tau = 1$, respectively. Simple calculations can be used to demonstrate that the estimate in (31) is unbiased for $E(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)$.

We use the expression 'observed' in referencing the quantity l_{Jt} in (30) since it is defined by replacing the relevant conditional expectations on the L.H.S. of (28) with *observed* measures, BV_t and MF_t. Hence, for given values of all unknowns, l_{Jt} can be viewed as the value of λ_{Jt} that is implied by the data. Our language is also prompted by the fact that the model in (27) mimics the structure of the generalized autoregressive heteroscedastic (GARCH) model for volatility of Bollerslev (1986), in which squared lagged returns are used as an 'observed' proxy of lagged volatility. (See also the model for trade durations in Engle and Russell, 1998). The fact that l_{Jt} is a *deterministic* function of unknown latent and static parameters, as is λ_{Jt} itself, renders this extension to the model easily manageable from a computational point of view. The absence of a random error term in (27) also avoids the (potential) need to price an additional random risk factor.

The full state space model, incorporating the dynamic jump risk premium, is obtained from (21) to (26) and with λ_J in (22) replaced by λ_{Jt} in (27). This full state space model obviously nests one in which the jump premium is held constant over time, with the predictive performance of the simpler model assessed in the empirical section.

4 MCMC Algorithm and Priors

4.1 Prior specification

To undertake a Bayesian analysis, prior probability distributions for the parameters of the model detailed in (21) to (27) are required. We begin our specification by noting that the variance jumps, which occur with probability δ_J on any day t, shift the intercept term in (23) by an amount Z_t^v. While the specification of the state space model requires neither explicit modelling of the price jumps, nor any assumption to be made about the relationship between variance and price jumps, we are able to exploit the Bayesian approach by quantifying prior beliefs about variance jumps - including their possible relationship with price jumps - via the prior distribution. For a given sample period (details of which are given in Section 5) we identify the presence of significant *price* jumps on any day t by the realizations of the statistic in (11). On the assumption that the intensity of price jumps provides some information about the intensity of the variance jumps, the prior for δ_J is then specified as a beta distribution with mean equal to the proportion of days throughout the sample on which significant price jumps are found to occur. The variance of this distribution is used as a tuning parameter in the algorithm.

In a similar fashion, the prior for μ_v is specified as inverse gamma, with a mean value proportionate to the average magnitude of

$$JV_t \cdot I\left(r_t < 0\right) \tag{32}$$

over the sample of days on which significant large price falls are in evidence (indicated by $I\left(r_t < 0\right)$, where r_t is the daily logarithmic return), with JV_t as defined in (12). That is, *a priori*, we assume that the magnitude of the average jump in variance is some proportion of the average magnitude of the square of *negative* price jumps, reflecting the prior belief of a negative correlation between the price and variance jumps (Eraker, Johannes and Polson, 2003, and Eraker, 2004). Once again, both the (*a priori*) proportionate relationship between μ_v and the average of the sample quantity in (32), plus the variance of the prior distribution for μ_v, are viewed as tuning parameters in the algorithm.

Priors for the remaining parameters are all non-informative, subject to appropriate constraints on the relevant parameter space in some cases, with further details provided in the discussion of the algorithm in the following section. The results of prior (and posterior) predictive analyses (Geweke, 2005), used to assess, in particular, the validity of the prior specifications for δ_J and μ_v, are documented in Appendix B.

4.2 Outline of the Gibbs-based algorithm

Given the complexity of the proposed state space model, the joint posterior distribution for all unknowns is analytically intractable. Hence, an MCMC algorithm is applied to produce draws from the joint posterior and those draws then used to estimate various quantities of interest, including predictive distributions, in the usual way. To reduce notation, we define the vectors $V_{1:t} = (V_1, V_2, \ldots, V_t)'$, $Z_{1:t}^v = (Z_1^v, Z_2^v, ..., Z_t^v)'$, $\Delta N_{1:t}^v = (\Delta N_1^v, \Delta N_2^v, ..., \Delta N_t^v)'$,

$BV_{1:t} = (BV_1, BV_2, \ldots, BV_t)'$ and $MF_{1:t} = (MF_1, MF_2, \ldots, MF_t)'$, for $t = 1, 2, ..., T$, and with $MF_{1:0}$ and $BV_{1:0}$ empty. Using this notation, the joint posterior density for all unknowns is given by $p\left(V_{1:T}, Z^v_{1:T}, \Delta N^v_{1:T}, \phi | MF_{1:T}, BV_{1:T}\right)$, where the vector of static parameters is given by $\phi = (\kappa, \theta, \sigma_{BV}, \sigma_{MF}, \sigma_v, \lambda_D, \mu_v, \delta_J, \lambda_{J0}, \alpha_1, \alpha_2)$. The joint posterior density satisfies

$$\begin{aligned}
& p\left(V_{1:T}, Z^v_{1:T}, \Delta N^v_{1:T}, \phi | MF_{1:T}, BV_{1:T}\right) \\
& \propto \left[\prod_{t=1}^{T} p\left(MF_t | MF_{1:t-1}, BV_{1:t-1}, V_t, \kappa, \theta, \sigma_{MF}, \lambda_D, \mu_v, \delta_J, \lambda_{J0}, \alpha_1, \alpha_2\right)\right. \\
& \left.\times\, p\left(BV_t | V_t, \sigma_{BV}\right) p\left(V_t | V_{t-1}, Z^v_t, \Delta N^v_t, \kappa, \theta, \sigma_v\right) p\left(Z^v_t | \Delta N^v_t, \mu_v\right) p\left(\Delta N^v_t | \delta_J\right)\right] p\left(\phi\right), \qquad (33)
\end{aligned}$$

where it is assumed that $V_0 = \theta + \frac{\mu_v \delta_J}{\kappa}$. The conditioning of MF_t on lagged values of $MF_{1:t-1}$ and $BV_{1:t-1}$ derives from the assumed structure for λ_{Jt} in (27).

The Gibbs-based MCMC algorithm is implemented in four main steps:

1. Generating $V_{1:T}$ from

$$\begin{aligned}
& p\left(V_{1:T} | Z^v_{1:T}, \Delta N^v_{1:T}, \phi, MF_{1:T}, BV_{1:T}\right) \\
& \propto \prod_{t=1}^{T} p\left(MF_t | MF_{1:t-1}, BV_{1:t-1}, V_t, \kappa, \theta, \sigma_{MF}, \lambda_D, \mu_v, \delta_J, \lambda_{J0}, \alpha_1, \alpha_2\right) \\
& \times\, p\left(BV_t | V_t, \sigma_{BV}\right) p\left(V_t | V_{t-1}, Z^v_t, \Delta N^v_t, \kappa, \theta, \sigma_v\right);
\end{aligned}$$

2. Generating $Z^v_{1:T}$ (T truncated normal random variables) from

$$p\left(Z^v_{1:T} | V_{1:T}, \Delta N^v_{1:T}, \phi, MF_{1:T}, BV_{1:T}\right) \propto \prod_{t=1}^{T} p\left(V_t | V_{t-1}, Z^v_t, \Delta N^v_t, \kappa, \theta, \sigma_v\right) p\left(Z^v_t | \Delta N^v_t, \mu_v\right),$$

3. Generating $\Delta N^v_{1:T}$ (T Bernoulli random variables) from

$$p\left(\Delta N^v_{1:T} | V_{1:T}, Z^v_{1:T}, \phi, MF_{1:T}, BV_{1:T}\right) \propto \prod_{t=1}^{T} p\left(V_t | V_{t-1}, Z^v_t, \Delta N^v_t, \kappa, \theta, \sigma_v\right) p\left(\Delta N^v_t | \delta_J\right),$$

4. Generating elements of ϕ from

$$\begin{aligned}
& p\left(\phi | V_{1:T}, Z^v_{1:T}, \Delta N^v_{1:T}, MF_{1:T}, BV_{1:T}\right) \\
\propto\ & \left[\prod_{t=1}^{T} p\left(MF_t | MF_{1:t-1}, BV_{1:t-1}, V_t, \kappa, \theta, \sigma_{MF}, \lambda_D, \mu_v, \delta_J, \lambda_{J0}, \alpha_1, \alpha_2\right)\right. \\
& \left.\times p\left(BV_t | V_t, \sigma_{BV}\right) p\left(V_t | V_{t-1}, Z^v_t, \Delta N^v_t, \kappa, \theta, \sigma_v\right)\right] p\left(\phi\right).
\end{aligned}$$

Obtaining draws of $V_{1:T}$ is the most challenging aspect of the simulation scheme, due to the presence of the state dependent errors in (21), (22) and (23). Whilst the algorithm of Stroud, Müller and Polson (2003) is used as the basis of our approach, we are not aware of this algorithm having been applied to a model with multiple sources of state dependence and a bivariate measure. Details of the mixture-based algorithm used to draw $V_{1:T}$ are provided in Appendix A.

Generation of the elements of ϕ occurs via embedded Metropolis-Hastings (MH) chains where necessary. Under the usual non-informative priors for the standard deviation parameters, σ_{BV} and σ_{MF}, simulation of these parameters is standard, via inverted gamma conditionals. The volatility of volatility parameter σ_v, is produced analogously, with the restriction $\sigma_v^2 < 2\kappa\theta$ ensuring the positivity of the latent variance process. The joint prior for κ, θ and λ_D is uniform, subject to $\kappa\theta > \sigma_v^2/2$ and $\lambda_D < 0$, with the associated univariate conditional posteriors being non-standard due to the fact that the conditional mean function in (22) is non-linear in all three parameters. We use the structure of the model to define Gaussian kernels and produce candidate draws for each of κ, θ and λ_D with separate MH sub-steps. Similarly, the parameters of the jump premium process are uniform *a priori*, subject to $\lambda_{J0} > 0$, $\alpha_1 > 0$, $\alpha_2 > 0$ and $\alpha_1 + \alpha_2 < 1$, resulting in a joint normal conditional truncated by the inequality constraints. Draws of the vector $\lambda_{J1:T} = (\lambda_{J1}, \lambda_{J2}, \ldots, \lambda_{JT})'$ are produced automatically from (the degenerate) $p\left(\lambda_{J1:T}|V_{1:T}, Z_{1:T}^v, \Delta N_{1:T}^v, \phi, MF_{1:T}, BV_{1:T}\right)$ via the conditionally deterministic relationship in (27), with $l_0 = \lambda_{J0}$. Given the form of independent, informative priors invoked for μ_v and δ_J (as described above), inverse gamma and beta candidates are adopted for the respective MH sub-steps for these two parameters.

5 Empirical Application

5.1 Data description and preliminary analysis

In this section we report results of the application of the algorithm to daily annualized variance measures constructed from intraday spot and option price data for the S&P500 index, for the period July 26, 1999 to December 31, 2008. The first 1961 observations (corresponding to daily variance measures for the period July 26, 1999 to June 21, 2007) are used as the initial dataset for estimation, with the remaining 386 observations being reserved for the evaluation of the one-step-ahead probabilistic forecasts, based on an expanding sample window. The evaluation period covers: the period immediately preceding the global financial crisis (June 22, 2007 to July 27, 2007); the early period of the crisis - during which defaults on U.S. sub-prime mortgages began to impact on the viability of financial institutions and the availability of credit (late July 2007 to August 2008); and the period of historically unprecedented high levels of stock market volatility (towards the end of 2008), as mortgage brokers and major financial institutions, including Lehman Brothers, filed for bankruptcy, and active intervention by the U.S. Federal Reserve ensued. All index data has been supplied by the Securities Industries Research Centre of Asia Pacific (SIRCA) on behalf of Reuters, with the raw index data having been cleaned using methods similar to those of Brownlees and Gallo (2006). The CBOE VIX volatility index for the S&P500 is constructed using the model-free methodology; see www.cboe.com for more details. Hence, the option-implied measure (corresponding to $\widehat{MF}_t^{QV}/22$ in (9)) is based on the publicly available $\left(\frac{VIX_t}{100}\right)^2$, associated with day t. All references to MF_t in this empirical section are as an *average* (annualized) variance measure represented by $(VIX_t/100)^2 - \widehat{E}^* \left[\left. \sum_{t<s\leq t+22}^{N_{t+22}} (Z_s^p)^2 \right| \mathcal{F}_t \right] /22$; that is MF_t is reported on the same scale as BV_t, rather than as an aggregate (annualized) measure over a trading month. Accordingly, the relevant scale parameter for the MF_t measurement error is reported as $\sigma_{MF}/22$.

The BV_t measure in (5) is based on fixed five minute sampling, with a 'nearest price' method used to construct artificial returns five minutes apart. Note that the various forms of microstructure noise-adjusted measures that have appeared in the literature have their prime motivation in the case of data on traded assets, rather than observations on a constructed index. However, one could argue that the presence of stale prices in the index at the point of any recorded up-date, plus the inherent discreteness in the underlying prices, induce a form of noise. Hence, as an additional form of noise adjustment (that is, over and above the use of relatively distant five minute intervals), we use the subsampling and averaging techniques advocated by Zhang, Mykland and Aït-Sahalia (2005).[7]

In Figure 1 we reproduce plots of BV_t and MF_t (based on $c = 1.5$ in (10)) for the initial estimation period (Panel A) and for the period reserved for the evaluation of the probabilistic forecasts (Panel B). The empirical regularity of the option-implied variance exceeding (in the main) the realized variance is evident in both panels. Despite the option-implied measure being much less noisy than the spot-based measure, both measures exhibit broadly similar fluctuations, with there being only a slight tendency for the peaks in MF_t to lag those in BV_t. The extra noise associated with the spot-based measure is to be expected. The values of BV_t track daily volatility of the underlying itself whilst, in contrast, MF_t is a daily measure of the (average) volatility expected by the option market to prevail over the next month. The magnitude and variability of the realized variance and the option-implied measures increased significantly during the global financial crisis, demonstrated here by the enormous

[7] The numerical results reported in this section have been produced using the JAVA programming language. The key aspects of the pre-filtering of the raw S&P500 index data are as follows: 1) The use of data recorded only between 9.30am and 4.05pm, with a view to eliminating the effect of 'off-market' transactions on the index values; 2) The omission of all days on which less than 3 hours of trading occurred; 3) The use of the median value for any index values recorded simultaneously; and 4) The omission of any inordinately large jumps in the index value that were reversed within a short time interval. The authors would like to acknowledge the excellent research assistance of Chris Tse in producing the variance data series used in the analysis.

increase in both series towards the end of the sample in Panel B, with note to be taken of the differing vertical scales used in the two panels. The *extent* of this increase is deliberately highlighted by the use of the raw variance scale in Panel B, rather than the log scale which would typically be used when data exhibits such extreme variability within one figure.

