Extracting S&P500 and NASDAQ volatility: The Credit Crisis of 2007-2008

Hedibert F. Lopes,[*]
&
Nicholas G. Polson[*]

University of Chicago

February 15, 2009

Chapter XVI of

The Handbook of Applied Bayesian Analysis

Eds: Tony O'Hagan & Mike West
Oxford University Press

Abstract

In this chapter we use particle filtering methods to estimate volatility and examine volatility dynamics for three financial time series during the early part of the current credit crisis. We compare estimates from a pure stochastic volatility model, a stochastic volatility model with jumps and a Garch model to each other and to the market volatilities implied by actual option prices. Our three time series are daily data for Standard and Poor's S&P500 index (or simply SP500), the Nasdaq NDX100 index (or simply NASDAQ) and the financial index XLF for 2007. We provide filtered volatility and parameter estimates for; pure stochastic volatility (SV), stochastic volatility jump (SVJ) and a Garch(1,1) model. Sequential model choice is a natural outcome of our application and we show how the evidence in support of the stochastic volatility jump model accumulates over time as market turbulence increases. Our findings have implications for derivative pricing, portfolio and risk management.

Keywords: Credit, Crisis, Particle Filtering, Learning, Stochastic Simulation, Posterior, Bayesian Inference, MCMC, Stochastic Volatility, Factor Models, Jumps.

[*]The University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL, 60637.

1 Introduction

Volatility and volatility dynamics are important for understanding the behavior of financial markets and market pricing. In this chapter we estimate volatility and volatility dynamics for daily data for the year 2007 for three market indicies: the Standard and Poor's S&P500, the Nasdaq NDX100 and the financial equity index called XLF. We study how the three series reflect the beginning of the credit crisis; the XLF index reflects the effect on financial companies that were among the earliest to be affected while the other two market indices reflect broader, economy wide effects. Three models of financial time series are estimated: a model with stochastic volatility, a model with stochastic volatility that also incorporates jumps in volatility and a Garch model. We compare our volatility estimates from these three models with subjective implied market volatility calculated from option prices, the VIX and VXN, for the S&P500, the Nasdaq NDX100, respectively. By sequentially computing marginal likelihoods or Bayes factor we can also evaluate the ability of the different models to capture the time series behavior over this turbulent period.

Volatility and volatility dynamics are central to many issues in financial markets including derivative prices, leverage ratios, credit spreads, and portfolio decisions. In times of low market volatility it is relatively straightforward to measure volatility and understand volatility dynamics. At other times, financial markets are affected by severe disruptions which may be largely isolated events like the market crash of 1987 or may be a series of events such as the Russian default and Long Term Capital Management Fund Crisis of 1998, or the current credit crisis. During such periods, apparent spikes in volatility and large movements in asset prices complicate estimation of volatility and volatility dynamics. This chapter describes how Bayesian particle filtering methods can track volatility in an online fashion under these circumstances. Our particle filtering methodology presented here provides sequential inference for financial time series in general and can also be applied to other problems such as longitudinal time series studies.

Sequential Bayesian methods based on particle filtering provide a natural solution to estimating volatility and volatility dynamics. Sequential inference is important as it provides estimates of current (spot) volatility and parameters for the evolution of the volatility dynamics given the currently available information. This includes the historical path of prices or returns available up until the current time. together with market beliefs about volatility dynamics implicit in option prices. These subjective market beliefs are measured by the so-called implied option volatilities and can be compared to model-based estimated volatility after accounting for the market price of volatility risk. We describe this further at the end of this section and in the appendicies.

Our primary filtering results are from two financial asset price models, one with stochastic volatility (SV) and one with stochastic volatility and jumps in volatility (SVJ). Jumps are transient by their nature and the resulting price and volatility dynamics can be very different depending upon whether jumps are included or not. In our estimation, we will identify the posterior distribution of the latent state variables at each point in time, denoted by $L_t = (V_t, J_t, Z_t)$ corresponding to the stochastic variance, jump time and the jump size. Our approach will also examine how quickly volatility estimation procedures react to changes in underlying asset returns.

We find a number of important empirical finance results. First, we compare our filtering results from the SV model and SVJ model with a Garch(1,1) model frequently applied to financial time series. We find that in periods of market turbulence that the Garch(1,1) model does not track option implied volatility as as well the two stochastic volatility models particularly for the more volatile NDX100 index. These facts also lead to differences in market prices of volatility risk between stochastic volatility models and Garch(1,1) throughout the period under investigation. Second, we examine the reaction of volatility estimates in periods of market stress and extreme movements can be surprisingly different even though SVJ is designed to capture jump effects. An attractive feature of the particle filtering approach is that we have a model diagnostic in the form of a sequential Bayes factor. We find that while the SVJ model outperforms the SV model for the entire period under investigation, the relative superiority of the SVJ model increases with market turbulence.

The methods used here build upon a number of recent papers that develop particle filtering algorithms for SV and SVJ models. Carvalho and Lopes (2007) develop sequential particle methods for Markov switching SV models and Carvalho et al. (2008) provide a general particle learning algorithm for SV models. Johannes et al. (2008) consider a general class of continuous-time asset price diffusion models and continuous-time stochastic volatility jump models. Particle filtering has an advantage over commonly used MCMC methods (Jacquier et al., 1994, 2004, Johannes and Polson,2004) which are computationally burdensome and inherently non-sequential.

The rest of the paper proceeds as follows. We begin by describing the applied problem and goal of the Bayesian analysis. Section 2 presents the estimation models. Section 3 discusses the simulation-based methodology for inference, including MCMC and particle filtering methods for filtering and parameter learning. Detailed analysis of our empirical findings appear in Section 4. Finally, Section 5 concludes.

1.1 Problem context and goals

Subjective beliefs about volatility are available from option prices in the form of implied volatilities. Model-based volatility estimates can be determined in an online fashion using historical return data and a specification for volatility dynmaics. Our goal then is to extract volatility estimates sequentially and compare them with market-based volatility measures. Our three underlying equity time series are daily data for the S&P500, the NDX100 and the XLF. We use two common market volatility indicies, the VIX and VXN indicies corresponding to the option implied volatilities of the benchmark indicies.

We begin by studying the volatility of the financial stock index, XLF. Figure 1 plots the XLF together with realized volatility estimates based on a simple rolling window of past returns. At the end of February, there are is a volatility spike, when the first credit crisis indicators in the form of sub-prime mortgage issues and collateral debt obligations (CDOs) came to light. The XLF stock index fell from 36 to 31 in the week at the end of February. The volatility spike on February 27th, 2007 occurred when the Dow Jones index dropped 320 points. This followed a fall of over 10% in the Chinese stock market. Our estimate of volatility, $\sqrt{V_t}$, of the XLF index moved from around 10% to nearly 30% very quickly in a matter of days.

In the next few months prices stabilized an volatility mean-reverted decaying back to the 20% range. The most dramatic change in volatility dynamics occurs just before the beginning of August and persists throughout the rest of the year. For example, the volatility spikes higher again in October back to the 40% range with the XLF dropping further from 36 at the beginning of October to 30 in December. We now describe theoretically and empirically the relationship between option prices and volatility.

1.2 Option prices and volatility

To study the relationship between option prices and volatility more closely we observe that a financial model describes asset price dynamics for the physical or objective measure $\mathbb{P}$ whereas derivatives are priced under $\mathbb{Q}$ the risk-neutral dynamics. These two different probability distributions underpin financial market pricing. In the next section we describe clearly the stochastic volatility dynamics describing these two distributions. Historical returns can be used to estimate the dynamics of the physical dynamics $\mathbb{P}$ and option returns can be used to assess the risk neutral dynamics $\mathbb{Q}$. Both returns and option prices will have information concerning the latent stochastic variance V_t. This is because, derivative prices depend crucially on the current volatility state, V_t, through a pricing formula

$$C\left(S_t, V_t\right) = e^{-r(T-t)} E_t^{\mathbb{Q}}\left[\max\left(S_T - K, 0\right) \mid S_t, V_t\right],$$

where S_t is the current value of the equity index and V_t is the current (spot) volatility. Appendix B provides explicit formulas for option pricing with stochastic volatility. This is related to the implied volatilities, VIX and VXN indices. The VIX index is computed from option prices and it also serves as a basis for volatility swaps which are a tradable asset.

We term these indices option market implied volatility denoted by IV_t. Formally, they are related to the expected future cumulative volatility via

$$IV_t = E_t^{\mathbb{Q}}\left[\int_t^{t+\tau} V_s ds\right].$$

If we include jumps, then

$$IV_t^J = E_t^{\mathbb{Q}}\left[\int_t^{t+\tau} V_s ds\right] + var_t^{\mathbb{Q}}\left[\sum_{n=N_t^s+1}^{N_{t+\tau}^s} Z_n^s\right],$$

where V_s denotes the path of volatility and N_t and Z_t denote the number of jumps and the jump sizes, respectively. As expected, the option implied volatility is providing a market-based prediction of future cumulative volatility.

Most approaches for estimating volatility and jumps rely exclusively on returns, ignoring the information embedded in option prices. In principle, options are a rich source of information regarding volatility, which explains the common use of Black-Scholes implied volatility as a volatility proxy for practitioners. In contrast to this is the common academic approach of only using returns to estimate volatility.

Options are highly informative about volatility, but the information content is dependent on the model specification and $\mathbb{P}$ and $\mathbb{Q}$-measure parameter estimates. Misspecification or poorly estimated risk premia could results in directionally biased estimates of V_t using option prices. This informativeness of index options is mitigated by the fact that these options contain relatively large bid-ask spreads, as noted by Bakshi et al. (1997) or **?**. This implies that while option prices are informative, they made be quite noisy as a measure of volatility. There is also empirical evidence for jumps in returns (Bates, 2000; Bates, 2006; Bakshi et al., 1997; and Eraker, 2004) as well as volatility (see Eraker et al., 2003).

Figure 2 compares the filtered XLF volatilities with the option implied volatilities of the broader indices. Not surprisingly, the volatilities are better tracked by the VIX than for the VXN. This is primarily due to the fact that financial companies comprised 25% of the SP500 index at the time. Even though the financial weighting in the Nasdaq is lower, it is surprising that the VXN volatility also tracks the movements in the XLF albeit not as well. This shows the influence of contagion effects of the credit crisis even to a sector of the market that has very little financial leverage. Implied volatility versus a 10-day moving average tracks very well, hence confirming the notion that implied volatility has information about future average volatility.

Over our data period the move in XLF prices was mirrored by the changes in the market option volatility indices (VIX and VXN). The sharp moves in market option implied volatility are hard to fully capture with a pure diffusive SV model. even though filtered volatility estimates move quicker than smooth estimates which incorporate future and past returns. Hence we will also study a stochastic volatility model that allows for jumps. However, we also here that these moves are even harder to capture with a deterministic time-varying volatility model like Garch (Rosenberg, 1972, Engle, 1982). We will see later than these type of return movements provides large statistical evidence in favor of the stochastic volatility jump model. For the next 3 months volatility mean reverts back to around the 10% level before the next shock hits at the end of July. In the next section we describe more clearly our methodology and comparison with the SP500 and NDX100 stock indices.

In principle, option prices also allow us to extract information about the current volatility state although this is not the focus of the study here. It is important to realise that in going from risk-neutral to physical measures we can assess the market price of volatility risk, see Polson and Stroud (2003), Eraker (2004) and Todorov (2007,2008) and Appendix B for further discussion. There is also related to a well-known volatility puzzle where option implied volatility is usually higher than estimated historical volatility. This is due to the fact that there exists a positive market price for taking volatility risk. Our particle methods will be able to measure this quantity.

2 Models

Our empirical analysis will focus on sequential volatility filtering in the 2007. Parameter estimates will be performed on a longer historical period of returns from 2002-2006. Our particle filtering approach will then be implemented for the year of 2007. We use daily return data for the Standard and Poor's SP500 stock index and the Nasdaq NDX 100 index, denoted here by SP500 and NASDAQ, respectively. We also study the behaviour of the XLF which is an equity index for the prices for US financial firms. The corresponding option volatility indices are the VIX and the VXN; again these are available on a continuous basis. We now describe the three models in question: the pure diffusive SV model, the SVJ model and the Garch(1,1) model. Firstly, we describe the three competing models that we use for price dynamics: pure diffusive SV model, SVJ model and Garch(1,1) model. Secondly, we describe the different volatility indices and how they are related to option prices and the market price of volatility risk.

2.1 Stochastic volatility (SV) model

A common model of asset price dynamics results in the following two equation describing the movements of an equity index S_t and its stochastic volatility V_t,

$$\begin{aligned} \frac{dS_t}{S_t} &= \mu + \sqrt{V_t} dB_t^{\mathbb{P}} \\ d\log V_t &= \kappa_v(\theta_v - \log V_t) + \sigma_v dB_t^V \end{aligned}$$

where μ is an expected rate of return and the parameters governing the volatility evolution are $\Theta = (\kappa_v, \theta_v, \sigma_v)$. The Brownian motions $(B_t^{\mathbb{P}}, B_t^V)$ are possibly correlated giving rise to a leverage effect.

This is a pure stochastic volatility (SV) model. The probabilistic evolution $\mathbb{P}$ describes what is known as the physical dynamics as opposed to the risk-neutral dynamics $\mathbb{Q}$ which is used for pricing. To analyze this model in a sequential fashion it is common to using an Euler discretization of the above model for continuously compounded returns (see details in Appendix A). In the subsequent analysis we will use daily time scale.

Let $Y_{t+1} = \log(S_{t+1}/S_t)$ be log-returns and transform volatility to logarithms as well, $X_{t+1} = \ln V_{t+1}$, then the Euler discretization is

$$\begin{aligned} Y_{t+1} &= e^{\frac{X_{t+1}}{2}} \varepsilon_{t+1} \\ X_{t+1} &= \alpha_v + \beta_v X_t + \sigma_v \eta_{t+1} \end{aligned}$$

Here ε_{t+1} and η_{t+1} are normally distributed, serially and contemporaneously independent shocks. The parameters that govern the evolution of volatility dynamics are transformed to $\alpha = \kappa_v\theta_v$ and $\beta_v = 1 - \kappa_v$.

The model specification is completed with independent prior distributions for the components of $\Theta = (X_0, \alpha_v, \beta_v, \sigma_v^2)$, i.e. $N(X_0, V_X)$, $N(\alpha_0, V_\alpha)$, $N(\beta_0, V_\beta)$ and $IG(\nu/2, \nu\bar{\sigma}_v^2/2)$, respectively, where N and IG denote the normal distribution and the inverse Gamma distribution. Table 3 provides posterior summaries for parameter estimates the stochastic volatility (SV) model for all three series, SP500, NASDAQ and XLF, for based on data from January 2002 to December 2006. Wee assumed relatively uninformative priors and set the hyperparameters in common choices from the literature, i.e. $V_X = 10$, $V_\alpha = V_\beta = 1$ and $\nu = 3$. The hyperparameters α_0, β_0 and $\bar{\sigma}_v^2$ are simple ordinary least square estimates based on a 10-day moving window procedure. Setting these hyperparameters at $(0, 0, 1)$, for instance, produced virtually the same posterior summaries form Table 3. The estimates (α_0, β_0) and $\bar{\sigma}_v^2$ are used as initial values for the MCMC algorithm. Other sound initial values also produced the same posterior results. Finally, the results are similar to the ones found in Eraker et al. (2003), with all three series exhibiting high persistent evolution.

2.2 Stochastic volatility jump (SVJ) model

The stochastic volatility jump (SVJ) model includes the possibility of jumps to asset prices. Now an equity index S_t and its stochastic variance V_t jointly solve

$$\begin{aligned} \frac{dS_t}{S_t} &= \mu + \sqrt{V_t} dB_t^{\mathbb{P}} + d\left(\sum_{s=N_t}^{N_{t+1}} Z_s\right) \\ d\log V_t &= \kappa_v(\theta_v - \log V_t) + \sigma_v dB_t^V \end{aligned}$$

where the additional term in the equity price evolution describes the jump process. Since data are observed in discrete time it is again common to use an Euler discretization of this continuous time process (Appendix A) . Specifically,

$$\begin{aligned} Y_{t+1} &= e^{\frac{X_{t+1}}{2}}\varepsilon_{t+1} + J_{t+1}Z_{t+1} \\ X_{t+1} &= \alpha_v + \beta_v X_t + \sigma_v \eta_{t+1} \\ J_{t+1} &\sim Ber(\lambda) \\ Z_{t+1} &\sim N(\mu_z, \sigma_z^2) \end{aligned}$$

with $Y_{t+1} = \log(S_{t+1}/S_t)$. The log-returns with ε_{t+1} and η_{t+1} normally distributed, serially and contemporaneously independent shocks. The parameter vector is $\Theta = (\lambda, \mu_z, \sigma_z, \alpha_v, \beta_v, \sigma_v)$.

Prior distributions are required for the initial volatility state, X_0, and for all parameters governing the dynamic of the volatilities, $(\alpha_v, \beta_v, \sigma_v^2)$. See Section 2.1 for explicit details. For the jump specification, we use a conditionally conjugate prior structure for parameters $(\lambda, \mu_z, \sigma_z^2)$, i.e. $\lambda \sim Beta(a, b)$, $\mu_z \sim N(c, d)$ and $\sigma_z^2 \sim IG(\nu/2, \nu\bar{\sigma}_z^2/2)$, respectively. We set $c = -3$ and $d = 0.01$ and $a = 2$ and $b = 100$ such that the prior mean and standard deviation of λ are around 0.02 and 0.014. ν and $\bar{\sigma}_z^2$ are set at 20 and 0.05, respectively, such that the prior mean and standard deviation of σ_z^2 are roughly 0.05 and 0.02. These prior specifications predict around five large negative jumps per year (roughly 250 business days) whose magnitude are around an additional three percent.