With a view to providing some initial empirical justification for the variance jump risk premium model specified in (27), we produce a preliminary diagnostic of sorts, namely the autocorrelation function of the 'observed' jump risk premium in (30), with $\widehat{E}(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)$ computed as per (31) and based on the observed BV_t and MF_t (with $c = 1.5$) for the full sample period, July 26, 1999 to December 31, 2008. In producing this preliminary diagnostic, BV_t is used as a proxy for the latent V_t, while the required static parameters in ϕ are set equal to the corresponding marginal posterior means produced from a modified MCMC algorithm where V_t is replaced by BV_t for all t. The resulting autocorrelation function for l_{Jt}, plotted in Figure 2, provides clear evidence both of serial correlation in l_{Jt}, and the type of exponential decay that characterizes a short memory autoregressive process, such as that postulated for the latent quantity λ_{Jt} in (27). Further support for the model specified for λ_{Jt} is provided via both the magnitude of the estimated parameters of (27), reported in Section 5.2, and the results of the forecast evaluation exercise documented in Section 5.4.

5.2 Empirical estimates

In Table 1 we report the results based on estimation of the model in (21) to (27), for the initial sample period, July 26, 1999 to June 21, 2007. Marginal posterior means (MPM) and 95% higher posterior density (HPD) interval estimates are calculated from $G = 50,000$ MCMC draws, following a 50,000 draw burn-in period. We report results based on three different values of c in (10), $c = 1.0, 1.2, 1.5$. The first value of c implies that the risk neutral expectation of jump variation is equivalent to the objective expectation, while the second and third values (respectively) imply that the risk neutral expectation is 20% and 50% larger than the objective expectation. The estimates are seen to be very robust overall to the value of c used in the specification of MF_t.[8] Most notably, for all values of c the results indicate a high level of persistence in the latent variance process (small estimates of κ) and reasonably large (positive) estimates of α_1 and α_2 in the model for λ_{Jt}. As anticipated, the (time series) average of λ_{Jt}, $\overline{\lambda}_J$, (as estimated here by the average of the T values of λ_{Jt}, $t = 1, 2, ..., T$, in turn averaged over the G MCMC draws), is positive,

[8] Prompted by a referee, we experimented with values of $0 < c \leq 1$, in which case the assumption of a positive premium for price jump risk is not imposed. Results for $c = 0$ were also produced, corresponding to a case in which the VIX measure is assumed to contain information about the (risk neutral expectation of) integrated volatility only. The qualitative nature of the in-sample results was not affected by these choices of values of c. Indeed, the overall robustness of our results to this choice parameter is consistent with that documented in Bollerslev, Gibson and Zhou (2011).

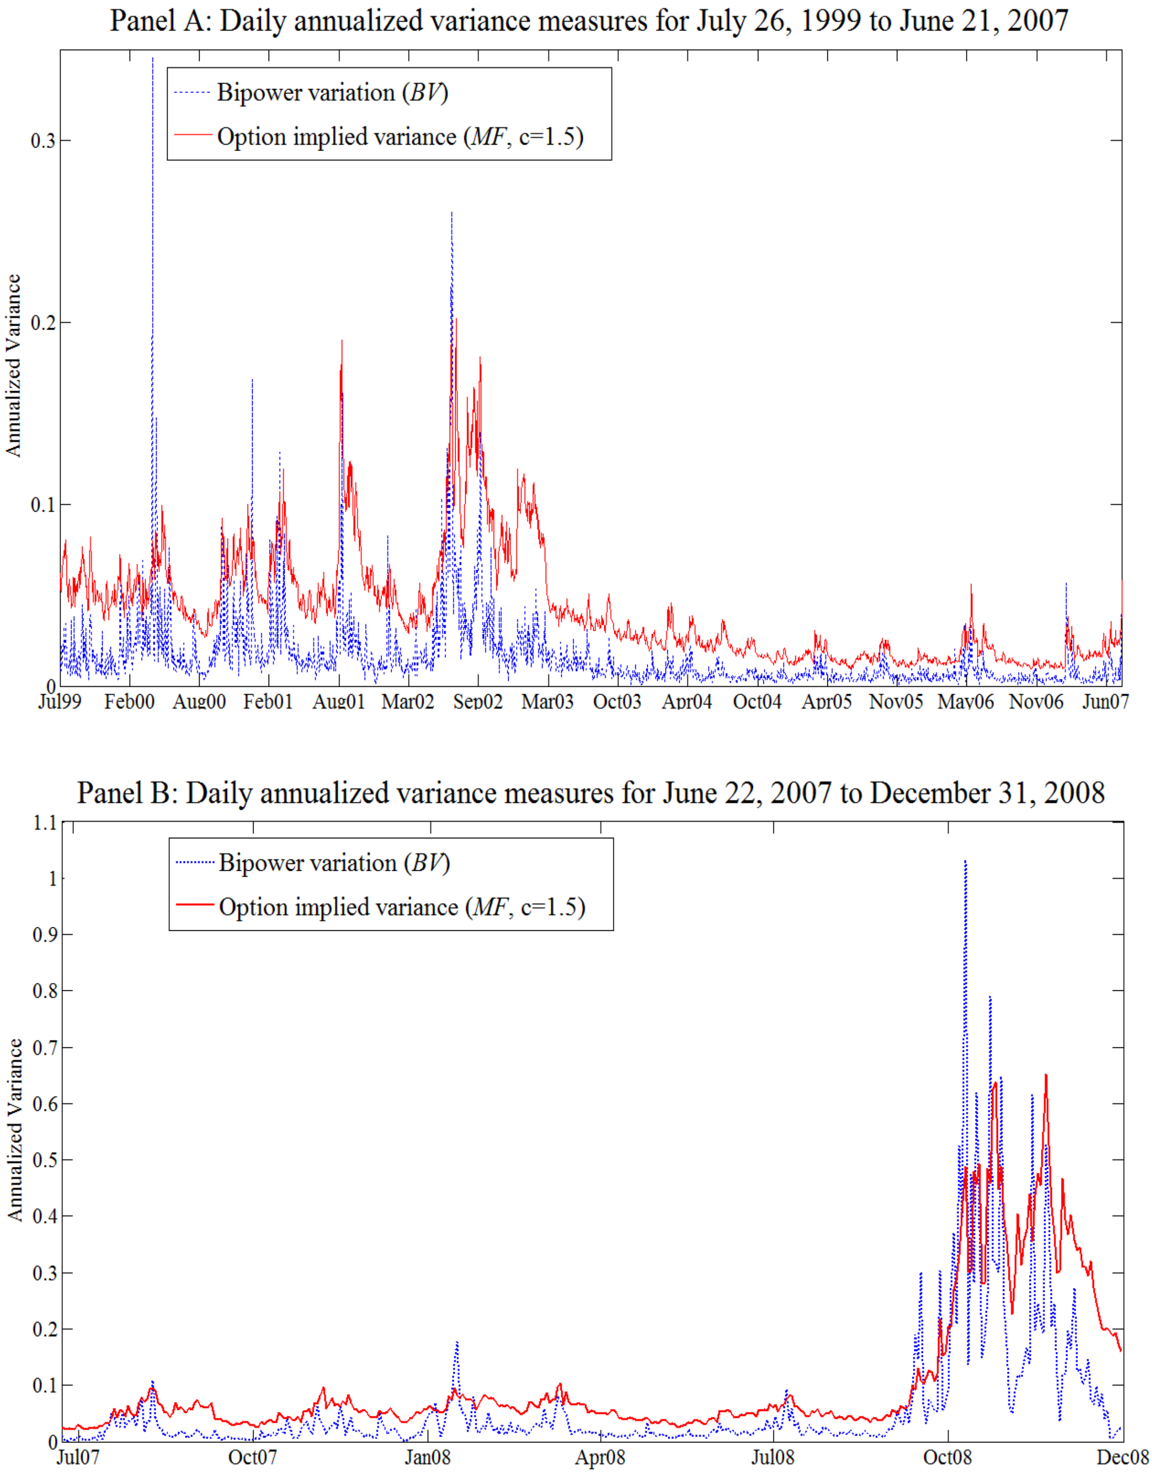

Figure 1: Plots of observed (annualized) variance measures: BV_t, a measure of daily integrated variance, and MF_t (with $c = 1.5$), a measure of the (average) risk neutral expection of integrated variance over the next trading month. Plots are given for the initial estimation period (Panel A) and the prediction evaluation period (Panel B). Note the differing vertical scales used in the two panels, with very extreme volatility observed towards the end of the evaluation period.

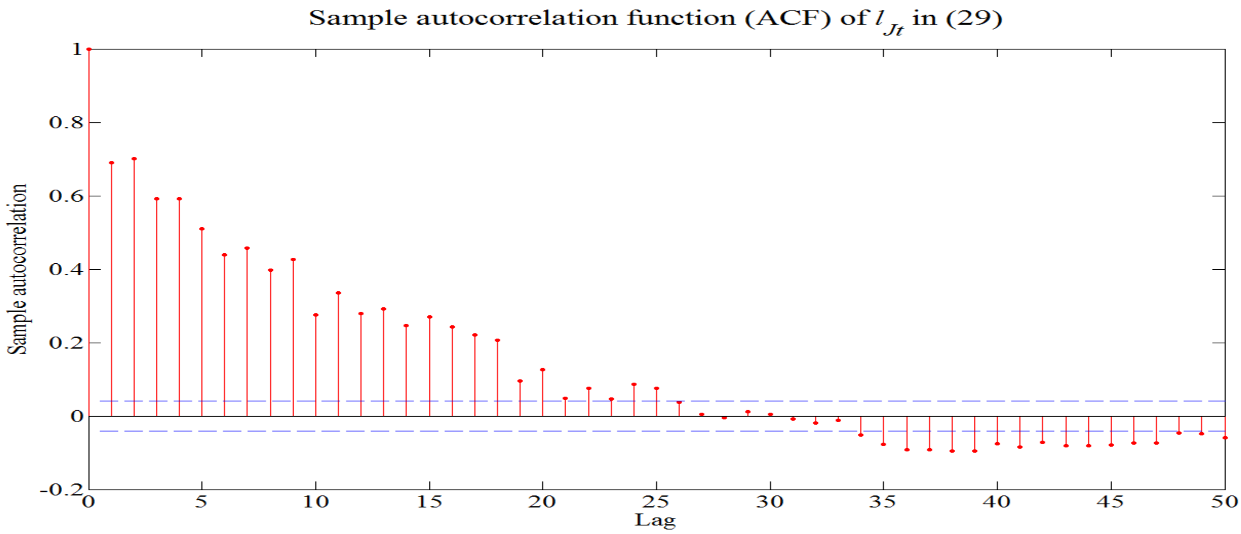

Figure 2: Sample autocorrelation function of the 'observed' jump variance risk premium, l_{Jt}, constructed from observed data and posterior point estimates of fixed parameters resulting from a modified MCMC analysis with $V_t \equiv BV_t$ for all t. Dashed lines denote 95% confidence bounds for each autocorrelation coefficient.

whilst the (constant) value of λ_D is negative. The magnitude of λ_D is small, consistent with values reported by Eraker (2004), indicating a slight increase in the persistence in volatility under the risk neutral measure. However, the magnitude of $\overline{\lambda}_J$ is large, in particular in comparison with the estimated magnitude of the mean of the actual variance jumps, μ_v, over the period. This result indicates the extreme sensitivity of the market to this aspect of the latent variance, with the implied risk neutral expectation of (variance) jump size (μ_v^*) being many orders of magnitude larger than the *actual* mean jump size. The point (and interval) estimates of δ_J indicate that the probability of a variance jump occurring on any one day ranges from approximately 2% to 5%. These magnitudes slightly exceed those reported in previous studies (based on different model specifications and for earlier sample periods), as summarized by Broadie, Chernov and Johannes (2007), where point estimates of the probability of variance jumps vary from 0.3% to 2%. The estimates of the mean of the variance jumps reported in Broadie *et al.*, μ_v, range from 0.018 to 0.037 (in annualized decimal form), and are thus are slightly higher than the magnitudes for this parameter recorded in Table 1, but are broadly consistent nevertheless.

Inefficiency factors are also reported in Table 1, estimated as the ratio of the variance of the sample mean of a set of MCMC draws of a given unknown, to the variance of the sample mean from a hypothetical independent sampler. The factors fall in the range of 3 to 190, with the average size of the variance jumps (μ_v) having the largest factor. Indeed, apart from this one parameter, all inefficiency factors are less than 73, and many are less than 10. Taking into account previously reported inefficiency factors for simpler volatility models, our MCMC algorithm can be deemed relatively efficient.[9] The acceptance rates for all parameters drawn using MH schemes range from 40% to 70%, with the acceptance rates for drawing $V_{1:T}$ (in blocks) ranging from 40% to 45%. Visual inspection of cumulative sum plots (Yu and Mykland, 1998) also indicated that the MCMC algorithm had converged, with μ_v being the slowest to converge, as is consistent with its high inefficiency factor.

5.3 Probabilistic forecasts of volatility and its risk premia

Bayesian predictions of the latent variance, V_{T+1}, and the dynamic variance jump risk premium, λ_{JT+1}, are produced by estimating the respective predictive densities, $p(V_{T+1}|MF_{1:T}, BV_{1:T})$ and $p(\lambda_{JT+1}|MF_{1:T}, BV_{1:T})$, using the MCMC draws. The predictive density of the diffusive risk premium, $p(\lambda_D V_{T+1}|MF_{1:T}, BV_{1:T})$, is derived from the resulting draws of the product, $\lambda_D V_{T+1}$, composed from the posterior draws of λ_D and V_{T+1}.

Recursive forecasts of V_{T+1}, λ_{JT+1} and $\lambda_D V_{T+1}$ (based on expanding windows) are produced over the evaluation period: June 22, 2007 to December 31, 2008. Given the robustness of the results to the value of c, we summarize the predictive results for $c = 1.5$ only. The three predictive densities are estimated using every 10^{th} of the $G = 50,000$ draws of all unknowns. The draws of ϕ are only updated (approximately) every 60 days throughout the evaluation period; the robustness of the static parameter estimates across sub-periods justifying this attempt to ease the computational burden associated with the production of the forecasts. In Figure 3, plots of both the marginal predictive means and 95% predictive intervals for V_{T+1}, λ_{JT+1} and $\lambda_D V_{T+1}$ are given respectively in Panels A, B and C. As the crisis deepens, both the level of the latent variance itself and the premium demanded by investors for variance risk (of both the diffusive and jump type), is seen to increase. The degree of uncertainty associated with all three latent variables, as measured by the width of the prediction intervals, is also *markedly* larger in the extreme crisis period, in late 2008, than in the earlier (and pre-) crisis period.

The predictive densities of λ_{JT+1} depicted in Panel B of Figure 3 highlight an interesting result, namely that the predictive mean of the time t variance jump risk premium is (very) occasionally negative. These occasions correspond to days in which particularly extreme movements in the spot market have occurred on the previous day(s), with the measure of spot volatility being significantly higher than its risk-neutral counterpart, contrary to the usual situation of the risk-neutral quantity exceeding the objective one. Multiple instances of

[9] For example, Kim, Shephard and Chib (1998), who estimate a stochastic volatility model using daily returns data, report inefficiency factors for certain of their algorithms in the range of 1.4 to 155. Liesenfeld and Richard (2006) report inefficiency factors in the range of 28 to 370, for an alternative stochastic volatility specification. Neither of these analyses involve option price data and the estimation of risk premia.

Table 1: Empirical results for the S&P500 stock index for July 26, 1999 to June 21, 2007, based on the measures, BV_t and MF_t. The value of c recorded is that used in the production of MF_t in (9).