This structure naturally leads to conditional posterior distributions that can be easily simulated to form a Gibbs sampler (see Eraker et al.,2003). Table 2 provides posterior summaries for parameter estimates the stochastic volatility jump (SVJ) model for all three series, SP500, NASDAQ and XLF. Jump probability estimates are all similar at about 0.04 or 10 jumps per year. The largest estimates jump sizes are-3.72% for the XLF, -2.14% for the SP500 and -1.98% for the NASDAQ. The table also provides posterior means, standard deviations and 5% and 95% quantiles for $(\alpha_v, \beta_v, \sigma_v^2)$ and $(\lambda, \mu_z, \sigma_z^2)$.

2.3 Garch model

The Garch(1,1) model is a time-varying volatility model that uses The evolution of returns and volatility is then given by

$$\begin{aligned} Y_{t+1} &= \sqrt{V_{t+1}}\varepsilon_{t+1} \\ V_{t+1} &= \alpha_v + \beta_v V_t + \gamma_v \varepsilon_t^2 \end{aligned}$$

This leads to a time-varying volatility sequence given the residuals from the observation equation. The parameters have the usual constraints: $\alpha_v > 0$, $\beta_v > 0$ and $\gamma_v > 0$ to ensure positive variance, and $\alpha_v + \beta_v + \gamma_v < 1$ for stationarity. This model is fundamentally different from stochastic volatility, since last period's return shock ε_t^2 is used as a regressor as opposed to a stochastic volatility term $\sigma_v \eta_{t+1}$. Other differences include assessments of tail-probabilities and predictives. MCMC approaches for Bayesian inference on Garch models and several of its variants appear in Müller and Pole (1994), Bauwens and Lubrano (1998), Vontros and Politis (2000), Wago (2004) and Ausín and Galeano (2007) amongst many others.

Table 3 shows the parameters α_v, β_v and γ_v based on the 2002-2006 period. We then sequentially apply the model through the time period of 2007. Of particular interest is the comparison in August 2007 at the beginning of the credit crisis. We will also compare with the implied volatility series to see what the implications are for market prices of volatility risk. We find that the Garch volatility effect γ_v on the squared residual ε_t^2 is largest for the XLF at 0.057 and smallest for the NASDAQ at 0.03. This is not surprising as the credit crisis affected leveraged finance companies as opposed to technology stocks that traditionally has low debt levels.

3 Sequential learning via particle filtering

Given a time series of observations, $Y_{1:T} = (Y_1, \ldots, Y_T)$, the usual estimation problem is to estimate the parameters, Θ, and the unobserved states, $L_{1:T}$, from the observed data. In our case, the latent variables include (1) the volatility states, (2) the jump times, and (3) the jump sizes. In a Bayesian setting, this information is summarized by the joint posterior distribution $p(\Theta, L_{1:T}|Y_{1:T})$. In turn this joint posterior can be used to find estimates of the current state variables and parameters estimates using $\hat{L}_t = E(L_t|Y_{1:t})$ and $\hat{\Theta}_t = E(\Theta|Y_{1:t})$. Samples from this distribution are usually obtained via MCMC methods by iteratively sampling from the complete conditional distributions, $p(L_{1:T}|\Theta, Y_{1:T})$ and $p(\Theta|L_{1:T}, Y_{1:T})$. From these samples, it is straightforward to obtain smoothed estimates of the parameters and states.

Alternatively, particle filtering provides sequential estimates from the set of joint distributions $p(L_{1:t}|Y_{t:t})$. Particle filtering (PF, Gordon et al., 1993) methods are a simulation-based approach to sequential Bayesian filtering. Doucet et al. (2001) provide a reference text for a detailed discussion of the theoretical properties and applications. We also use a variant of the PF known as the auxiliary particle filter (APF, Pitt and Shephard, 1999).

The sequential estimation procedure is implemented in what follows, where Θ is assumed known to simplify the presentation. In our application, Θ is estimated based on daily data on SP500, NASDAQ and XLF from January 2002 to December 2006. Carvalho et al. (2008) provides more details about the following resample-propagate scheme including parameter learning.

3.1 Extracting state variables $L_t = (V_t, J_t, Z_t)$

The optimal filtering problem is solved by the sequential computation of the set of posteriors $p(V_t, J_t, Z_t|Y_{1:t})$. By Bayes rule these posteriors satisfy the recursion

$$p(V_{t+1}, J_{t+1}, Z_{t+1}|Y_{1:t+1}) \propto p(Y_{t+1}|V_{t+1}, J_{t+1}, Z_{t+1})p(V_{t+1}, J_{t+1}, Z_{t+1}|Y_{1:t})$$

where the normal likelihood $p(Y_{t+1}|V_{t+1}, J_{t+1}, Z_{t+1})$ has mean $J_{t+1}Z_{t+1}$ and variance V_{t+1}, while the prior $p(V_{t+1}, J_{t+1}, Z_{t+1}|Y_{1:t})$ is given by

$$p(V_{t+1}, J_{t+1}, Z_{t+1}|Y_{1:t}) \propto p(V_{t+1}|V_t)p(J_{t+1}, Z_{t+1})p(V_t, J_t, Z_t|Y_{1:t}).$$

The state evolution gives rise to a normal density $p(V_{t+1}|V_t)$ with mean $\alpha_v + \beta_v V_t$ and variance σ_v^2. Finally, as the jumps as transient and conditionally i.i.d. with $Pr(J_{t+1} = 1) = \lambda$ and normal density $p(Z_{t+1}|J_{t+1} = 1)$ with mean μ_z and variance σ_z^2. It would also be straightforward to allow for the jump probability to depend on the volatility state through a conditional $p(J_{t+1} = 1|V_t)$ (see Johannes et al., 1998). These equations will form the basis of our particle filtering algorithm that we develop in what follows. Following Carvalho et al. (2008) we now describe our resample-propagate filtering scheme.

3.2 Resample-propagation filter

Let the current filtering posterior be denoted by $p(L_t|Y_{1:t})$. The next likelihood is $p(Y_{t+1}|L_{t+1})$ and the state evolution is $p(L_{t+1}|L_t)$. Bayes rule links these to the next filtering distribution through Kalman updating. This takes the form of a smoothing step and a prediction step

$$\begin{aligned} p(L_t|Y_{1:t+1}) &\propto p(Y_{t+1}|L_t)p(L_t|Y_{1:t}) \\ p(L_{t+1}|Y_{1:t+1}) &= \int p(L_{t+1}|L_t)p(L_t|Y_{1:t+1})dL_t \end{aligned}$$

where $Y_{1:t}$ are the continuously compounded log-returns. Specifically, for extracting volatility and jumps as latent states we let $L_t = (V_t, J_t, Z_t)$. Uncertainty about these quantities is summarized via the filtered posterior distribution $p(L_t|Y_{1:t})$. Particle methods will represent this distribution as

$$p^N(L_t|Y_{1:t}) = \frac{1}{N}\sum_{i=1}^{N} \delta_{L_t^{(i)}}.$$

for particles $L_t^{(1)}, \ldots, L_t^{(N)}$. The previous filtered distribution is represented by its particle approximation and the key is how to propagate particle forward. From the updating formulas we can approximate

$$p^N(L_{t+1}|Y_{1:t+1}) = \sum_{i=1}^{N} w_t^{(i)} p(L_{t+1}|L_t^{(i)}, Y_{t+1})$$

where weights are given by

$$w_t^{(i)} = p\left(Y_{t+1}|L_t^{(i)}\right) \Big/ \sum\nolimits_{\ell=1}^{N} p\left(Y_{t+1}|L_t^{(\ell)}\right) \qquad i = 1, \ldots, N.$$

Hence after we have re-sampled the initial particles with weights proportional to $w_t^{(i)}$ we then propagate new particles using $p(L_{t+1}|L_t^{(i)}, Y_{t+1})$. This then leads us to the following simulation algorithm.

Resample-Propagate filter

1. *Resample:* For $i = 1, ..., N$, compute

$$w_t^{(i)} = p\left(Y_{t+1}|L_t^{(i)}\right) \Big/ \sum\nolimits_{\ell=1}^{N} p\left(y_{t+1}|L_t^{(\ell)}\right),$$

   draw

$$z(i) \sim Mult\left(N; w_t^{(1)}, \ldots, w_t^{(N)}\right),$$

   and set $L_t^{(i)} = L_t^{z(i)}$ for $i = 1, ...N$.

2. *Propagate:* For $i = 1, ..., N$, draw

$$L_{t+1}^{(i)} \sim p\left(L_{t+1}|L_t^{(i)}, Y_{t+1}\right).$$

For the SV and SVJ models, it is common to include V_{t+1} into the definition of the latent state variable L_t. This is due to the nonlinearities that the volatility induces into the asset price dynamics. In the SVJ model the state variables describing the jump process (J_t, Z_t) can be marginalised out. This is due to the fact that they are independent of V_t and we have the integral decomposition

$$p(Y_{t+1}|V_{t+1}) = \int p(Y_{t+1}|V_{t+1}, J_{t+1}, Z_{t+1}) p(J_{t+1}, Z_{t+1}) dJ_{t+1} dZ_{t+1}.$$

Therefore, in practice, we propagate forward the volatility $V_{t+1} \sim p(V_{t+1}|V_t)$ and then attach it to the current particle before using the Resample-Propagate filter described above.

3.3 Sequential model choice

An important by-product of sequential Monte Carlo methods is the ability to easily compute approximate marginal predictive densities and then Bayes factors. Let $\mathcal{M}$ denote a given model, then the sequential marginal predictive for any t can be approximated via

$$p(Y_{t+1}|Y_{1:t}, \mathcal{M}) = \frac{1}{N} \sum_{i=1}^{N} p(Y_{t+1}|L_t^{(i)}, \mathcal{M})$$

where $L_t^{(i)}$ is the particle for the latent volatility state of model $\mathcal{M}$. This approximation allows one to sequentially compute the Bayes factors (West, 1986), namely $\mathcal{BF}_{1:t}$, for competing models $\mathcal{M}_0$ and $\mathcal{M}_1$

$$\mathcal{BF}_{1:t} = \frac{p(Y_{1:t}|\mathcal{M}_1)}{p(Y_{1:t}|\mathcal{M}_0)}.$$

The Bayes factor is related to the Bayesian posterior probabilities of the models being true via the following identity

$$\frac{p(\mathcal{M}_1|Y_{1:t})}{p(\mathcal{M}_0|Y_{1:t})} = \mathcal{BF}_{1:t} \times \frac{p(\mathcal{M}_1)}{p(\mathcal{M}_0)}$$

where $p(\mathcal{M}_1)/p(\mathcal{M}_0)$ is typically set equal to one to denote a priori equal weight on either model. One advantage of this approach is that we can interpret the relative posterior model probabilities whether or not the "true" data generating process is either on of the models under consideration.

A sequential decomposition of the joint distribution is also available and we write

$$p(Y_{1:t}|\mathcal{M}) = p(Y_1|\mathcal{M}) \prod_{j=1}^{t-1} p(Y_{j+1}|Y_{1:j}, \mathcal{M})$$

Now we can use the particle approximation in an on-line fashion to compute the quantities $p(Y_{j+1}|Y_{1:j}, \mathcal{M})$. Hence we can sequentially compute the Bayes factor for the full data sequence by compounding

$$\mathcal{BF}_{1:T} = \prod_{j=1}^{T-1} \mathcal{BF}_{j:j+1}.$$

When competing models (and priors) are nested, which is the case in the SV versus SVJ case, for example, a further simplification is the implementation of the Savage-Dickey density ratio (Verdinelli and Wasserman, 1995). More specifically, let $\Theta = (\lambda, \mu_z, \sigma_z^{-2})$, ie. the parameters of the SVJ model (the full model or $\mathcal{M}_1$) that are not in the SV model (the reduced model or $\mathcal{M}_0$), then

$$\mathcal{BF}_{1:t} = \frac{p(\Theta = \Theta_0|Y_{1:t}, \mathcal{M}_1)}{p(\Theta = \Theta_0|\mathcal{M}_1)} \approx \frac{1}{N} \sum_{i=1}^{N} \frac{p(\Theta = \Theta_0|Y_{1:t}, L_{1:t}^{(i)}, \mathcal{M}_1)}{p(\Theta = \Theta_0|\mathcal{M}_1)}$$

where $\Theta_0 = (0, 0, 0)$. The computation of $p(\Theta = \Theta_0|Y_{1:t}, L_{1:t}^{(i)}, \mathcal{M}_1)$ is the main challenge since

$$p(\Theta = \Theta_0|Y_{1:t}, L_{1:t}^{(i)}, \mathcal{M}_1) \propto p(Y_{1:t}|L_{1:t}^{(i)}, \Theta = \Theta_0, \mathcal{M}_1) p(\Theta = \Theta_0|L_{1:t}^{(i)}, \mathcal{M}_1).$$

Hence this provides a computational tractable approach for calculating Bayes factors from particle filtering output.

Figure 6 provides the sequential log-Bayes factor $\mathcal{BF}(SV, SVJ)$ for comparing the pure SV model with an SVJ model. We provide the Bayes factor diagnostic for all three series, SP500, NASDAQ and XLF. Not surprisingly, the large negative shock in February 2007 leads to substantial evidence in favor of SVJ over SV. However, for the rest of the year as there are no extreme shocks to returns the evidence decays back and at the end of the period slightly favors the pure SV model. Providing these estimates within a pure MCMC framework would be computationally expensive, see, for example, the discussion in Chapter 7 of Gamerman and Lopes (2006).

4 Empirical results

We implement our particle filtering methodology to find volatility estimates for all the models on a daily basis in 2007. Special focus is on the period at the onset of the credit crisis, namely August 2007. We compare volatility estimates and implied prices of volatility risk for each model. Table 4 provides this comparison for the month of August. The volatilities are filtered and a direct comparison with the VIX and VXN for each model gives a measurement of the price of volatility risk for each model. The SVJ model tracks the movements more closely than the pure SV model. The SVJ model implies a relatively constant λ_v until the end of the month when all model imply a very small λ_v. It is also interesting to see how the sequential Bayes factor discriminates between these models. see Figure 6 and the subsequent discussion.

Periods of high volatility risk premia occur at the end of July for both SV and SVJ models while lower premia occur for the Garch(1,1) model. A high premia is empirically justified by subsequent increases and volatility spikes in the August period. Remember that one can interpret the risk premia are market expectations of future changes in average volatility.

This difference can be explained as follows. In the last week of July the Dow Jones index dropped from $14,000$ to $13,000$ with a sequence of negative shocks. The Garch(1,1) model therefore estimates volatility at 23% on August 1st as negative shocks feed directly into the Garch(1,1) model via $\gamma\varepsilon_t^2$. Both SV and SVJ models attribute some of theirs negative shocks to the stochastic volatility error term of the jump component rather than directly to $\sqrt{V_t}$. To initialize our estimate of volatility we use 17% for the SV model.

Our comparisons are summarized in more detail in Figures 3 to 5. They provide sequential comparisons of volatility estimates relative to the VIX and VXN indices. Specifically, Figure 3 compares the underlying indices SP500 and NASDAQ with their respective implied volatility series. Figure 4 compares Garch(1,1) with SV and SVJ against the option implied series VIX and VXN.

Figure 5 provides a direct comparison between the models at the beginning of the credit crisis in August 2007. For reasons described before, the Garch(1,1) model does not react as quickly as a stochastic volatility model. It also starts the month of August at higher estimates of V_t or equivalently lower estimates of the volatility risk premia. The filtered Garch(1,1) estimates are also considerably smoother than the SV and SVJ estimates, this is because the SV model have the extra flexibility in the random variance term to adapt to large shocks.

Figure 5 also compares the three models with their option implied volatilities. This allows the researcher to gauge the movement in the market price of volatility risk, λ_v. It is also useful to look at the sequential bayes factors as a measurement of the market price of volatility risk is conditional on the model. Within the SV model, the average level of risk neutral volatility is $\theta^\star = \kappa\theta/(\kappa + \lambda_v)$. See Duffie et al. (2000) and Pan (2002) for a discuss with the SVJ model. By the end of August the initial shock has dissipated and all three models imply that there is little price of volatility risk.

Figure 6 shows that just before August 2007 the Bayes factor favors the stochastic volatility jump model. Hence we would expect the volatility estimates from the SVJ model to more closely track the option implied volatility VXN. This is indeed the case. The market volatility risk premia is effectively constant for this model over this data period, except at the very end of the period where the implied option volatility decay quickly and the estimated volatility does not. This is also coincident with the Bayes factor decaying back in favor of the pure SV model for the NASDAQ index. By the end of 2007, the odds favor the pure SV model over the SVJ model for the NASDAQ index.

For the XLF, most of the evidence for jumps is again contained in the February move. The sequential Bayes factor tends to lie in between the strong evidence for the SP500 and weaker evidence for the NASDAQ index. The story for the SP500 is different. Figure 6 shows that after the February shock, the SVJ model is preferred to the SV model for the whole period. When comparing with VIX the jump model seems to track the option implied volatility with an appropriate market price of volatility risk.

5 Conclusions

In this chapter we proposed and implemented sequential particle filtering methods to study the US credit crisis which began in 2007. The purpose of our study was to show how model-based volatility estiamtes can be compared to model-based ones in a dynamic setting. Online daily volatility for S&P500, NDX100 and XLF are estimated for SV, SVJ and Garch(1,1) models. Market-based volatilities are based on the implied volatility indicies for the S&P500 and NDX100 indices, namely the VIX and VXN. We show that particle filtering methods naturally allow for online volatility, jump and sequential model comparison.

We find a number of empirical results. First, tracking volatility in turbulent periods is much harder than low volatility periods. The inclusion of the possibility of jumps can change current estimates of volatility dramatically. The pure stochastic volatility (SV) and Garch(1,1) models perform significantly worse in periods of market stress, both in terms of tracking subjective market-based volatility and relative marginal likelihoods versus the SVJ model. Second, by calculating sequential marginal likelihoods for our data period, we see that the stochastic volatility jump model is clearly preferred. Not surprisingly, this evidence accumulates mainly on a few days where the stock returns are most extreme.

Extensions include multivariate modeling of the return series possibly with the use of a factor stochastic volatility models as proposed by Aguilar and West (2000), Lopes (2000) and Lopes and Migon (2002). Lopes and Carvalho (2007) show how to perform sequential inference for this class of models. More flexible volatility dynamics can also be modeled. Recently, Carvalho et al. (2009) introduce a class of affine *shot-noise* continuous-time stochastic volatility model that accounts for both fast moving, rapidly mean-reverting shot-noise volatility and slow-moving square-root diffusive volatility.