Parameter	$c = 1.0$ MPM (95% HPD) *Inefficiency Factor*	$c = 1.2$ MPM (95% HPD) *Inefficiency Factor*	$c = 1.5$ MPM (95% HPD) *Inefficiency Factor*
κ	0.017 (0.013, 0.021) *20.09*	0.018 (0.014, 0.023) *32.14*	0.017 (0.013, 0.021) *15.76*
θ	0.0088 (0.0053, 0.0122) *16.89*	0.0087 (0.0054, 0.0118) *14.09*	0.0089 (0.0051, 0.0125) *15.42*
σ_v	0.0082 (0.0076, 0.0087) *40.68*	0.0081 (0.0075, 0.0089) *72.73*	0.0089 (0.0083, 0.0095) *39.31*
σ_{BV}	0.61 (0.59, 0.63) *6.20*	0.60 (0.57, 0.62) *7.52*	0.59 (0.57, 0.62) *7.38*
$\sigma_{MF}/22$	0.27 (0.19, 0.46) *3.52*	0.25 (0.18 0.44) *3.57*	0.24 (0.18, 0.44) *3.68*
μ_v	0.0076 (0.0060, 0.0098) *190.68*	0.0087 (0.0068, 0.0111) *159.10*	0.0080 (0.0063, 0.0104) *188.44*
δ_J	0.037 (0.025, 0.052) *45.87*	0.037 (0.024, 0.051) *43.03*	0.038 (0.025, 0.053) *46.26*
λ_D	-0.0071 (-0.0173, -0.0003) *4.62*	-0.0080 (-0.0188, -0.0003) *4.78*	-0.0076 (-0.0177, -0.0003) *4.60*
λ_{J0}	0.077 (0.004, 0.159) *3.58*	0.088 (0.004, 0.183) *3.58*	0.081 (0.004, 0.169) *3.59*
α_1	0.38 (0.02, 0.88) *4.04*	0.39 (0.02, 0.88) *4.07*	0.38 (0.02, 0.88) *4.01*
α_2	0.28 (0.01, 0.67) *3.89*	0.26 (0.01, 0.64) *3.89*	0.27 (0.01, 0.61) *3.85*
$\overline{\lambda}_J$	0.37	0.36	0.35

negative (point) predictions for λ_{JT+1} occur during the period of mid-September to October 2008, following the collapse of Lehman Brothers and subsequent interventions by the U.S. Federal Reserve, marking the depth of the financial crisis. Economically speaking, negative values of λ_{JT+1} are somewhat counter-intuitive, since they imply, all other things being equal, that investors are not willing to pay a positive premium for option cover of large jumps in the volatility of the underlying. As such, negative values for point predictions of this quantity could be viewed as reflecting a deficiency in our model. However, the overall behaviour of the probabilistic predictions of λ_{JT+1} is broadly consistent with plausible behaviour on the part of investors, and the upper 95% bounds associated with the *distributional* predictions - extending as they do into the positive region of the support of λ_{JT+1} at *all* points in time - lend further support to the proposition that the model is a reasonable simplification of reality.[10]

The one-step-ahead Bayesian predictives for the observable variance quantities are

$$p(BV_{T+1}|MF_{1:T}, BV_{1:T}) = \int p(BV_{T+1}|V_{T+1}, \phi)p(V_{T+1}|V_T, \phi)p(V_T, \phi|MF_{1:T}, BV_{1:T})dV_{T+1}dV_T d\phi$$

and

$$\begin{aligned} p(MF_{T+1}|MF_{1:T}, BV_{1:T}) &= \int [p(MF_{T+1}|V_{T+1}, \lambda_{JT+1}, \phi)p(\lambda_{JT+1}|\lambda_{JT}, \phi, MF_{1:T}, BV_{1:T}) \\ &\quad \times p(\lambda_{JT}, \phi|MF_{1:T}, BV_{1:T})p(V_{T+1}|V_T, \phi) \\ &\quad \times p(V_T, \phi|MF_{1:T}, BV_{1:T})d\lambda_{JT}dV_{T+1}dV_T d\phi] \,. \end{aligned}$$

Given the assumption of (conditionally) Gaussian measurement errors, these (marginal) predictives are estimated by averaging the Gaussian conditional predictives, $p(BV_{T+1}|V_{T+1}, \phi)$ and $p(MF_{T+1}|V_{T+1}, \lambda_{JT+1}, \phi)$ respectively, over the G MCMC draws of all unknowns, including V_{T+1} and λ_{JT+1}. These marginal predictive distributions are summarized in Figure 4, with the observed values also displayed. Over the full evaluation period, the empirical coverage of the 95% prediction intervals for BV_{T+1} and MF_{T+1} are 94.82% and 92.28%, respectively, with these figures highlighting the accuracy with which our approach predicts the observable variance measures (most notably the objective measure), even during the height of the financial crisis.

[10] We experimented with different ways of restricting the model to impose positive values of λ_{Jt} for each t. Given that the results overall were not qualitatively different from those based on the unrestricted model, we decided to report the latter results. Indeed, the occasional empirical occurrence of spot volatility exceeding the option implied measure in periods of extreme market stress is an interesting empirical feature and one that we felt should be allowed to influence risk premia estimates. We are not aware of any comparable empirical analyses that include the unprecedented period of high volatility in the latter half of 2008. For example, in their analysis of equity and variance risk premia, Bollerslev and Todorov (2009) deliberately omit the late 2007 to 2008 period to 'prevent our estimates of risk premia and investor fears to be unduely influenced by the recent financial market crises'.

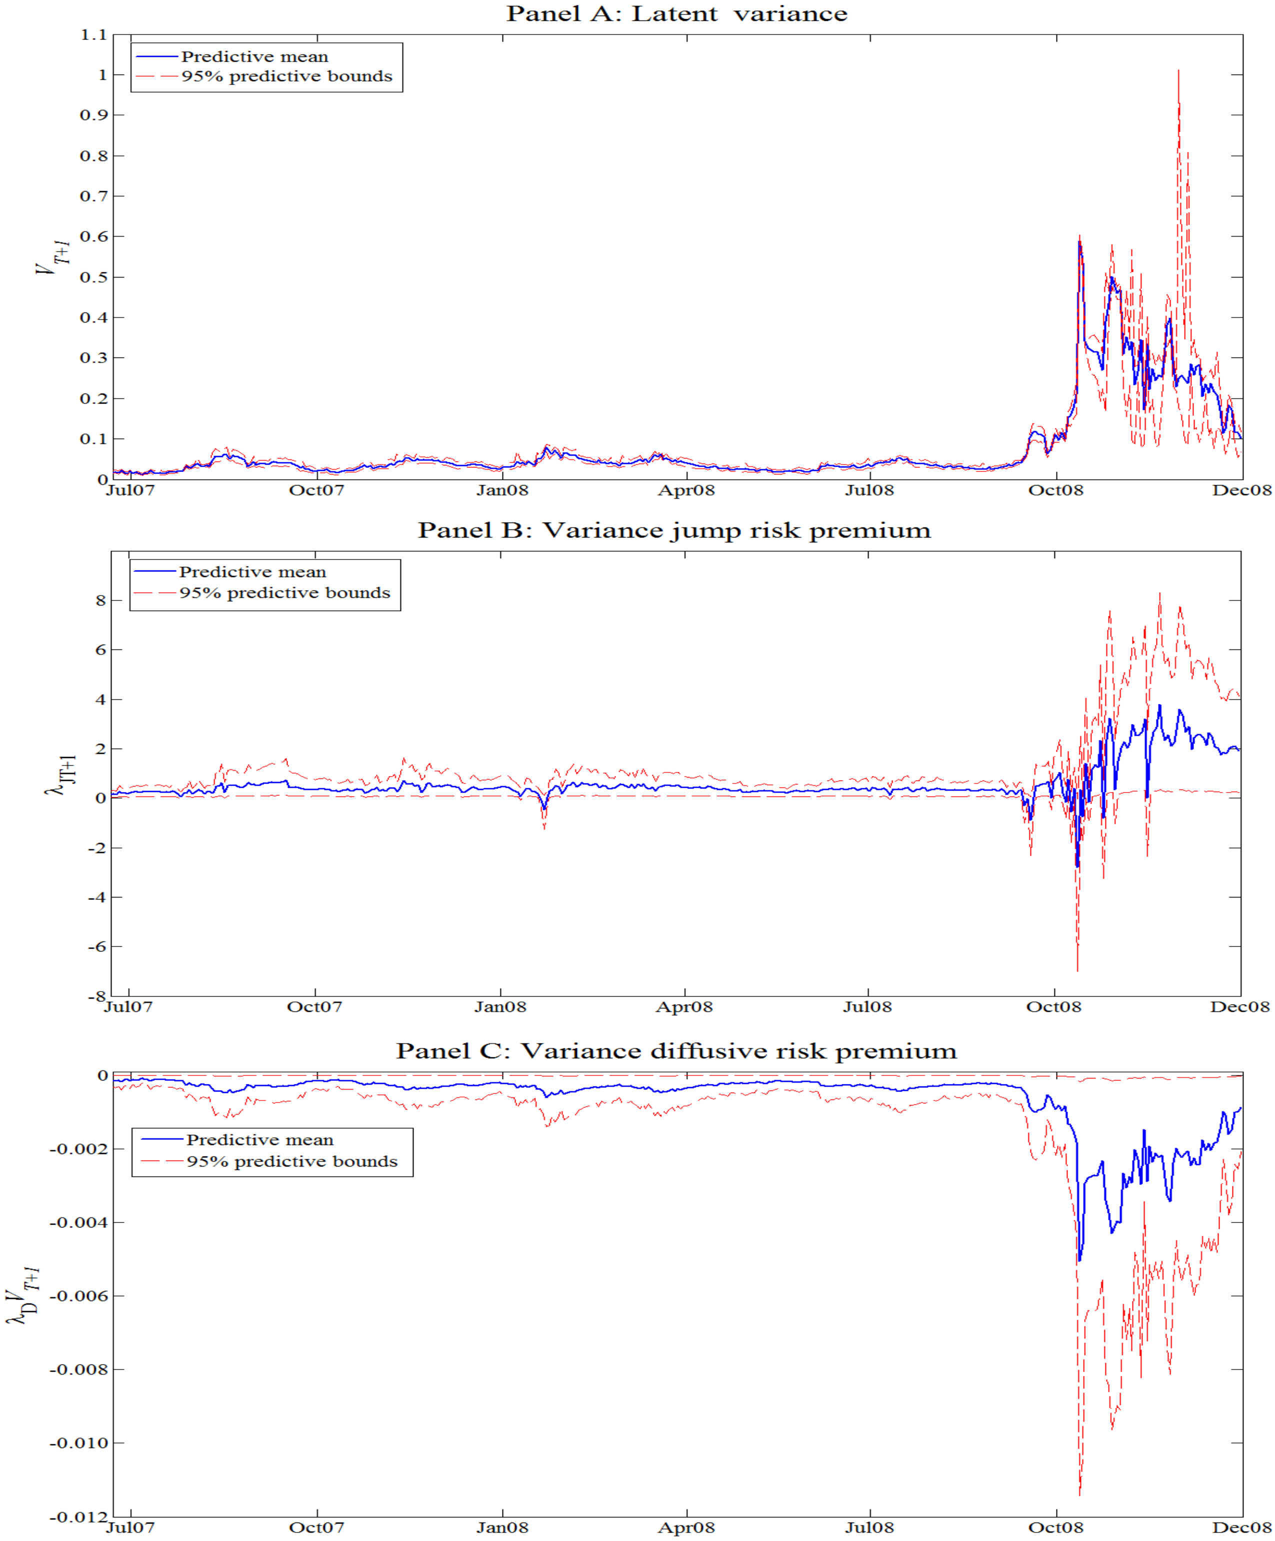

Figure 3: One-step-ahead predictive means and 95% predictive intervals for V_{T+1} (Panel A), λ_{JT+1} (Panel B) and $\lambda_D V_{T+1}$ (Panel C), over the evaluation period June 22, 2007 to December 31, 2008. The scale of all quantities relate to an annualized variance.

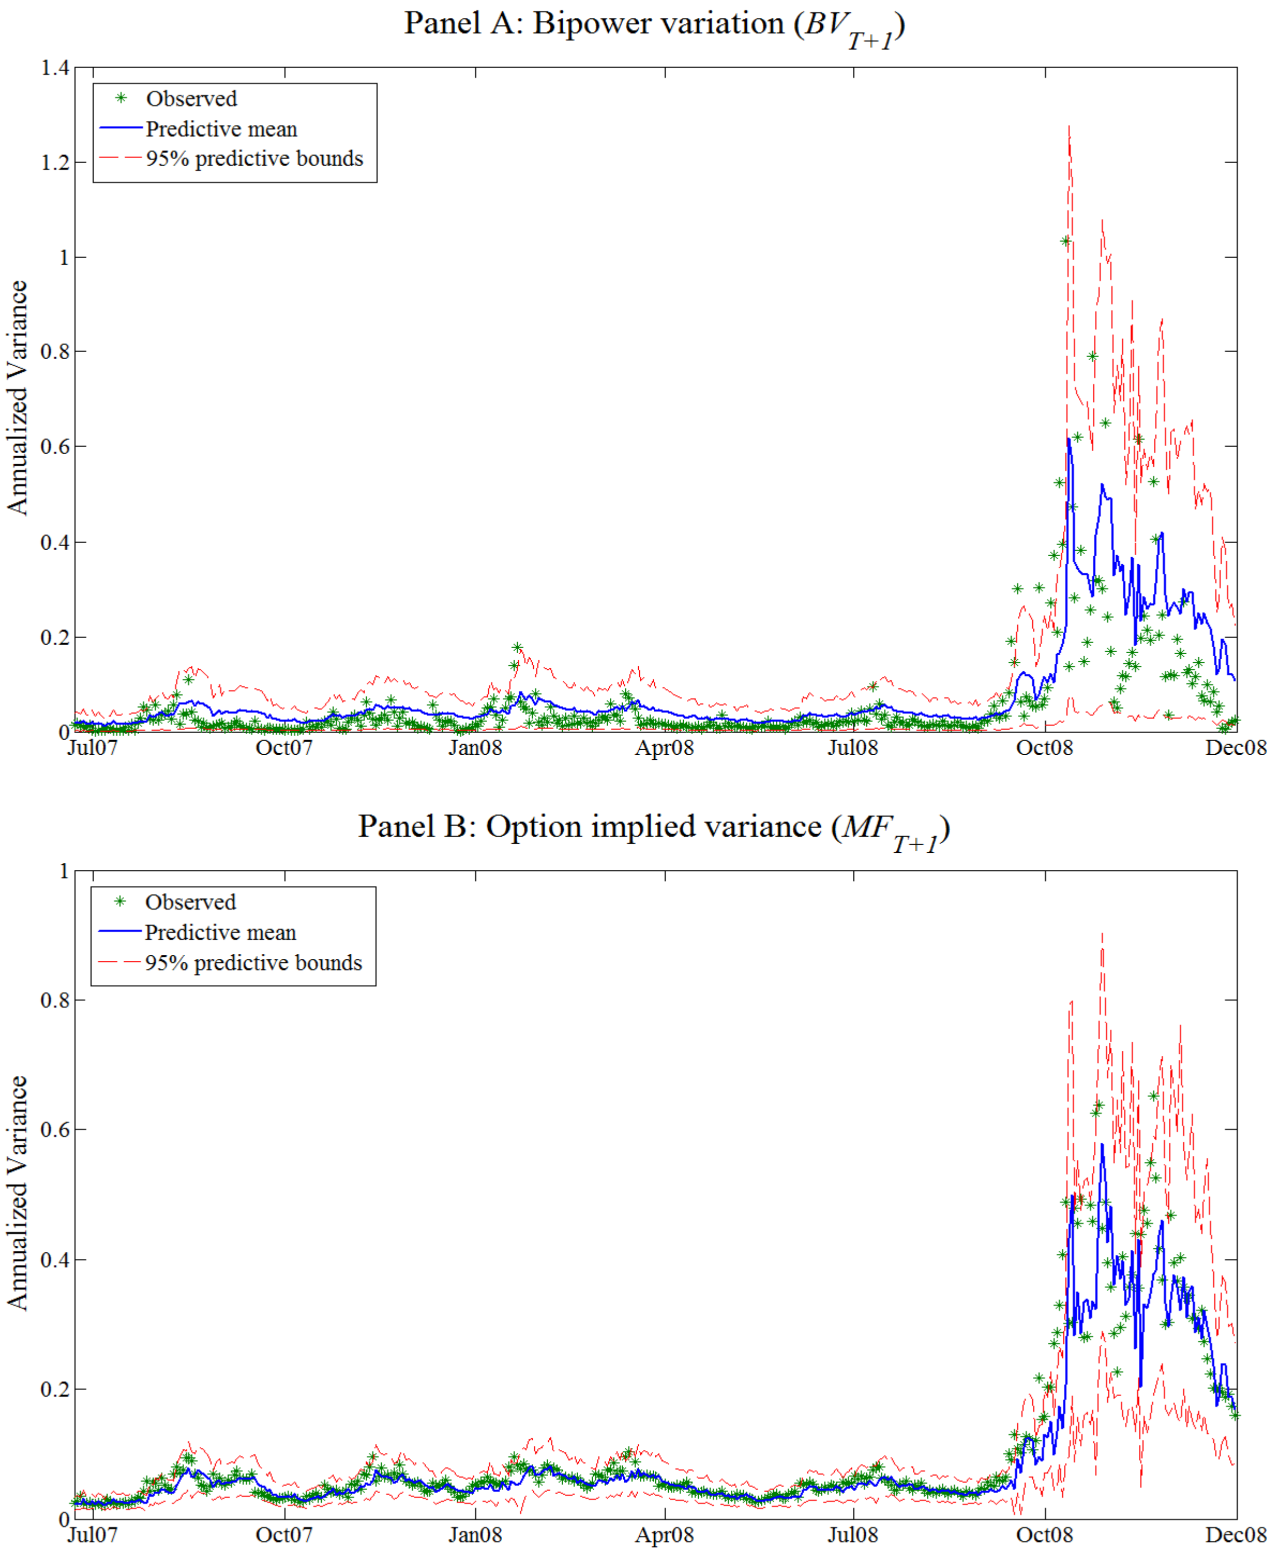

Figure 4: One-step-ahead predictive means, 95% predictive intervals and observed values of BV_{T+1} (Panel A) and MF_{T+1} (Panel B) for the evaluation period: June 22, 2007 to December 31, 2008. (As noted in the text, MF_t is reported throughout Section 5 on the same scale as BV_t.) The empirical coverage of the 95% prediction intervals for BV_{T+1} and MF_{T+1} are 94.82% and 92.28%, respectively.

5.4 Prediction comparison

Arguably the most novel aspects of our model are the use of the dual variance measures in a state space setting, plus the specification of a dynamic model for the jump risk premium. Hence, it is of interest to compare the predictive performance of our full model with that of two nested versions, namely one in which the jump risk premium is held constant, and a second in which the option-based measurement equation is omitted. In what follows we denote the full model, given by (21) to (27) by FULL-SSM; the model with a constant jump risk premium, given by (21) to (27), but with $\alpha_1 = \alpha_2 = 0$ in (27), as CONS-SSM; and the model with the BV measure only, given by (21) and (23) to (26), as BV-SSM. In the case of the third model, only the objective parameters can be estimated, and the only observable variance measure able to be predicted is BV_t.

Given the current practice of forecasting observed variance measures via univariate, observation-driven time series models, it is also of interest to compare the accuracy of our method with such simpler alternatives. In particular, it is of interest to compare the predictive accuracy of our short-memory model with that of models that cater for the long memory features of observed variance measures in general, including the measures analysed herein. With an extensive forecast comparison exercise being beyond the scope of this paper, we compare our state space-based forecasts with forecasts from just a small set of univariate models for each of BV_t and MF_t. Specifically, for both measures (expressed in logarithmic form) we consider the heterogeneous autoregressive (HAR) specification introduced by Corsi (2009) and used by Andersen, Bollerslev and Diebold (2007) and Bollerslev, Kretschmer, Pigorsch and Tauchen (2009), amongst many others. Using generic notation Y_t to denote a variance measure (either BV_t or MF_t here) and $\ln Y_{t,t+h} \equiv h^{-1} \sum_{i=1}^{h} \ln Y_{t+h}$ to denote the average of the daily values of $\ln Y_t$ over the horizon of h days (with $\ln Y_{t,t+1} = \ln Y_{t+1}$), the HAR model is given by

$$\ln Y_{t+1} = \beta_0 + \beta_D \ln Y_t + \beta_W \ln Y_{t-5,t} + \beta_M \ln Y_{t-22,t} + \varepsilon_{t+1}. \tag{34}$$

The horizons used to define the regressors in (34) correspond to the previous day, previous trading week and previous trading month respectively; the heuristic rationale of the model being that stock market volatility is determined by the behaviour of distinct groups of investors with different investment horizons. Whilst not formally a long-memory model, (34) has been shown to be capable of producing the slow autocorrelation decay that characterizes a long memory series.[11]

[11] The option-implied measure MF_t, is a measure of the risk neutral expectation of integrated variance over the next trading month. When $Y_t = MF_t$ in (34), $\ln Y_{t-h,t}$ (for $h = 5$, 22) captures the market's expectation of (monthly) integrated variance as recorded over the last h days, including day t. Hence, $\ln MF_{t-h,t}$ still

In order to best replicate typical statistical practice in the area, we estimate (34) using quasi-maximum likelihood estimation (QMLE), based on Gaussian innovations, both with and without a first order GARCH component in ε_{t+1} (denoted HAR1-GARCH and HAR1 respectively.) In addition, in order to capture some possible dependence of a measure computed from the prices from one market (spot or option) on the measure computed using the other data source, we augment the regressor set in (34) to include lagged versions of the alternative measure. Using generic notation X_t to denote the variance measure computed from the alternative data source to that used to compute Y_t, and with $\ln X_{t,t+h} \equiv h^{-1} \sum_{i=1}^{h} \ln X_{t+h}$, we specify the model

$$\begin{aligned} \ln Y_{t+1} &= \beta_0 + \beta_{D,Y} \ln Y_t + \beta_{W,Y} \ln Y_{t-5,t} + \beta_{M,Y} \ln Y_{t-22,t} \\ &\quad +\beta_{D,X} \ln X_t + \beta_{W,X} \ln X_{t-5,t} + \beta_{M,X} \ln X_{t-22,t} + \varepsilon_{t+1}, \end{aligned} \tag{35}$$

again both with and without a GARCH component in ε_{t+1} (denoted HAR2-GARCH and HAR2 respectively).

In order to keep the exercise manageable, we document here only empirical predictive coverage for the eight competing HAR models (four each for forecasting BV_t and MF_t respectively), supplemented by further performance statistics for HAR specifications in the applications in Sections 5.5.1 and 5.5.2. The results recorded in Panel A of Table 2 show the empirical coverage of both the BV_t and MF_t 95% prediction intervals produced by the HAR models to be uniformly lower than the corresponding coverage associated with the full state space model (FULL-SSM). The coverage accuracy of the alternative state space models - recorded in Panel B - is also inferior to that of the full model.

For the three state space models (FULL-SSM, CONS-SSM and BV-SSM) we supplement the coverage statistics in Table 2 with results based on the cumulative difference in logarithmic scores (CLS), computed over the evaluation period, as

$$CLS_{BV}(k) = \sum_{t=T+1}^{T+k} \ln \left[\frac{\widehat{p}_{FULL}(BV_t^o | MF_{1:t-1}, BV_{1:t-1})}{\widetilde{p}_{ALT}(BV_t^o | MF_{1:t-1}, BV_{1:t-1})} \right] \text{ and} \tag{36}$$

$$CLS_{MF}(k) = \sum_{t=T+1}^{T+k} \ln \left[\frac{\widehat{p}_{FULL}(MF_t^o | MF_{1:t-1}, BV_{1:t-1})}{\widetilde{p}_{ALT}(MF_t^o | MF_{1:t-1}, BV_{1:t-1})} \right], \tag{37}$$

for $k = 1, 2, ..., 386$. The notation $\widehat{p}_{FULL}(BV_t^o | MF_{1:t-1}, BV_{1:t-1})$ in (36) is used to denote the estimated predictive density of BV_t produced by the full state space model (FULL-SSM), and evaluated at the observed value of BV_t^o, whilst $\widetilde{p}_{ALT}(BV_t^o | MF_{1:t-1}, BV_{1:t-1})$ denotes the corresponding quantity produced under one of the two alternative state space models (i.e.

represents a (log) volatility measure (albeit a forward looking one) calculated over a horizon of h days and, thus, can be viewed as playing a similar role in producing pseudo long memory to $\ln BV_{t-h,t}$.

Table 2: Empirical coverage of the 95% prediction intervals for observable BV_t and MF_t quantities generated from univariate HAR models (Panel A) and alternative state space models (Panel B). The models in Panel A are defined throughout Section 5.4, whilst those in Panel B are defined in the first paragraph of Section 5.4.

	Competing Model	BV Predictions Empirical Coverage	MF Predictions Empirical Coverage
Panel A	HAR1	90.41%	87.05%
	HAR1-GARCH	91.19%	90.93%
	HAR2	90.67%	86.27%
	HAR2-GARCH	92.75%	91.19%
Panel B	FULL-SSM	94.82%	92.23%
	CONS-SSM	91.71%	99.74%
	BV-SSM	90.67%	N/A

either CONS-SSM or BV-SSM). Comparable notation for the predictive densities for MF_t is used in (37). A positive value for CLS_{BV} (CLS_{MF}) indicates that the full state space model outperforms the competitor, in terms of producing accurate probabilistic predictions of BV_t (MF_t) over the evaluation period.[12]

Over the entire evaluation period, and in comparison with CONS-SSM, $CLS_{BV} = 10.72$ and $CLS_{MF} = 127.82$, indicating that the full state space model provides superior predictions of both BV_t and MF_t to the model in which the jump premium parameter is held fixed. Plots of both $CLS_{BV}(k)$ and $CLS_{MF}(k)$ for $k = 1, 3, ..., 386$ are produced in Panels A and B of Figure 5. These plots enable one to visualize relative predictive performance as each new pair of BV_t and MF_t values are observed, and with reference to key market events (see Geweke and Amisano, 2010). From the plot in Panel A of Figure 5, CONS-SSM is seen to dominate FULL-SSM over the early period of the crisis, in terms of prediction of BV_t, but with the full model producing superior predictions in the presence of each high volatility event highlighted. Most notably, the relative predictive performance of FULL-SSM improves enormously in the most extreme period of the crisis, subsequent to the Lehman Brothers collapse, leading to the positive value overall for CLS_{BV}. Somewhat in contrast, in Panel B the probabilistic predictions of MF_t produced from FULL-SSM are seen to be almost *uniformly* more accurate than those produced by CONS-SSM, apart from a slight reversal

[12] As detailed in Geweke and Amisano (2010), the CLS can be viewed as a decomposition of an (estimated) generalized Bayes factor. In brief, the individual components in the sum in (36) (and (37)) are produced by repeated application of the MCMC simulator to expanding samples, beginning with the initial estimation sample of size T, rather than the accumulation of data occurring from $t = 1$ to $t = T$, as would the case for a Bayes factor.

of the relative performance of the two models during late September to mid October, 2008. It was during this latter period that the cluster of negative point forecasts of λ_{JT+1} flagged earlier, occurred, possibly leading to the inferior performance of the full model - with the time-varying jump premium - over this time. Certainly after the initial period of the US Federal Reserve intervention in October 2008 the full model begins to regain its superior position relative to the restricted form. The graph in Panel C of Figure 5 highlights the fact that it is the inclusion of the option-based information *per se* in the state space model that produces the most substantial improvement in the accuracy of the probabilistic predictions of BV_t, with CONS-SSM versus BV-SSM yielding CLS_{BV} values for the full period exceeding 1000. Most notably, a marked shift in favour of the dual-measure model occurs after the Merrill Lynch failure, with its dominance maintained thereafter.

5.5 Applications of the volatility and risk premia forecasts

Whilst the overall accuracy of the full state space model-based forecasts of the observed variance measures themselves is certainly worthy of note, a more convincing testimony to the (predictive) worth of any model comes from its ability to produce accurate forecasts of financial quantities into which volatility is an *input*. With this in mind, we conduct two exercises. First, in Section 5.5.1 we document the accuracy of estimates of the settlement prices of VIX futures, as based on forecasts of the latent variance produced by the full model. Secondly, in Section 5.5.2, we *augment* that model by an additional measurement equation based on daily returns (r_t), to produce recursive predictive distributions for r_{T+1}, from which the daily 5% and 1% VaR quantiles are extracted.

Finally, in the spirit of that Bollerslev, Gibson and Zhou (2011), in Section 5.5.3 predictions of the latent variance and the variance risk *premia* are used to extract a sequence of one-step-ahead predictions of an approximation to the relative risk aversion of the representative agent. Whilst it is not possible to conduct an assessment of accuracy in this case, due to the absence of 'observations' on risk aversion, it is still of interest to document this outcome of our model and to calibrate the results with comparable results recorded in the literature.