Acknowledgements

We are grateful to the Editors.

References

Aguilar, O. and M. West (2000). Bayesian dynamic factor models and variance matrix discounting for portfolio allocation. *Journal of Business and Economic Statistics 18*, 338–57.

Ausín, M. C. and P. Galeano (2007). Bayesian estimation of the Gaussian mixture Garch model. *Computational Statistics & Data Analysis 51*, 2636–52.

Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *Journal of Finance 52*, 2003–49.

Bates, D. S. (2000). Post-'87 crash fears in S&P500 futures options. *Journal of Econometrics 94*, 181–238.

Bates, D. S. (2006). Maximum likelihood estimation of latent affine processes. *Review of Financial Studies 19*, 909–65.

Bauwens, L. and M. Lubrano (1998). Bayesian inference on Garch models using the Gibbs sampler. *Econometrics Journal 1*, C23–C46.

Black, F. (1976). Studies of stock price volatility changes. In *1976 Meetings of the Business and Economics Statistics Section*, pp. 177–81. Proceedings of the American Statistical Association.

Carvalho, C., M. Johannes, H. F. Lopes, and N. Polson (2008). Particle learning and smoothing. Working paper, University of Chicago Booth School of Business.

Carvalho, C., M. Johannes, H. F. Lopes, and N. Polson (2009). Stochastic volatility shot noise. Working paper, University of Chicago Booth School of Business.

Carvalho, C. and H. F. Lopes (2007). Simulation-based sequential analysis of markov switching stochastic volatility models. *Computational Statistics & Data Analysis 51*, 4526–42.

Doucet, A., N. de Freitas, and N. Gordon (2001). *Sequential Monte Carlo Methods in Practice*. New York: Springer.

Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica 68*, 1343–76.

Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of u.k. inflation. *Econometrica 50*, 987–1008.

Eraker, B. (2004). Do equity prices and volatility jump? reconciling evidence from spot and option prices. *Journal of Finance 59*, 1367–1403.

Eraker, B., M. Johannes, and P. Polson (2003). The impact of jumps in volatility and returns. *Journal of Finance 59*, 227–60.

Gamerman, D. and H. F. Lopes (2006). *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. Baton Rouge: Chapman & Hall/CRC.

Gordon, N., D. Salmond, and A. F. M. Smith (1993). Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings F-140*, 107–13.

Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*, 327–43.

Jacquier, E., N. G. Polson, and R. P. E. (1994). Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics 20*, 69–87.

Jacquier, E., N. G. Polson, and R. P. E. (2004). Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. *Journal of Econometrics 122*, 185–212.

Johannes, M., R. Kumar, and N. G. Polson (1998). State dependent jump models: How do US equity indices jump? Working paper, University of Chicago Booth School of Business.

Johannes, M. and N. Polson (2004). MCMC methods for financial econometrics. Princeton, NJ. Princeton University Press.

Johannes, M., N. G. Polson, and J. Stroud (2008). Optimal filtering of jump diffusions: extracting latent states from asset prices. *Review of Financial Studies (to appear)*.

Lopes, H. F. (2000). *Bayesian analysis in latent factor and longitudinal models*. Ph. D. thesis, Institute of Statistics and Decision Sciences, Duke University.

Lopes, H. F. and C. M. Carvalho (2007). Factor stochastic volatility with time varying loadings and Markov switching regimes. *Journal of Statistical Planning and Inference 137*, 3082–91.

Lopes, H. F. and H. S. Migon (2002). *Case Studies in Bayesian Statistics*, Volume VI of *Lecture Notes in Statistics 167,* Chapter Comovements and contagion in emergent markets: stock indexes volatilities, pp. 285–300. New York: Springer.

Müller, P. and A. Pole (1994). Monte carlo posterior integration in Garch models. *Sankhya, Series B 60*, 127–44.

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

Pitt, M. and N. Shephard (1999). Filtering via simulation: auxiliary particle filter. *Journal of the American Statistical Association 94*, 590–99.

Polson, N. G. and J. Stroud (2003). Bayesian inference for derivative prices. *Bayesian Statistics 7*, pp. 641–50. Oxford: Clarendon Press.

Rosenberg, B. (1972). The behaviour of random variables with nonstationary variance and the distribution of security prices. Working Paper 11, Research Programme in Finance, Graduate School of Business Administration, Institute of Business and Economic Research, University of California, Berkeley.

Todorov, V. (2007). Variance risk premium dynamics. Working paper, Department of Economics, Duke University, Durham.

Verdinelli, I. and L. Wasserman (1995). Computing Bayes factor using a generalization of the Savage-Dickey density ratio. *Journal of the American Statistical Association 90*, 614–8.

Vontros, I., D. P. and D. I. Politis (2000). Full Bayesian inference for Garch and EGarch models. *Journal of Business & Economic Statistics 18*, 187–98.

Wago, H. (2004). Bayesian estimation of smooth transition Garch model using Gibbs sampling. *Mathematics and Computers in Simulation 44*, 63–78.

West, M. (1986). Bayesian model monitoring. *Journal of the Royal Statistical Society (Series B) 48*, 70–8.

Appendix A: Time discretization and DLMs

A major advance in financial modelling was the development of continuous-time models for the evolution of stock prices and their volatility. The simplest model is a geometric Brownian motion Black-Scholes model where stock prices S_t solves

$$dS_t = S_t(\mu dt + \sigma dB_t)$$

for an instantaneous expected return μ and volatility σ. Here B_t is a standard Brownian motion with distribution $B_t \sim \mathcal{N}(0, t)$. A natural extension is the inclusion of stochastic volatility (SV). Transforming to a log-scale and replacing the volatility by $\sqrt{V_t}$ gives a model of the form

$$d \ln S_t = \mu dt + \sqrt{V_t} dB_t$$

where typically V_t solves its own stochastic evolution such as a square-root Ornstein-Uhlenbeck process.

In practice, we observe data on a discrete time scale and it is common to use a time discretisation of these models. The most common discretisation is an Euler scheme. Specifically, on a time interval $(t, t+\Delta)$ we have the evolution

$$\ln S_{t+\Delta} - \ln S_{t+\Delta} = \mu\Delta + \sqrt{V_t}\sqrt{\Delta}\epsilon_t$$

where $\epsilon_t \sim \mathcal{N}(0, 1)$. With loss of generality we take $\Delta = 1$. Moreover, writing the log-returns $R_{t+1} = \ln(S_{t+1}/S_t)$ we we now have a dynamic linear models (DLM) with a hidden state that can be filtered from the partially observed data.

Appendix B: Informational content of returns and option prices

Option prices provide information about state variables and parameters via the pricing equation. An option price is given by

$$C\left(S_t, V_t, \theta\right) = e^{-r(T-t)} \mathbb{E}_t^{\mathbb{Q}} \left[\max\left(S_t - K, 0\right) | S_t, X_t, \Theta \right]$$

where the expectation is taken under $\mathbb{Q}$ the risk-neutral probability measure. here we assume that the risk-free rate r is constant. This simplifies by letting $A = \{S_T \geq K\}$ denote the event that the stock ends in the money. Then the option price is given by

$$\begin{aligned} e^{-rt}\mathbb{E}_t^{\mathbb{Q}}\left(\max\left(S_T - K, 0\right)\right) &= e^{-rt}\mathbb{E}_t^{\mathbb{Q}}\left(S_T \mathbb{I}_A\right) - e^{-rt} K \mathbb{E}\left(\mathbb{I}_A\right) \\ &= e^{-rt}\mathbb{E}_t^{\mathbb{Q}}\left(S_T \mathbb{I}_A\right) - e^{-rt} K \mathbb{P}(A) \end{aligned}$$

A common approach is to use a leverage stochastic volatility model (Heston, 1993). Here the underlying equity price S_t evolves according to a stochastic volatility model with square-root dynamics for variance V_t, and correlated errors. The logarithmic asset price and volatility follow an affine process with $Y_t = (\log S_t, V_t) \in \Re \times \Re^+$ satisfying

$$\begin{aligned} dY_t &= \left(r - \frac{1}{2}V_t\right) dt + \sqrt{V_t} dB_{1,t} \\ dV_t &= \kappa(\theta - V_t)dt + \sigma_v \sqrt{V_t} dZ_t \end{aligned}$$

where $Z_t = \rho B_{1,t} + \sqrt{1-\rho^2} B_{2,t}$. The correlation, ρ, or so-called leverage effect (Black, 1976) is important to explain the empirical fact that volatility increases faster as equity prices drop. The parameters (κ, θ) govern the speed of mean reversion and the long-run mean of volatility and σ_v measures the volatility of volatility. Under the risk-neutral measure they become $(\kappa/\kappa + \lambda_v)\theta$ and $\kappa + \lambda_v$ where λ_v is the market price of volatility risk.

From affine process theory, the discounted transform

$$\psi(u) = e^{-rt}\mathbb{E}\left(e^{uY_T} | Y_t\right) = e^{\alpha(t,u) + uY_t + \beta(t,u)V_t}$$

where α, β satisfy Riccati equations. We can then use transform analysis (Duffie et al., 2000) and compute prices by inverting a fast fourier transform. Specifically, there exists a pair of probabilities $\mathbb{P}_j(V_t, \Theta, \Lambda), j = 1, 2$, such that the call price $C(Y_t^S, V_t, \Theta, \Lambda)$ is given by:

$$C(Y_t^S, V_t, \Theta, \Lambda) \;=\; S_t \mathbb{P}_1(V_t, \Theta, \Lambda) - K e^{-r\tau} \mathbb{P}_2(V_t, \Theta, \Lambda)$$

Specifically, $\mathbb{P}_j(V_t, \Theta, \Lambda) \;=\; Pr_j\left(\ln(S_T/K)|V_t, \Theta, \Lambda\right)$, where these probabilities can be determined by inverting a characteristic function

$$\mathbb{P}_j(V_t, \Theta, \Lambda) \;=\; \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re\left[\frac{e^{-i\phi \ln(K)} f_j(V_t, \Theta, \Lambda)}{i\phi}\right] d\phi$$

where

$$f_j(V_t, \Theta, \Lambda) \;=\; e^{C(\tau;\phi) + D(\tau;\phi)V_t + i\phi \ln(S_t)},$$

and the coefficients $C(\tau; \phi)$ and $D(\tau; \phi)$ are defined in Heston (1993).

$$C(\tau;\phi) \;=\; r\phi i\tau + \frac{a}{\sigma^2}\left\{(b_j - \rho\sigma\phi i + d)\tau - 2\log\left[\frac{1 - g e^{d\tau}}{1 - g}\right]\right\},$$

$$D(\tau;\phi) \;=\; \frac{b_j - \rho\sigma\phi i + d}{\sigma^2}\left[\frac{1 - e^{d\tau}}{1 - g e^{d\tau}}\right],$$

with parameters

$$g \;=\; \frac{b_j - \rho\sigma\phi i + d}{b_j - \rho\sigma\phi i - d},$$

$$d \;=\; \sqrt{(\rho\sigma\phi i - b_j)^2 - \sigma^2(2u_j\phi i - \phi^2)}.$$

The correlation and volatility of volatility parameters affect the skewness and kurtosis of the underlying return distribution and also how the correlation ρ affects the probabilities $\mathbb{P}_j(V_t, \Theta, \Lambda)$. For example, as is typically the case, when the correlation is negative then the underlying risk neutral distribution has a skewed left tail which in turn decreases the price of out-of-the-money call options while increasing the out-of-the-money put options.

Tables and figures

SP500	Mean	StDev	2.5%	97.5%
α_v	-0.0031	0.0029	-0.0092	0.0022
β_v	0.9949	0.0036	0.9868	1.0011
σ_v^2	0.0076	0.0026	0.0041	0.0144

NASDAQ	Mean	StDev	2.5%	97.5%
α_v	-0.0003	0.0023	-0.0046	0.0045
β_v	0.9968	0.0024	0.9914	1.0011
σ_v^2	0.0046	0.0015	0.0024	0.0084

XLF	Mean	StDev	2.5%	97.5%
α_v	-0.0020	0.0032	-0.0082	0.0040
β_v	0.9924	0.0042	0.9830	0.9996
σ_v^2	0.0115	0.0036	0.0064	0.0203

Table 1: *Stochastic volatility (SV) model.* Mean and StDev are posterior mean and standard deviation, respectively. 2.5% and 97.5% are posterior percentiles. Time span: 1/02/2002 - 12/29/2006.

SP500	Mean	StDev	2.5%	97.5%
α_v	-0.0117	0.0070	-0.0262	0.0014
β_v	0.9730	0.0084	0.9551	0.9886
σ_v^2	0.0432	0.0082	0.0302	0.0613
λ	0.0025	0.0017	0.0003	0.0066
μ_z	-2.7254	0.1025	-2.9273	-2.5230
σ_z^2	0.3809	0.2211	0.1445	0.9381

NASDAQ	Mean	StDev	2.5%	97.5%
α_v	0.0079	0.0065	-0.0042	0.0215
β_v	0.9785	0.0073	0.9631	0.9916
σ_v^2	0.0390	0.0071	0.0275	0.0553
λ	0.0031	0.0021	0.0004	0.0082
μ_z	-3.9033	0.1001	-4.0906	-3.7054
σ_z^2	0.6420	0.3856	0.2445	1.6314

XLF	Mean	StDev	2.5%	97.5%
α_v	-0.0044	0.0064	-0.0176	0.0083
β_v	0.9728	0.0085	0.9554	0.9880
σ_v^2	0.0472	0.0091	0.0324	0.0677
λ	0.0026	0.0018	0.0003	0.0071
μ_z	-3.2676	0.0997	-3.4593	-3.0700
σ_z^2	0.8983	0.5490	0.3467	2.2171

Table 2: *Stochastic volatility with jump (SVJ) model.* Mean and StDev are posterior mean and standard deviation, respectively. 2.5% and 97.5% are posterior percentiles. Time span: 1/02/2002 - 12/29/2006.

Index	α_v	β_v	γ_v	$\alpha_v + \beta_v + \gamma_v$
SP500	0.0042	0.9440	0.0501	0.9983
NADSAQ	0.0035	0.9652	0.0319	1.0006
XLF	0.0072	0.9354	0.0573	1.0000

Table 3: *Garch model.* Parameter estimates. Time span: 1/02/2002 - 12/29/2006.

	SP&500			NASDAQ			XLF				
Day	SVJ	SV	Garch	SVJ	SV	Garch	SVJ	SV	Garch	VIX	VXN
1	18.1	16.5	17.4	18.6	14.9	16.5	25.9	22.6	21.3	23.8	23.6
2	17.0	16.1	17.0	18.3	14.9	16.4	24.2	22.0	20.9	25.1	25.4
3	21.1	18.3	16.5	20.2	15.8	16.2	31.0	26.3	20.3	23.4	24.5
6	23.2	19.5	18.6	20.1	15.8	16.9	37.9	29.9	23.5	22.8	24.5
7	21.9	19.1	19.5	19.3	15.7	16.8	35.4	29.3	26.4	24.6	27.1
8	21.7	19.2	18.9	19.1	15.7	16.6	35.0	29.6	25.6	24.0	26.6
9	25.3	21.1	18.7	22.0	17.0	16.5	37.6	31.4	25.4	26.2	29.2
10	23.6	20.7	20.4	21.3	16.9	17.5	34.7	30.6	26.1	27.4	29.9
13	21.8	20.2	19.7	20.4	16.7	17.3	32.2	29.8	25.3	25.3	27.5
14	22.8	20.7	19.0	21.4	17.2	17.1	31.3	29.5	24.5	25.0	28.0
15	22.7	20.7	19.2	22.7	17.8	17.4	29.1	28.8	24.1	24.8	27.4
16	21.1	20.2	19.0	22.2	17.8	17.8	31.9	30.0	23.4	24.9	27.0
17	23.7	21.3	18.4	23.2	18.3	17.7	36.5	32.6	24.1	26.5	28.7
20	22.0	20.8	19.4	21.8	18.1	18.0	34.4	31.9	25.5	20.4	23.3
21	20.4	20.3	18.7	21.3	18.0	17.6	31.7	30.8	24.9	20.0	22.5
22	19.8	20.3	18.1	21.3	18.1	17.5	29.0	29.7	24.2	20.4	22.0
23	18.4	19.8	17.9	20.1	17.9	17.4	26.5	28.5	23.4	19.0	20.7
24	18.3	19.6	17.3	20.4	18.0	17.1	24.0	27.4	22.6	19.4	20.8
27	17.7	19.3	17.1	19.5	17.8	17.2	23.3	26.8	21.9	18.6	20.8
28	21.1	20.6	16.8	22.8	18.9	16.9	28.1	28.8	21.5	17.6	20.6
29	22.7	21.5	18.3	25.9	20.4	17.9	26.8	28.2	22.4	17.0	20.6
30	21.2	21.0	18.9	24.4	20.0	18.9	24.6	27.0	22.0	18.0	21.0
31	20.7	20.7	18.3	23.8	19.9	18.5	24.0	26.5	21.4	17.8	21.1

Table 4: SVJ, SV and Garch comparison. Columns 2 to 10 are annualized standard deviations.