5.5.1 Pricing VIX futures

The CBOE introduced trading on futures contracts based on the VIX in March 2004. For a given futures contract with maturity date T_M, the settlement price quoted on day t, used for marking-to-market purposes, reflects the market's expectation of the VIX value at the time of maturity. The quantity $\tau(\frac{VIX}{100})^2$ ($=\widehat{MF}^{QV}$ in (9)) at time T_M is, in turn, a representation of the risk neutral expectation of $Q\mathcal{V}_{T_M,T_M+\tau}$, namely, quadratic variation over the period

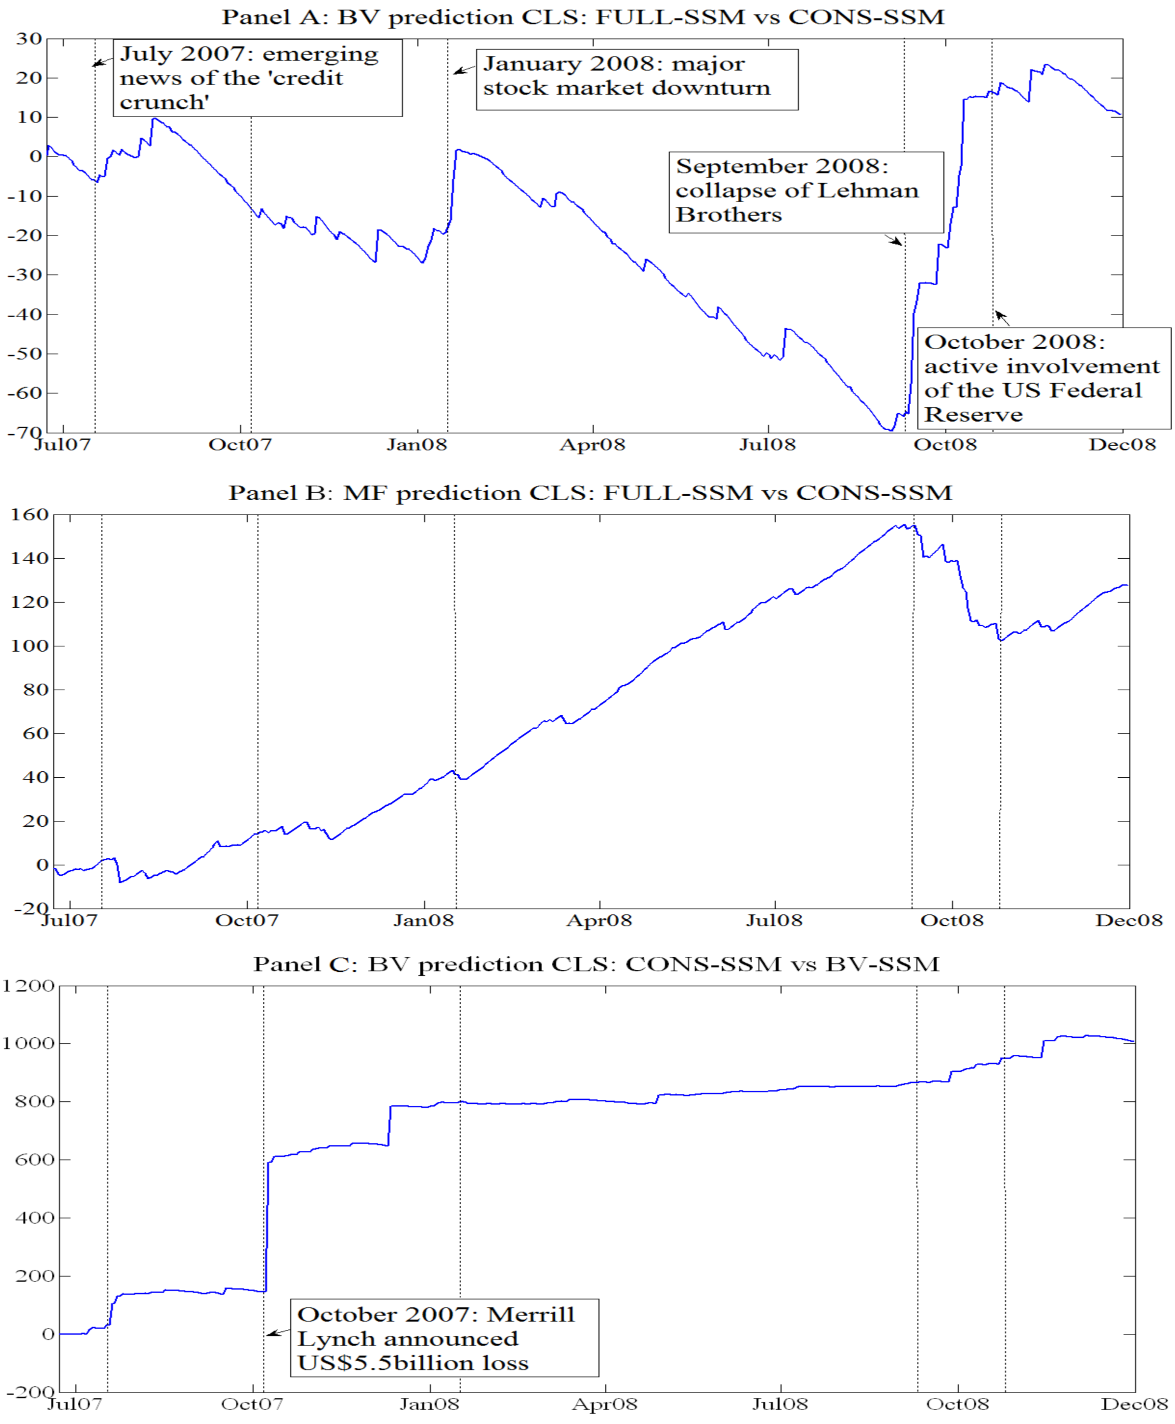

Figure 5: Cumulative logarithmic scores for the predictive performance of the alternative models. Panel A: prediction of BV_t; FULL-SSM against CONS-SSM; Panel B: prediction of MF_t; FULL-SSM against CONS-SSM; Panel C: prediction of BV_t; CONS-SSM against BV-SSM. All three models are defined in the first paragraph of Section 5.4. The dotted lines extending over all three panels indicate the timing of particular key market events, with identifying details of each marked event shown accordingly.

τ dated from T_M. It is thus of interest to ascertain the degree to which a model-based prediction of $E^*\left(Q\mathcal{V}_{T_M,T_M+\tau}|\mathcal{F}_{T_M}\right)$ accords with observed settlement prices on the VIX.

Under our model,

$$\begin{aligned}
E\left[E^*\left(Q\mathcal{V}_{T_M,T_M+\tau}|\mathcal{F}_{T_M}\right)|\mathcal{F}_t\right] &= E\left[E^*(\mathcal{V}_{T_M,T_M+\tau}|\mathcal{F}_{TM})|\mathcal{F}_t\right] \\
&\quad +E\left[E^*\left(\left.\sum_{T_M<s\leq T_M+\tau}^{N_{T_M+\tau}}(Z_s^p)^2\right|\mathcal{F}_{T_M}\right)|\mathcal{F}_t\right] \\
&= a_\tau^* E\left(V_{T_M}|\mathcal{F}_t\right)+b_\tau^*+\tau\delta_J\left[\mu_v+E(\lambda_{JT_M}|\mathcal{F}_t)\right] \\
&\quad +E\left[E^*\left(\left.\sum_{T_M<s\leq T_M+\tau}^{N_{T_M+\tau}}(Z_s^p)^2\right|\mathcal{F}_{T_M}\right)|\mathcal{F}_t\right],
\end{aligned}$$

where $E\left(V_{T_M}|\mathcal{F}_t\right)$ and $E(\lambda_{JT_M}|\mathcal{F}_t)$ are (T_M-t)-step-ahead predictions of the stochastic variance and the variance jump risk premium, respectively. These two quantities can be obtained via our assumed models for V_t and λ_{Jt}, (23) and (27) respectively, whilst

$$E\left[E^*\left(\left.\sum_{T_M<s\leq T_M+\tau}^{N_{T_M+\tau}}(Z_s^p)^2\right|\mathcal{F}_{T_M}\right)|\mathcal{F}_t\right]=E\left[cE\left(\left.\sum_{T_M<s\leq T_M+\tau}^{N_{T_M+\tau}}(Z_s^p)^2\right|\mathcal{F}_{T_M}\right)|\mathcal{F}_t\right]$$

is produced from the univariate time series modelling of price jump variation (as described in Section 2), with $c=1.5$ imposed.

Using posterior draws of all unknowns, draws of $\sqrt{E\left[E^*\left(Q\mathcal{V}_{T_M,T_M+\tau}|\mathcal{F}_{T_M}\right)|\mathcal{F}_t\right]}$ can be produced for $t=T,T+1,...,T+385$, over the evaluation period: June 22, 2007 to December 31, 2008, for futures contracts with the closest maturity date (i.e. the smallest value of T_M at any given point-in-time t, with T_M never being more than one month from day t), with a 95% probability interval produced from the draws in the usual way. The 95% intervals constructed in this way (with the draws appropriately adjusted to the scaling of the VIX) cover approximately 82% of the observed daily settlement values over the evaluation period. Whilst this represents an underestimate of the nominal coverage level, we note that the corresponding coverage from each of the four HAR-based comparator models (outlined in Section 5.4 and used to produce predictions of $Y_t=(VIX/100)^2$ at maturity date T_M, with $X_t=RV_t$ in (35)) do not exceed 71%.

5.5.2 Value at risk (VaR) prediction

Predicting the one-day-ahead 5% and 1% VaR for the market portfolio associated with the S&P500 index is equivalent to calculating the 5% and 1% quantile respectively for the predictive distribution for the portfolio return. However, given our primary aim of producing volatility (and risk premia) predictions that are independent of a specific model for the asset price, our state space model as it stands does not generate predictions of the asset return.

As such, we demonstrate here a method for augmenting the inferences drawn from the model based on the variance measures to produce, in turn, forecast distributions for the future return. Specifically, draws from the posterior distribution of the volatility model, conditional upon the spot- and option-based variance measures, are *resampled* R times to reflect *additional* conditioning on observed end-of-day returns. The model for the logarithmic price return, $r_t = p_t - p_{t-1}$, conditional on V_t, V_{t-1}, Z_t^v, and ΔN_t, where ΔN_t denotes the assumed common value of the discretized processes dN_t^p and dN_t^v, is based on an initial Euler approximation of (1) with $\mu_t = \mu - \frac{V_t}{2}$,

$$\begin{aligned} r_t &= \left(\mu - \frac{V_t}{2}\right) + \frac{\rho}{\sigma_v}\left(V_t - \kappa\theta - (1-\kappa)V_{t-1} - Z_t^v \Delta N_t\right) \\ &\quad + \left(\mu_p + \rho_J Z_t^v\right)\Delta N_t + \sqrt{V_t\left(1-\rho^2\right) + \sigma_p^2 \Delta N_t}\xi_{4t}, \end{aligned} \tag{38}$$

with $\{\xi_{4t}\} \overset{iid}{\sim} N(0,1)$, assumed to be independent of $\{\xi_{3t}\}, \{\xi_{2t}\}$ and $\{\xi_{1t}\}$ in (21), (22) and (23) respectively. In producing (38), we invoke the additional assumptions that $Z_t^p | Z_t^v \sim N(\mu_p + \rho_J Z_t^v, \sigma_p^2)$ and $\delta = \delta_J$. Under this specification, random jumps in the (logarithmic) price and the variance occur contemporaneously at rate δ_J, but with magnitudes determined respectively by a normal and an exponential distribution. The magnitudes of jumps in the price and variance processes are assumed to be correlated, governed by ρ_J. The two Brownian increments dB_t^p and dB_t^v are correlated with a coefficient ρ; however dB_t^i and dJ_t^i are assumed to be independent, for $i = \{p, v\}$. This type of model is often referred to in the literature as the stochastic volatility with contemporaneous jumps (SVCJ) model (e.g. Duffie, Pan and Singleton, 2000, Eraker, Johannes and Polson, 2003, Eraker, 2004, Broadie, Chernov and Johannes, 2007).

In order to render the resampling method computationally efficient, we make the simplifying assumption that $\sigma_p^2 = 0$, and introduce an unknown scale parameter, σ_r, absorbing all constant factors in the error variance, so that the error term in (38) collapses to $\sigma_r \sqrt{V_t}\xi_{4t}$. In addition, we introduce an additional regression parameter, β_v, resulting in a final model for returns given by

$$r_t = \mu + \beta_v V_t + \rho\left(\frac{V_t - \kappa\theta - (1-\kappa)V_{t-1} - Z_t^v \Delta N_t}{\sigma_v}\right) + \mu_p \Delta N_t + \rho_J Z_t^v \Delta N_t + \sigma_r \sqrt{V_t}\xi_{4t}. \tag{39}$$

Augmentation of the posterior density to include conditioning on daily returns involves multiplying the joint posterior in (33) by the product of the conditional distributions for r_t implied by (39), with normalization required to render this augmented posterior proper. To accomplish this task, draws of $V_{1:T}$, $Z_{1:T}^v$, $\Delta N_{1:T}$, θ, κ and σ_v , produced via the application of the MCMC algorithm described in Section 4, are resampled (as per the description in Appendix C), with draws of the return-*specific* parameters in (38), μ, β_v, ρ, μ_p, ρ_J and

σ_r, then produced by exploiting the regression structure in (39) in the usual way. With standard informative priors used for the return-specific parameters, the resampling method exploits the (closed-form) solution for the marginal likelihood for the vector of returns $r_{1:T} = (r_1, r_2, ..., r_T)'$, conditional on $V_{1:T}$, $Z^v_{1:T}$, $\Delta N_{1:T}$, θ, κ and σ_v. The draws are used to estimate the marginal predictive of the future return, $p(r_{T+1}|r_{1:T}, MF_{1:T}, BV_{1:T})$, from which the 5% and 1% VaR quantiles are calculated.

In Table 3 we report the empirical coverage statistics for the 5% and 1% VaR predictions produced by the (augmented) state space approach based on resampling $R = 5000$ times from the original $G = 50,000$ draws. We also report the p-values associated with tests of correct unconditional coverage and independence of exceedances of the VaR (Christoffersen, 1998). The former test assesses whether the empirical coverage differs significantly from the nominal level, while the latter tests for independence in the sequence of returns that exceed the VaR. An acceptable series of VaR predictions should fail to reject both of these hypotheses. Results are reported for the overall evaluation period, a period that excludes the final four months of 2008, and for that four month extreme period alone. As in Section 5.5.1, we compare the full state space model forecasts with distributional forecasts produced from the four HAR models described in Section 5.4, here with $Y_t = RV_t$, and $X_t = (VIX_t/100)^2$. The HAR-based variance forecasts of RV_{T+1}, which appropriately capture both diffusive *and* jump variation in the asset price at time $T + 1$, are inserted, along with the historical average return, into a conditionally normal model for the future return. Point predictions of the VaRs are then produced in the usual way, as the lower 5% and 1% quantiles of this predictive distribution for the return. To save space, we report in Table 3 the results for the HAR2-GARCH model only, as these results were uniformly superior to those of the other three HAR-based models considered.

With reference to both the overall period and the 'pre-extreme' period, the state space approach is seen to provide more accurate coverage than the HAR2-GARCH model, with empirical coverage that is a good deal closer to the nominal level. The empirical coverage of the state space model over the pre-extreme period in particular, for both the 5% and 1% VaR, is also not significantly different from its nominal level (at the 5% significance level), in contrast to the HAR2-GARCH model. In addition, the state space approach accurately captures the dynamics in returns - over both periods (overall and pre-extreme) and for both VaR levels - as the null hypothesis of independence in the exceedances is not rejected in all four cases. The HAR2-GARCH model, on the other hand, fails the independence test in the overall period for both the 5% and 1% VaR and for the 5% VaR in the pre-extreme period. The state space and HAR2-GARCH models perform comparably over the most extreme period, with excessive empirical coverages, indicating that both models have underestimated downside risk during this period of unprecedented instability.