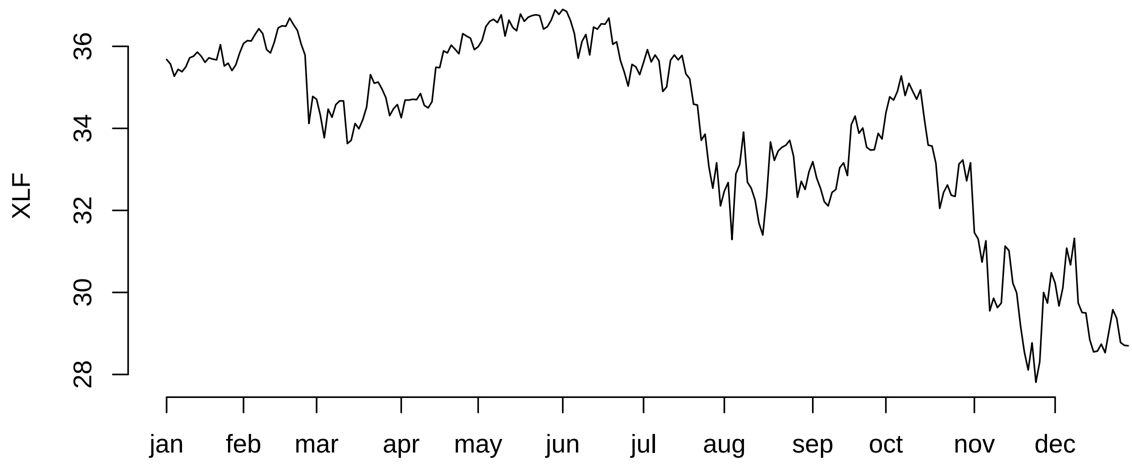

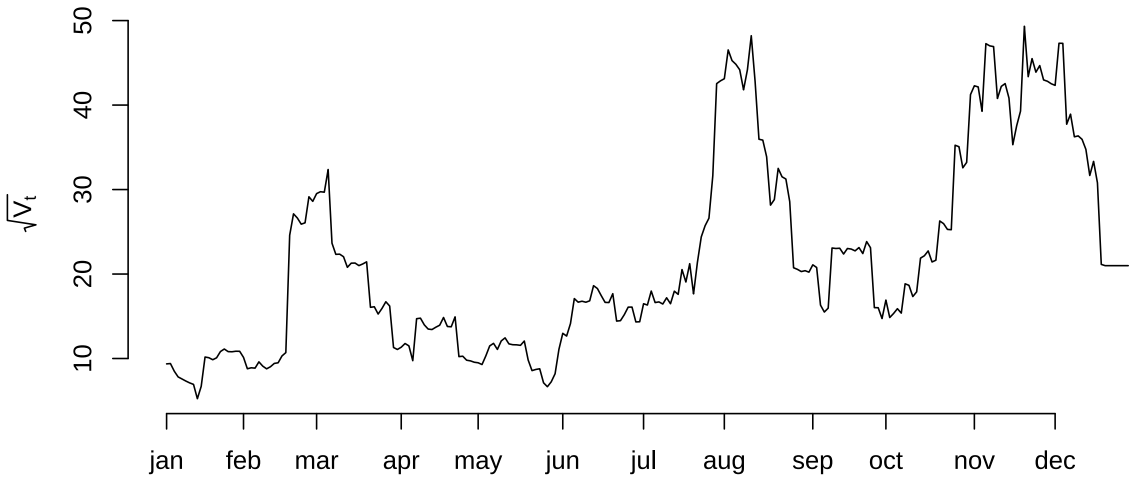

Figure 1: *Year 2007:* XLF underlying price (top) and rolling window annualized standard deviations (bottom). Volatility increases just before August 2007.

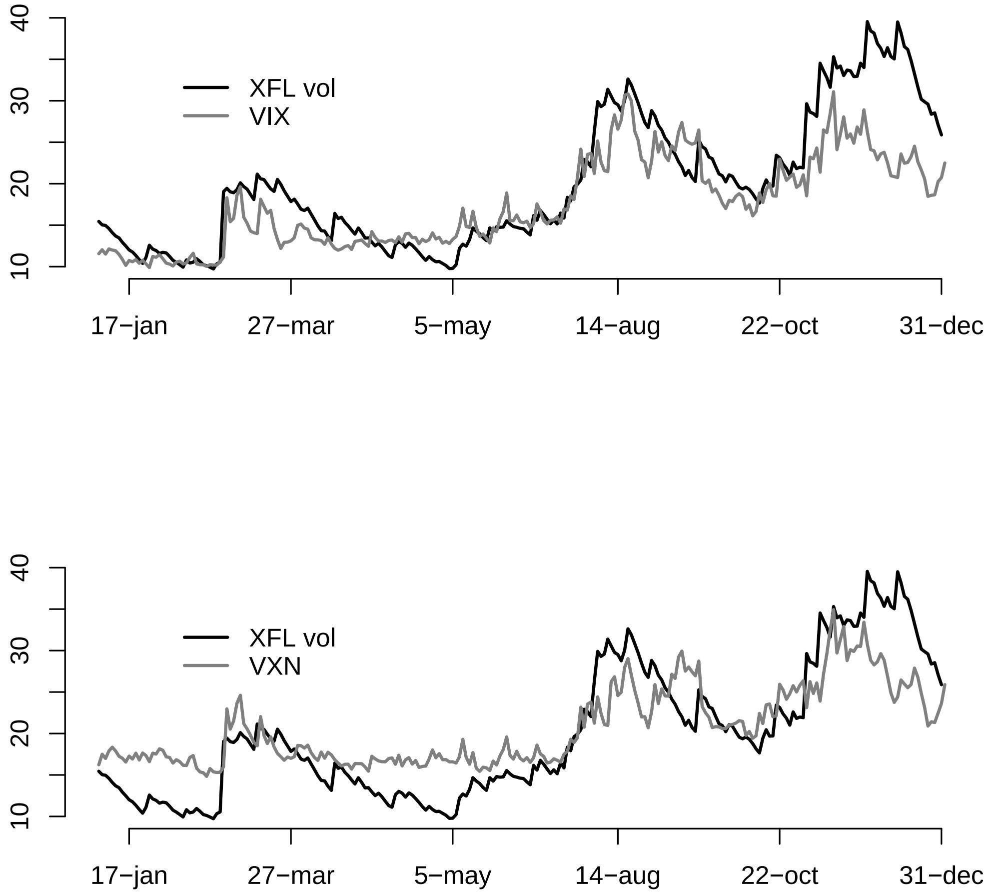

Figure 2: *Year 2007:* XLF volatility versus VIX and VXN.

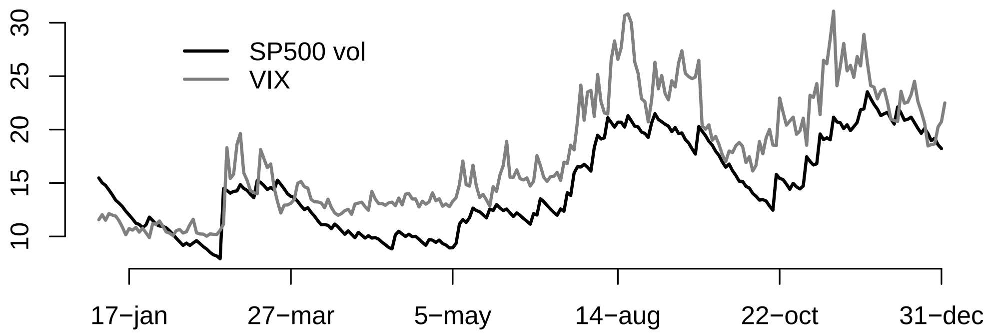

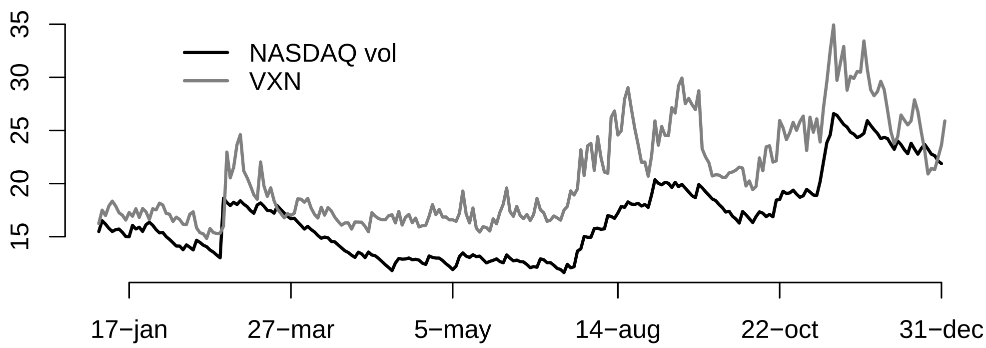

Figure 3: *Year 2007:* SP500 and NASDAQ volatilities and VIX and VXN indices.