To complement these VaR predictive results, we assess the probability integral transform (PIT) of the predictive return distribution from both the augmented state space and HAR2-GARCH models. The PIT, depicted in Figure 6 together with the 95% confidence bounds, indicates that both models have quite poor calibration in the lower tail of the return distribution, with both rejecting the null hypothesis of uniformity based on a Chi-squared test. Nevertheless, a visual inspection of Figure 6 highlights that the augmented state space model displays better calibration properties overall.

Table 3: Empirical coverage and p-values for the unconditional coverage (UC) and independence (IND) tests for 5% and 1% one-step-ahead VaR predictions. The overall evaluation period constitutes the 386 trading days from June 22, 2007 to December 31, 2008. Both the augmented state space model and the HAR2-GARCH model for returns are detailed in the text of Section 5.5.2.

		5% VaR		*1% VaR*	
		Augmented State Space	HAR2-GARCH	Augmented State Space	HAR2-GARCH
Pre-extreme period	Empirical Coverage	7.31%	12.96%	0.33%	6.98%
June 22, 2007	UC Test	0.08	0.00	0.18	0.00
to Aug 31, 2008	IND Test	0.06	0.00	0.93	0.08
Extreme period	Empirical Coverage	17.65%	20.00%	10.59%	9.41%
Sept 01, 2008	UC Test	0.00	0.00	0.00	0.00
to Dec 31, 2008	IND Test	0.60	0.31	0.14	0.19
Overall period	Empirical Coverage	10.10%	14.51%	2.59%	7.51%
	UC Test	0.00	0.00	0.0088	0.0000
	IND Test	0.58	0.00	0.47	0.03

5.5.3 Prediction of risk aversion

In modelling both variance risk premia (diffusive and jump) as dynamic processes, we are effectively modelling time variation in the risk aversion of the representative investor. Following Bollerslev, Gibson and Zhou (2011), we link the risk aversion parameter, γ, associated with the power utility function for a representative investor, to both $\lambda_D V_t$ and λ_{Jt}. Appendix D shows that the following approximate relationship holds between risk aversion in period $T+1$, γ_{T+1}, the two risk premia, $\lambda_D V_{T+1}$ and λ_{JT+1}, and certain other unknowns in the model,

$$\gamma_{T+1} \simeq \frac{\lambda_D V_{T+1} - \delta_J \lambda_{JT+1}}{\rho \sigma_v V_{T+1} + \rho_J \delta_J \mu_v^2}. \tag{40}$$

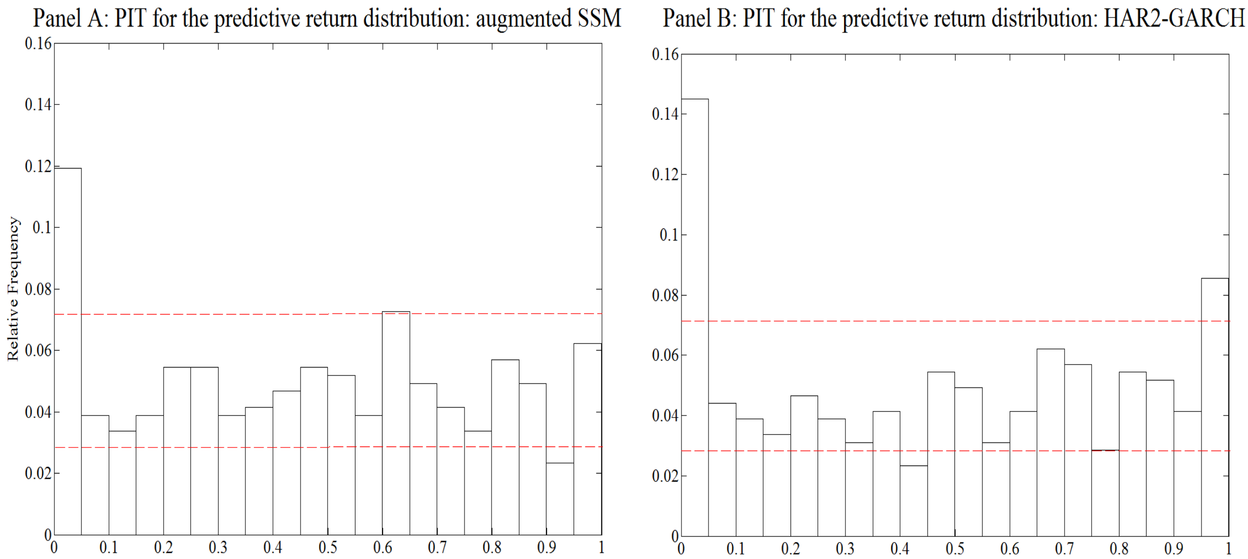

Figure 6: Histogram and 95% confidence bounds from the probability integral transform (PIT) for the observed return relative to the predictive distributions generated by the augmented state space model (SSM) (Panel A) and the HAR2-GARCH model for RV_t (Panel B).

Thus, draws of all unknowns on the right hand side of (40) (with draws of the full set of these unknowns, including ρ and ρ_J, obtained as part of the resampling exercise in Section 5.5.2) can be used to produce draws of (an approximation to) γ_{T+1} and those draws used, in turn, to estimate the predictive density of γ_{T+1}. Similar to the plots presented in Figure 3, a plot summarizing the predictive distribution of γ_{T+1} is produced in Figure 7. In the earlier period, in particular, the mean values of γ_{T+1} fluctuate around values that are broadly consistent with the range of estimates - produced via very different means - that have been reported for this parameter in the literature. (See Cochrane, 2005, Bakshi and Madan, 2006 and Bollerslev, Gibson and Zhou, 2011, for some recent discussion and documentation of these values.) However, as is to be expected, given the relationship between γ_{T+1} and the variance risk premia, the implied risk aversion of the representative investor increases *dramatically* during the extreme crisis period towards the end of 2008, together with a widening of the predictive bounds. In particular, the widening bounds translate into very uncertain estimates of the risk aversion level associated with a rational representative investor. (See Bollerslev and Todorov, 2009, for related discussion). Due to the nature of the relationship defined in (40), we also observe negative values of the measures of risk aversion when the variance jump risk premium parameter, λ_{JT+1}, is negative. These occasions, as discussed earlier, correspond to occasional days during the very height of the financial crisis, in which the spot measure of volatility was significantly higher than the option implied measure. On such

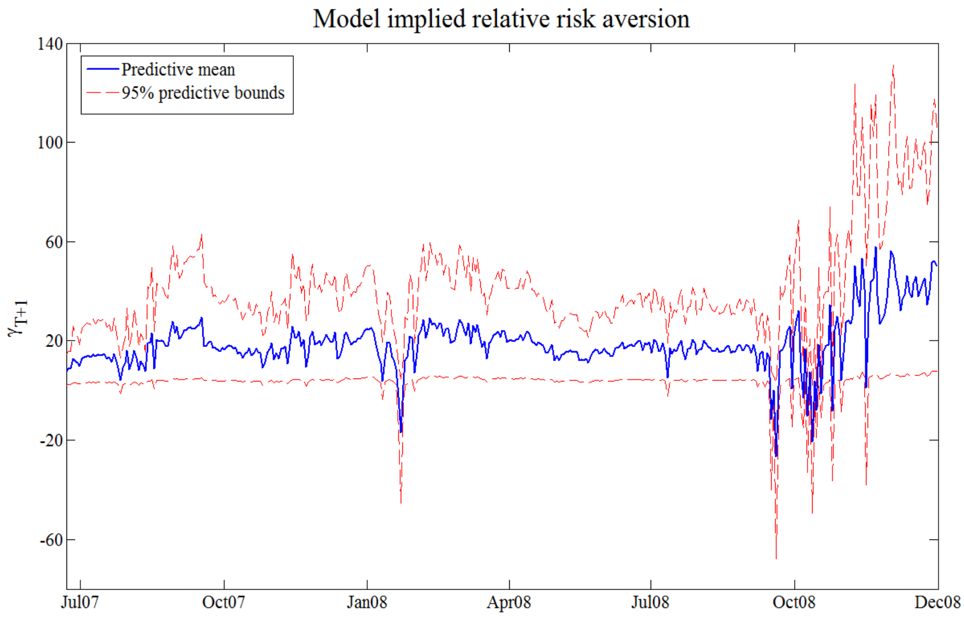

Figure 7: One-step-ahead predictive mean and 95% predictive interval for γ_{T+1} (the relative risk aversion parameter for a representative investor with power utility) for June 22, 2007 to December 31, 2008.

days, the equilibrium conditions required for the derivation of (40) may be questioned.

6 Conclusions

This paper is the first to use non-parametric option- and spot-based volatility measures in a non-linear state space framework for the purpose of producing probabilistic predictions of volatility and its risk premia. The absence of price measures in the state space model means that detailed parametric modelling of the asset price, including any premium for price jump risk, can be avoided. The resultant forecasts of volatility and its premia are thus robust to any reasonable assumptions about the asset price process. The premium for variance jump risk is modelled as a dynamic process driven by a past function of the measurements, with a linear (in variance) premium adopted for diffusive risk. The distributional forecasts then quantify the uncertainty associated with future values - including extreme values - of volatility itself and volatility risk of both types.

An empirical investigation using spot and option-based measures for the S&P 500 index sheds light on the changes in the predictive distributions of all quantities of interest over the period of the recent global financial crisis. Our state space approach is shown to provide very accurate predictive coverage over the full out-of-sample period, for both observable variance measures, and superior predictions to those of alternative univariate observation-driven models. Comparison with more restrictive forms of the state space model also illustrates that the use of the dual measures, plus the specification of a dynamic model for the jump risk premium, improves the accuracy with which the observable quantities can be forecast, in particular during the most volatile period of the recent crisis. The distributional forecasts of all latent quantities are used, in turn, to produce probabilistic predictions of related quantities - namely the market prices of futures written on the VIX index and the VaR on the market portfolio. Via a particular form of representative agent model, we link the dynamic risk premia to the risk aversion parameter, enabling probabilistic forecasts of the (approximate) risk aversion of a representative investor to be produced. Our model quantifies the changes in investor risk aversion over the financial crisis period, where, as might be anticipated, risk aversion is seen to increase dramatically in both level and variability as the crisis deepens. Most notably, the results point to the fact that *extreme* values for this behavioural parameter would have been predicted, with non-negligible probability, during the height of the stock market turmoil.

References

[1] Aït-Sahalia, Y. and Mancini, L. 2008. Out of Sample Forecasts of Quadratic Variation, *Journal of Econometrics,* 147: 17-33.

[2] Andersen, T.G., Bollerslev, T. and Diebold, F.X. 2007. Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility, *The Review of Economics and Statistics*, 89: 701-720.

[3] Andersen, T.G., Bollerslev, T. and Diebold, F.X. 2010. Parametric and Non-parametric Volatility Measurement, *Handbook of Financial Econometrics* (Eds. Y. Aït-Sahalia and L.P. Hansen), North Holland, Amsterdam.

[4] Andersen, T.G., Bollerslev, T., Diebold, F.X. and Labys, P. 2003. Modelling and Forecasting Realized Volatility, *Econometrica*, 71: 579-625.

[5] Bakshi, G. & Madan, D. 2006. A Theory of Volatility Spreads, *Management Science*, 52: 1945–1956.

[6] Barndorff-Nielsen, O.E. and Shephard, N. 2002. Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models, *Journal of the Royal Statistical Society B*, 64: 253-280.

[7] Barndorff-Nielsen, O.E. and Shephard, N. 2004. Power and Bipower Variation with Stochastic Volatility and Jumps, *Journal of Financial Econometrics*, 2: 1-37.

[8] Barndorff-Nielsen, O.E. and Shephard, N. 2006. Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation, *Journal of Financial Econometrics*, 4: 1-30.

[9] Bates, D.S. 2000. Post-87 Crash Fears in the S&P 500 Futures Option Market, *Journal of Econometrics*, 94: 181-238.

[10] Bekaert, G., Engstrom, E., and Xing, Y. 2009. Risk, Uncertainty and Asset Prices, *Journal of Financial Economics*, 91: 59-82.

[11] Blair, B.J, Poon, S-H. and Taylor, S.J. 2001. Forecasting S&P100 Volatility: the Incremental Information Content of Implied Volatilities and High Frequency Index Returns, *Journal of Econometrics*, 105: 5-26.

[12] Bollerslev, T. 1986. Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics*, 31: 307-327.

[13] Bollerslev, T., Gibson, M. and Zhou, H. 2011. Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities, *Journal of Econometrics*, 160: 235–245.

[14] Bollerslev, T., Kretschmer, U., Pigorsch, C. and Tauchen, G. 2009. A Discrete-Time Model for Daily S&P500 Returns and Realized Variations: Jumps and Leverage Effects, *Journal of Econometrics*, 150: 151-166.

[15] Bollerslev, T., Sizova, N. and Tauchen, G. 2012. Volatility in Equilibrium: Asymmetries and Dynamic Dependencies, *Review of Finance*, 16: 31-80. Advanced access, doi:10.1093/rof/rfr005, 1-50.

[16] Bollerslev, T., Tauchen, G. and Zhou, H. 2009. Expected Stock Returns and Variance Risk Premia, *Review of Financial Studies*, 22: 4463-4492.

[17] Bollerslev, T. and Todorov, V. 2009. Tails, Fears and Risk Premia, *CREATES Research Paper No. 2009-26, School of Economics and Management*, *University of Aarhus*, *Denmark.*

[18] Brandt, M.W. and Wang, K.Q. 2003. Time-Varying risk Aversion and Unexpected Inflation. *Journal of Monetary Economics*, 50: 1457-1498.

[19] Breeden, D.T. 1979. An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities, *Journal of Financial Economics*, 7:265-196.

[20] Britten-Jones, M. and Neuberger, A. 2000. Option Prices, Implied Price Processes and Stochastic Volatility, *The Journal of Finance*, LV: 839-866.

[21] Broadie, M., Chernov, M. and Johannes, M. 2007. Model Specification and Risk Premia: Evidence from Futures Options, *The Journal of Finance*, LXII: 1453-1490.

[22] Brownlees, C.T. and Gallo, G.M. 2006. Financial Econometric Analysis at Ultra-High Frequency: Data Handling Concerns, *Computational Statistics and Data Analysis*, 51: 2232-2245.

[23] Busch, T, Christensen, B.J., and Nielsen, M.O. 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.

[24] Carr, P. and Wu, L. 2009. Variance Risk Premiums, *Review of Financial Studies,* 22: 1311-1341.

[25] Carter, C.K. and Kohn, R. 1994. On Gibbs Sampling for State Space Models, *Biometrika*, 81: 541-553.

[26] Chernov, M. and Ghysels, E. 2000. A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation, *Journal of Financial Economics*, 56: 407-458.

[27] Christoffersen, P. F. 1998. Evaluating Interval Forecasts, *International Economic Review*, 39: 841-862.

[28] Cochrane, J.H. 2005. *Asset Pricing*, Second Edition, Princeton University Press.

[29] Corradi, V. and Swanson, N. 2006. Predictive Density and Conditional Confidence Interval Accuracy Tests, *Journal of Econometrics*, 135: 187–228.

[30] Corsi, F. 2009. A Simple Approximate Long Memory Model of Realized Volatility, *Journal of Financial Econometrics*, 7: 174-196.

[31] Cox, J.C., J.E. Ingersoll and Ross, S.A. 1985. An Intertemporal General Equilibrium Model for Asset Prices, *Econometrica*, 53: 363–384.

[32] Creal, D.D. 2008. Analysis of Filtering and Smoothing Algorithms for Levy-driven Stochastic Volatility Models, *Computational Statistics and Data Analysis*, 52: 2863-2876.

[33] Duan, J.C. and Yeh, C.Y. 2010. Jump and Volatility Risk Premiums Implied by VIX, *Journal of Economic Dynamics and Control*, 34: 2232-2244.

[34] Duffie, D., Pan J. and Singleton, K. 2000. Transform Analysis and Asset Pricing for Affine Jump-Diffusions, *Econometrica*, 68: 1343-1376.

[35] Elerian, O., Chib, S. and Shephard, N. 2001. Likelihood Inference for Discretely Observed Non-linear Diffusions, *Econometrica,* 69: 959-993.

[36] Engle, R.F. and Russell, J.R. 1998. Autoregressive Conditional Duration: A New Approach for Irregularly Spaced Transaction Data, *Econometrica*, 66: 987-1007.

[37] Eraker, B. 2001. MCMC Analysis of Diffusion Models with Applications to Finance, *Journal of Business and Economic Statistics*, 19: 177-199.

[38] Eraker, B. 2004. Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices, *The Journal of Finance*, LIX: 1367-1403.

[39] Eraker, B. 2008. The Volatility Premium, *Working Paper*, available for download from http://www.marginalq.com/eraker/.

[40] Eraker, B., Johannes, M. and Polson, N. 2003. The Impact of Jumps in Volatility and Returns, *The Journal of Finance*, LVIII: 1269-1300.

[41] Forbes C.S., Martin, G.M. and Wright J. 2007. Inference for a Class of Stochastic Volatility Models Using Option and Spot Prices: Application of a Bivariate Kalman Filter, *Econometric Reviews, Special Issue on Bayesian Dynamic Econometrics*, 26: 387-418.

[42] Frühwirth-Schnatter, S. 1994. Data Augmentation and Dynamic Linear Models, *Journal of Time Series Analysis*, 15: 183-202.

[43] Garcia, R., Lewis, M., Pastorello, S. and Renault, E. 2011. Estimation of Objective and Risk-Neutral Distributions Based on Moments of Integrated Volatility, *Journal of Econometrics*, 160: 22-32.

[44] Geweke, J. 2005. *Contemporary Bayesian Econometrics and Statistics*, Wiley, New York.

[45] Geweke, J. and Amisano, G. 2010. Comparing and Evaluating Bayesian Prediction Distributions of Asset Returns, *International Journal of Forecasting*, 26: 216-230.

[46] Gneiting, T. 2008. Editorial: Probabilistic forecasting, *Journal of the Royal Statistical Society (A)*, 171: 319-321.

[47] Gneiting, T., Balabdaoui, F. and Raftery, A. 2007. Probabilistic Forecasts, Calibration and Sharpness, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69: 243–268.

[48] Gordon, S. and St-Amour, P. 2004. Asset Returns and State-dependent Risk Preferences, *Journal of Business and Economic Statistics*, 22: 241-252.

[49] Guo, D. 1998. The Risk Premium of Volatility Implicit in Currency Options, *Journal of Business and Economic Statistics*, 16: 498-507.

[50] Jacquier, E. and Miller, S. 2010. The Information Content of Realized Volatility, *Working Paper*, HEC, University of Montreal.

[51] Jiang, G.J. and Tian, Y.S. 2005. The Model-Free Implied Volatility and its Information Content, *The Review of Financial Studies*, 18: 1305-1342.

[52] Johannes, M., Polson, N.G. and Stroud, J.R. 2009. Optimal Filtering of Jump-Diffusions: Extracting Latent States from Asset Prices, *Review of Financial Studies*, 22: 2759-2799.

[53] Jones, C. 2003. The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets, *Journal of Econometrics*, 116: 181-224.

[54] Kim, S., Shephard, N. and Chib, S. 1998. Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models, *Review of Economic Studies*, 65: 361-393.

[55] Koopman, S.J., Jungbacker, B. and Hol, E. 2005. Forecasting Daily Variability of the S&P100 Stock Index using Historical, Realized and Implied Volatility Measurements, *Journal of Empirical Finance*, 12: 445-475.

[56] Liesenfeld, R. and Richard, J-F. 2006. Classical and Bayesian Analysis of Univariate and Multivariate Stochastic Volatility Models, *Econometric Reviews*, 25: 335-360.

[57] Liu, J.S. 2001. *Monte Carlo Strategies in Scientific Computing*, Springer, New York.

[58] Liu, C. and Maheu, J. 2009. Forecasting Realized Volatility: A Bayesian Model-Averaging Approach, *Journal of Applied Econometrics,* 24: 709-733.

[59] Martens, M., van Dijk, D. and de Pooter, M. 2009. Forecasting S&P500 Volatility: Long Memory, Level Shifts, Leverage Effects, Day-of-the-Week Seasonality, and Macroeconomic Announcements, *International Journal of Forecasting*, 25: 282-303.

[60] Martens, M. and Zein, J. 2004. Predicting Financial Volatility: High-Frequency Time Series Forecasts Vis-a-Vis Implied Volatility, *Journal of Futures Markets*, 24: 1005-1028.

[61] Martin, G.M., Reidy, R. and Wright, J. 2009. Does the Option Market Produce Superior Forecasts of Noise-Corrected Volatility Measures? *Journal of Applied Econometrics*, 24: 77-104.

[62] McCabe, B.P.M., Martin, G.M. and Harris, D.G. 2011. Efficient Probabilistic Forecasts for Counts, *Journal of the Royal Statistical Society, Series B*, 73: 253-272.

[63] Pan, J. 2002. The Jump-risk Premia Implicit in Options: Evidence from an Integrated Time-series Study, *Journal Of Financial Economics*, 63: 3-50.

[64] Polson, N. G. and Stroud, J.R. 2003. Bayesian Inference for Derivative Prices, *Bayesian Statistics*, 7: 641-650.

[65] Pong, S., Shackleton, M.B., Taylor, S.J. and Xu, X. 2004. Forecasting Currency Volatility: a Comparison of Implied Volatilities and AR(FI)MA Models, *Journal of Banking and Finance*, 28: 2541-2563.

[66] Smith. A.F.M. and Gelfand, A.E. 1992. Bayesian Statistics without Tears: A Sampling Resampling Perspective, *The American Statistician*, 46: 84-88.

[67] Stroud, J.R., Müller, P. and Polson, N.G. 2003. Nonlinear State-space Models with State-Dependent Variances, *Journal of the American Statistical Association*, 98: 377-386.

[68] Tauchen G. and Zhou, H. 2011. Realized Jumps on Financial Markets and Predicting Credit Spreads, *Journal of Econometrics*, 160: 102–118

[69] Todorov, V. 2010. Variance Risk-Premium Dynamics: The Role of Jumps, *Review of Financial Studies*, 23: 345-383.

[70] Tokdar, S.T. and Kass, R.E. 2010. Importance Sampling: A Review, *Wiley Interdisciplinary Reviews: Computational Statistics*, 2: 54-60.

[71] Wu, L. 2011. Variance Dynamics: Joint Evidence from Options and High-frequency Returns, *Journal of Econometrics*, 160: 280-287.

[72] Yu, B. and Mykland, P. 1998. Looking at Markov Samplers through CUMSUM Path Plots: A Simple Diagnostic Idea, *Statistics and Computing*, 8: 275-286.

[73] Zellner, A. 1986. On Assessing Prior Distributions and Bayesian Regression Analysis with g-Prior Distributions, in *Bayesian Inference and Decision Techniques: Essays in Honours of Bruno de Finetti* (eds. P.K. Goel and A. Zellner), North Holland, Amsterdam.

[74] Zhang, L., Mykland, P.A. and Aït-Sahalia, Y. 2005. A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High Frequency Data, *Journal of the American Statistical Association*, 100: 1394-1411.

Appendix A: Generation of $V_{1:T}|Z^v_{1:T}, \Delta N^v_{1:T}, \phi, MF_{1:T}, BV_{1:T}$. As the state variable, V_t, appears in the error terms of the state equation in (23) and the measurement equations in (21) and (22), a closed form representation of the conditional posterior distribution for the stochastic variance vector, $V_{1:T}$, is not available. In this paper we extend an approach suggested by Stroud, Müller and Polson (2003) and augment the state space model with two mixture indicator vectors: $g^v_{1:T} = (g^v_1, g^v_2, ..., g^v_T)'$ associated with the latent variance vector, and $g^o_{1:T} = (g^o_1, g^o_2, ..., g^o_T)'$ associated with the observed bivariate variance measure. Each element of a mixture indicator vector takes on the value of an integer from 1 to K, and defines a suitable linearization of the relevant state or observation equation. The mixture indicator vectors are then used to establish a candidate draw, within a MH subchain, given the previously sampled latent variance vector, $V_{1:T}$, conditional upon all other parameters ϕ, jump occasions $\Delta N^v_{1:T}$ and variance jump sizes $Z^v_{1:T}$. After constructing an expanded state variable, $\{g^v_{1:T}, g^o_{1:T}, V_{1:T}\}$, and producing MCMC draws from $p\left(g^v_{1:T}, g^o_{1:T}, V_{1:T}|Z^v_{1:T}, \Delta N^v_{1:T}, \phi, MF_{1:T}, BV_{1:T}\right)$, by discarding the draws of $g^v_{1:T}$ and $g^o_{1:T}$ the remaining latent volatilities are seen as draws from the desired full conditional posterior distribution, $p\left(V_{1:T}|Z^v_{1:T}, \Delta N^v_{1:T}, \phi, MF_{1:T}, BV_{1:T}\right)$.

Prior to sampling, a grid of values over the space of plausible stochastic variance values, $\{\overline{\mu}_1, \overline{\mu}_2, ..., \overline{\mu}_K\}$, along with a set of associated bandwidth parameters, $\{\overline{\sigma}_1, \overline{\sigma}_2, ..., \overline{\sigma}_K\}$, are established, determining prior probability weights for both mixture indicators. The prior weight for the state variance indicator, $g^v_t = k$, is proportional to a normal density for the lagged variance, V_{t-1}, centred on $\overline{\mu}_k$ and with associated variance $\overline{\sigma}^2_k$. In a similar way, the prior weight for the observation indicator, $g^o_t = k$, is proportional to a normal density for the current variance, V_t, centred on $\overline{\mu}_k$ and with variance $\overline{\sigma}^2_k$. To keep the expressions uncluttered, in what follows the requisite conditioning on $\{Z^v_{1:T}, \Delta N^v_{1:T}, \phi, MF_{1:T}, BV_{1:T}\}$ has been suppressed.

The MH algorithm for the expanded state variable $\{g^v_{1:T}, g^o_{1:T}, V_{1:T}\}$ proceeds as follows:

1. Sample $g^v_{1:T}$ given $\{g^o_{1:T}, V_{1:T}\}$ from T independent multinomial distributions

$$p^a\left(g^v_t = k|g^o_{1:T}, V_{1:T}\right) = \frac{\left(\overline{\sigma}_k\sqrt{\overline{\mu}_k}\right)^{-1}\exp\left\{-\frac{(V_t - E[V_t|V_{t-1}])^2}{2\sigma^2_v\overline{\mu}_k} - \frac{(\overline{\mu}_k - V_{t-1})^2}{2\overline{\sigma}^2_k}\right\}}{\sum_{j=1}^K\left(\overline{\sigma}_j\sqrt{\overline{\mu}_j}\right)^{-1}\exp\left\{-\frac{(V_t - E[V_t|V_{t-1}])^2}{2\sigma^2_v\overline{\mu}_j} - \frac{(\overline{\mu}_k - V_{t-1})^2}{2\overline{\sigma}^2_j}\right\}},$$

for each $t = 1, 2, ..., T,$ where $E\left[V_t|V_{t-1}\right] = \kappa\theta + (1-\kappa)V_{t-1} + Z^v_t\Delta N^v_t.$

2. Sample $g^o_{1:T}$ given $\{g^v_{1:T}, V_{1:T}\}$ from T independent multinomial distributions, with

$$p^a(g^o_t = k|g^v_{1:T}, V_{1:T}) = \frac{(\overline{\sigma}_k\overline{\mu}^2_k)^{-1}\exp\left\{-\frac{(BV_t - V_t)^2}{2\sigma^2_{BV}\overline{\mu}^2_k} - \frac{(MF_t - E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t))^2}{2\sigma^2_{MF}\overline{\mu}^2_k} - \frac{(\overline{\mu}_k - V_t)^2}{2\overline{\sigma}^2_k}\right\}}{\sum_{j=1}^K(\overline{\sigma}_j\overline{\mu}^2_j)^{-1}\exp\left\{-\frac{(BV_t - V_t)^2}{2\sigma^2_{BV}\overline{\mu}^2_j} - \frac{(MF_t - E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t))^2}{2\sigma^2_{MF}\overline{\mu}^2_j} - \frac{\left(\overline{\mu}_j - V_t\right)^2}{2\overline{\sigma}^2_j}\right\}},$$

for $t = 1, 2, ..., T,$ where $E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t) = a^*_\tau V_t + b^*_\tau + \tau\left[\mu_v + \lambda_{Jt}\right]\delta_J.$

3. Sample $\widetilde{V}_{1:T}$ given $\{g^v_{1:T}, g^o_{1:T}\}$ using an FFBS sampling algorithm (see Carter and Kohn, 1994, or Frühwirth-Schnatter, 1994) and a candidate state space model designed to have a smoothed state distribution that closely approximates that of the augmented model, given $\{g^v_{1:T}, g^o_{1:T}\}$. The candidate state space model for each of day $t = 1, 2, ..., T-1$ is comprised of a (4×1) observation vector, $(\ BV_t,\ \ MF_t,\ \ \overline{\mu}_{g^o_t},\ \ \overline{\mu}_{g^v_{t+1}}\ )'$, whose distribution conditional on the state variable V_t is normal with mean vector equal to $(\ V_t,\ \ (b^*_\tau + \tau\left[\mu_v + \lambda_{Jt}\right]\delta_J + a^*_\tau V_t),\ \ V_t,\ \ V_t\ )'$, and diagonal covariance matrix with the corresponding nonzero elements given by $\sigma^2_{BV}\overline{\mu}^2_{g^o_t}$, $\sigma^2_{MF}\overline{\mu}^2_{g^o_t}$, $\overline{\sigma}^2_{g^o_t}$ and $\overline{\sigma}^2_{g^v_{t+1}}$, respectively. For day T, the (3×1) candidate state space model is slightly modified, as the last element of the observation vector is not required. The candidate state space model uses the scalar state equation

$$V_t = (1-\kappa)V_{t-1} + (\kappa\theta + Z^v_t\Delta N^v_t) + \omega_t, \qquad \omega_t \sim N\left(0, \sigma^2_v\overline{\mu}_{g^v_t}\right),$$

for all $t = 1, 2..., T,$ with $V_0 = \theta + \frac{\mu_v\delta_J}{\kappa}$.

4. Accept $\widetilde{V}_{1:T}$ with probability determined by the MH acceptance rule

$$\min\left\{1, \prod_{t=1}^T\left[\frac{C(V_{t-1})p^a(V_t|V_{t-1})p(\widetilde{V}_t|\widetilde{V}_{t-1})C(V_t)p^a(BV_t, MF_t|V_t)p(BV_t, MF_t|\widetilde{V}_t)}{C(\widetilde{V}_{t-1})p^a(\widetilde{V}_t|\widetilde{V}_{t-1})p(V_t|V_{t-1})C(\widetilde{V}_t)p^a(BV_t, MF_t|\widetilde{V}_t)p\left(BV_t, MF_t|V_t\right)}\right]\right\},$$

where $V_{1:T}$ denotes the previous draw in the Markov chain. Here, $C(V_t)$ is the normalizing constant resulting from the linear approximation, with

$$C\left(V_t\right) = \sum_{k=1}^K\left(2\pi\overline{\sigma}^2_k\right)^{-1/2}\exp\left\{-\frac{1}{2}\left(\frac{\overline{\mu}_k - V_t}{\overline{\sigma}_k}\right)^2\right\}, \text{ for all } t = 0, 1, 2, ..., T,$$

the candidate and target state transition densities are

$$\begin{aligned} p^a\left(V_t|V_{t-1}\right) &= C\left(V_{t-1}\right)^{-1}\sum_{k=1}^{K}\left[\left(\overline{\sigma}_k\sigma_v\sqrt{\overline{\mu}_k}\right)^{-1}\right. \\ &\quad \left.\times\exp\left\{-\frac{\left(V_t-\left(\kappa\theta+Z_t^v\Delta N_t^v\right)-\left(1-\kappa\right)V_{t-1}\right)^2}{2\sigma_v^2\overline{\mu}_k}-\frac{\left(\overline{\mu}_k-V_{t-1}\right)^2}{2\overline{\sigma}_k^2}\right\}\right], \end{aligned}$$

and

$$p\left(V_t|V_{t-1}\right)=\left(\sigma_v\sqrt{V_{t-1}}\right)^{-1}\exp\left(-\frac{\left(V_t-\left(\kappa\theta+Z_t^v\Delta N_t^v\right)-\left(1-\kappa\right)V_{t-1}\right)^2}{2\sigma_v^2V_{t-1}}\right),$$

respectively, and the candidate and target measurement densities are, respectively,

$$\begin{aligned} p^a(BV_t,MF_t|V_t) &= C\left(V_t\right)^{-1}\sum_{j=1}^{K}\left(\sigma_{BV}\sigma_{MF}\overline{\sigma}_j\overline{\mu}_j^2\right)^{-1} \\ &\quad \times\exp\left\{-\frac{\left(BV_t-V_t\right)^2}{2\sigma_{BV}^2\overline{\mu}_j^2}-\frac{\left(MF_t-E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)\right)^2}{2\sigma_{MF}^2\overline{\mu}_j^2}-\frac{\left(\overline{\mu}_j-V_t\right)^2}{2\overline{\sigma}_j^2}\right\}\Bigg], \end{aligned}$$

and

$$p(BV_t,MF_t|V_t)=\left(\sigma_{BV}\sigma_{MF}V_t^2\right)^{-1}\exp\left\{-\frac{\left(BV_t-V_t\right)^2}{2\sigma_{BV}^2V_t^2}-\frac{\left(MF_t-E^*(\mathcal{V}_{t,t+\tau}|\mathcal{F}_t)\right)^2}{2\sigma_{MF}^2V_t^2}\right\}.$$

5. Discard $g_{1:T}^v$ and $g_{1:T}^o$.

Appendix B: Prior and Posterior Predictive Analysis In order to assess the appropriateness of the prior specification, in particular the informative priors specified for the variance jump parameters, δ_J and μ_v, we conduct both prior and posterior predictive analyses; see Geweke (2005) for a detailed outline. The predictive analyses are based on a comparison of the predictive distributions of pertinent 'model checking' functions with observed values, where the predictive distributions are produced by repeated simulation from the model, including the joint prior. Simulation from the non-informative components of the joint prior occurs by using very broad, but bounded, supports. The checking functions are chosen to be functions of the observed variance measures that reflect important features of the data that we want the model to capture. As in Section 5.1, here MF_t refers to $(VIX_t/100)^2-\widehat{E}^*\left[\sum_{t<s\leq t+22}^{N_{t+22}}(Z_s^p)^2\middle|\mathcal{F}_t\right]/22$.

Since our primary aim is to model (for the purpose of forecasting) the dynamics in volatility and its risk premia, plus to cater appropriately for large volatility jumps, we use as checking functions the first order sample autocorrelations of BV_t, MF_t and the 'observed' conditional risk premium: $\overline{BV}_{t,t+22} - MF_t$ (denoted respectively by ρ_{BV}, ρ_{MF} and ρ_{RP}) plus measures of skewness (Sk) and kurtosis (Ku) calculated from the full sample of BV_t and MF_t values. The symbol $\overline{BV}_{t,t+22}$ here denotes the forward average of the daily values of BV_t over the one month (or 22 trading days) subsequent to period t.

Whilst there are a number of different ways in which the observed values of the checking function can be assessed relative to the estimated predictives, we take the simplest (albeit somewhat frequentist) approach here, calculating 'predictive p-values', namely the probability of obtaining a value for the checking function as or more extreme than the observed value. In Table 4, we report prior and posterior predictive p-values for two different models. Model M_1 (Panel A) refers to both the full SSM model and the prior specification described in Section 4.1 (and as used in producing all numerical results), whereas model M_2 (Panel B) replicates M_1, except that the prior standard deviation for both δ_J and μ_v is doubled. In the table, each checking function is denoted by z with its corresponding observed value denoted by z^o. For each model M_i, $i = 1, 2$, the prior predictive p-value equals the probability that $z > z^o$ conditioned only on that model M_i, whilst the corresponding posterior predictive p-value further conditions on the observed data. The results demonstrate that the model is capable of producing the observed quantities under both the prior and posterior predictive distributions, with non-zero (and non-unit) p-values. Crucially, the posterior results indicate that, with the possible exception of ρ_{MF}, the model (and prior) are able to produce, in artificial replications, values for these key statistics that tally well with the statistics calculated from the empirical data, and that this conclusion is robust to the precise specification of the informative priors for the jump parameters. In addition, despite the inflation of the prior variances undertaken in the experimental setting, the marginal posterior densities for all fixed parameters remained largely unchanged, indicating that the joint posterior density is dominated by information from the data.

Appendix C: Sampling/Importance Resampling from the Augmented State Space Model Given the specification for returns in (39), the joint posterior density for *all* unknowns, conditional on the three vectors of observed data, $r_{1:T}$, $MF_{1:T}$, and $BV_{1:T}$, can be decomposed as follows

$$\begin{aligned} & p\left(\mu, \beta_v, \rho, \mu_p, \rho_J, \sigma_r, V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \phi | r_{1:T}, MF_{1:T}, BV_{1:T}\right) \\ \propto\; & p\left(\mu, \beta_v, \rho, \mu_p, \rho_J, \sigma_r, | r_{1:T}, V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v\right) \\ & \quad \times p\left(r_{1:T} | V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v\right) p\left(V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \phi | MF_{1:T}, BV_{1:T}\right). \quad (41) \end{aligned}$$

In order to sample from (41), we first draw from $p\left(V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \phi | MF_{1:T}, BV_{1:T}\right)$ using the MCMC algorithm for the volatility model as described in Section 4. Then, given the

Table 4: Prior and posterior predictive p-values for model checking functions

z	Panel A: Prior specification from Section 4.1		Panel B: More diffuse informative priors for δ_J and μ_v	
	$\Pr(z > z^o \mid M_1)$	$\Pr(z > z^o \mid BV, MF, M_1)$	$\Pr(z > z^o \mid M_2)$	$\Pr(z > z^o \mid BV, MF, M_2)$
ρ_{BV}	0.02	0.61	0.01	0.40
ρ_{MF}	$2.2e^{-3}$	0.01	$6.0e^{-3}$	0.01
ρ_{RP}	0.01	0.36	0.01	0.33
Sk_{BV}	$2.0e^{-4}$	0.15	$1.2e^{-3}$	0.22
Sk_{MF}	$1.2e^{-3}$	0.95	$4.0e^{-3}$	0.84
Ku_{BV}	$4.0e^{-4}$	0.16	$1.2e^{-3}$	0.25
Ku_{MF}	$2.6e^{-3}$	0.95	$6.6e^{-3}$	0.87

vector of observed returns, $r_{1:T}$, and using the draws of $(V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v)$, a draw is produced from the *candidate* density,

$$q\left(\mu, \beta_v, \rho, \mu_p, \rho_J, \sigma_r, V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \phi\right) \propto p\left(V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \phi | MF_{1:T}, BV_{1:T}\right) \\ \times p\left(\mu, \beta_v, \rho, \mu_p, \rho_J, \sigma_r | r_{1:T}, V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v\right). \quad (42)$$

Under a standard Gaussian-inverse gamma $g-$prior framework (Zellner, 1986), the return-specific parameters, $\mu, \beta_v, \rho, \mu_p, \rho_J$, and σ_r, are sampled from the resulting Gaussian-inverse gamma posterior, conditional upon the draws of the latent variables, $V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}$, and the static parameters contained in ϕ, specifically θ, κ and σ_v. Given that, up to a constant of proportionality, the candidate density in (42) differs from the target density (41) only by the factor $p\left(r_{1:T} | V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v\right)$, a feasible sampling/importance resampling (SIR) algorithm (Smith and Gelfand, 1992) to correct the initial draws is implemented, so that a sample from the joint posterior distribution in (41) is obtained. Crucially, given the assumed structure of (39), including the assumption of normality for the innovations, and the use of a conjugate prior structure, $p\left(r_{1:T} | V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v\right)$ is known in closed form. Hence, given G available draws of $(V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v)$ from the MCMC algorithm, the draws are from the discrete distribution defined by the weights

$$w^{(i)} = \frac{p\left(r_{1:T} | V^{(i)}_{1:T}, \Delta N^{(i)}_{1:T}, Z^{v(i)}_{1:T}, \theta^{(i)}, \kappa^{(i)}, \sigma^{(i)}_v\right)}{\sum_{k=1}^{G} p\left(r_{1:T} | V^{(k)}_{1:T}, \Delta N^{(k)}_{1:T}, Z^{v(k)}_{1:T}, \theta^{(k)}, \kappa^{(k)}, \sigma^{(k)}_v\right)}, \; i = 1, 2, ..., G.$$

The R reweighted draws of $V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v$ are each used to produce a draw from $p\left(\mu, \beta_v, \rho, \mu_p, \rho_J, \sigma_r | r_{1:T}, V_{1:T}, \Delta N_{1:T}, Z^v_{1:T}, \theta, \kappa, \sigma_v\right)$, which together are then used to produce the conditional (Gaussian) predictive distribution for r_{T+1}. The (marginal) predictive, $p(r_{T+1}|r_{1:T}, MF_{1:T}, BV_{1:T})$ is produced by averaging the R conditional predictives in the usual way.[13]

Appendix D: Transformation to Risk Aversion The equilibrium frameworks of Breeden (1979) and Cox, Ingersoll and Ross (1985) lead to factor risk premiums that are equal to the negative of the covariance between changes in the factor and the rate of change in the marginal utility of wealth. For the latent variance model adopted in this paper, this implies that

$$(\lambda_D V_t - \delta_J \lambda_J)\, dt = -cov\left(\frac{dw_t}{w_t}, dV_t\right), \tag{43}$$

where $dV_t = dV_t^d + dJ_t^v$, with $dV_t^d = \kappa[\theta - V_t]dt + \sigma_v \sqrt{V_t} dB_t^v$ representing the diffusive component of the volatility process. Thus, the variance risk premium can be decomposed into the variance diffusive risk premium and the variance jump risk premium,

$$\lambda_D V_t dt = -cov\left(\frac{dw_t}{w_t}, dV_t^d\right) \text{ and } -\delta_J \lambda_J dt = -cov\left(\frac{dw_t}{w_t}, dJ_t^v\right), \tag{44}$$

respectively, where w_t denotes the marginal utility of wealth for the representative investor. Following Bollerslev, Gibson and Zhou (2011), for the purposes of this illustration we adopt the canonical power utility function,

$$U_t = e^{-\beta t}\frac{W_t^{1-\gamma}}{1-\gamma}, \qquad \gamma > 0,$$

where β denotes the constant subjective discount rate and γ is the risk aversion parameter, from which it follows that,

$$w_t = e^{-\beta t} W_t^{-\gamma}.$$

Proxying wealth by the value of the market portfolio and recognizing that, in our empirical setting, the price process $P_t = \exp(p_t)$ refers to a market stock index, where p_t is defined in (1), Ito's lemma yields

$$dw_t = -\gamma P_t^{-1} w_t dP_t + \gamma(1+\gamma) P_t^{-2} w_t (dP_t)^2,$$

that is,

$$\frac{dw_t}{w_t} = -\gamma P_t^{-1} dP_t + \gamma(1+\gamma) P_t^{-2} (dP_t)^2. \tag{45}$$

[13] The sample second moment of the normalised resampling weights, a useful indicator of importance sampling efficiency, must by construction lie between $\frac{1}{R}$ and 1. For our empirical implementation of the resampling procedure with $R = 5000$, we achieved values for the sample second moment around 0.61, indicating that our resampling algorithm is relatively efficient (see Liu, 2001, and Tokdar and Kass, 2010 for related discussion). The efficiency of such a resampling algorithm could, of course, potentially be improved, but such an exploration is beyond the scope of this paper.

Substituting (45) into (44) and maintaining the assumptions adopted in Section 5.5.2, yields

$$\lambda_D V_t dt = \gamma P_t^{-1} cov\left(dP_t, dV_t^d\right) = \gamma \rho \sigma_v V_t dt$$

and

$$\delta_J \lambda_J dt = -\gamma P_t^{-1} cov\left(dP_t, dJ_t^v\right) \simeq -\gamma \delta_J \rho_J \mu_v^2 dt, \tag{46}$$

which imply, in turn, that

$$\lambda_D V_t - \delta_J \lambda_J \simeq \gamma \rho \sigma_v V_t + \gamma \delta_J \rho_J \mu_v^2. \tag{47}$$

The approximation in (46), based on a first order Taylor series expansion of $\exp\left(Z_t^p\right)$, is valid for small but nonzero jump sizes. Finally, solving (47) for γ, we obtain

$$\gamma \simeq \frac{\lambda_D V_t - \delta_J \lambda_J}{\rho \sigma_v V_t + \rho_J \delta_J \mu_v^2}.$$

We are, of course, adopting a dynamic model for λ_J, which, combined with dynamics of the stochastic variance V_t, implies a dynamic model for γ.