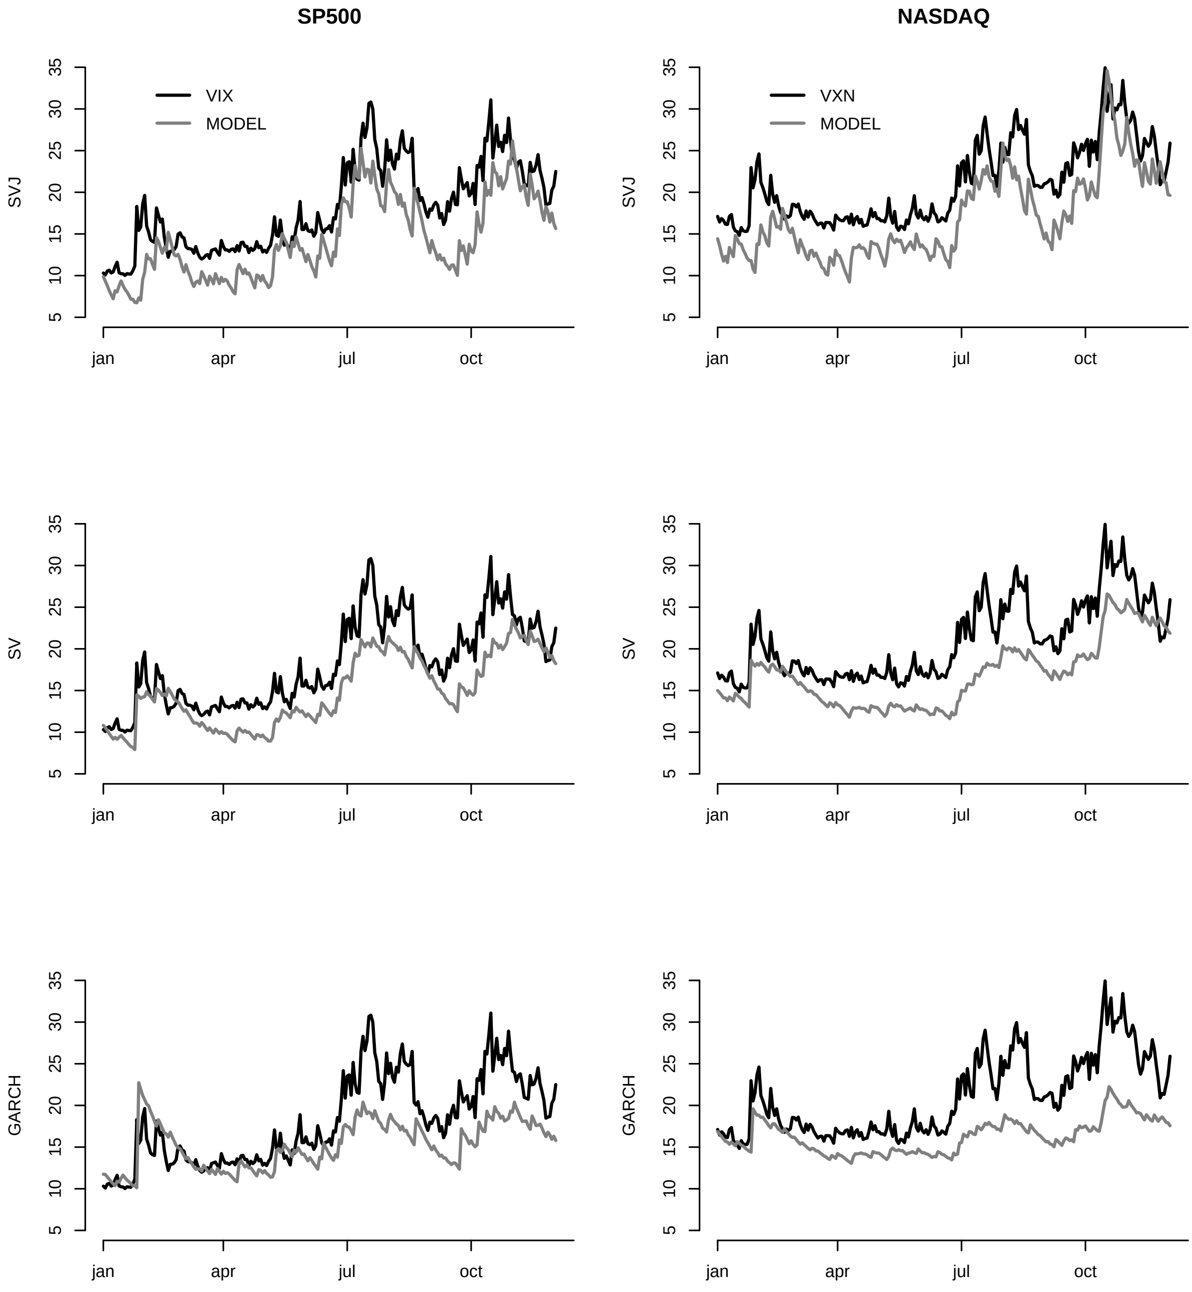

Figure 4: *Year 2007:* Comparing Garch, SV and SVJ models.

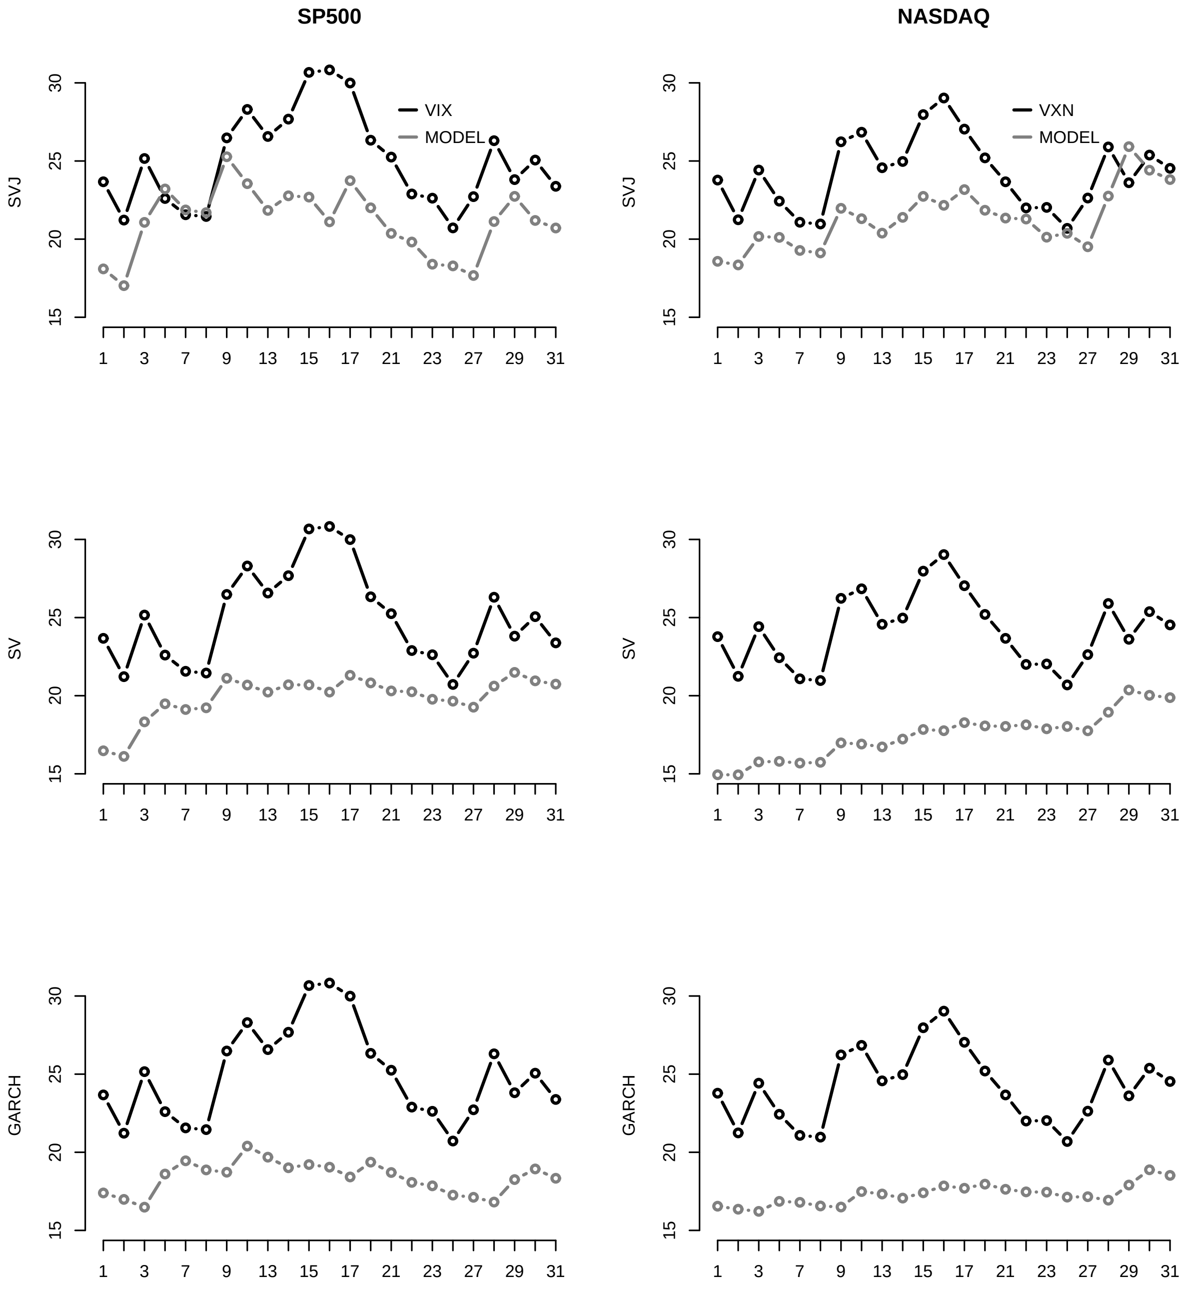

Figure 5: *August 2007:* Comparing Garch, SV and SVJ models.

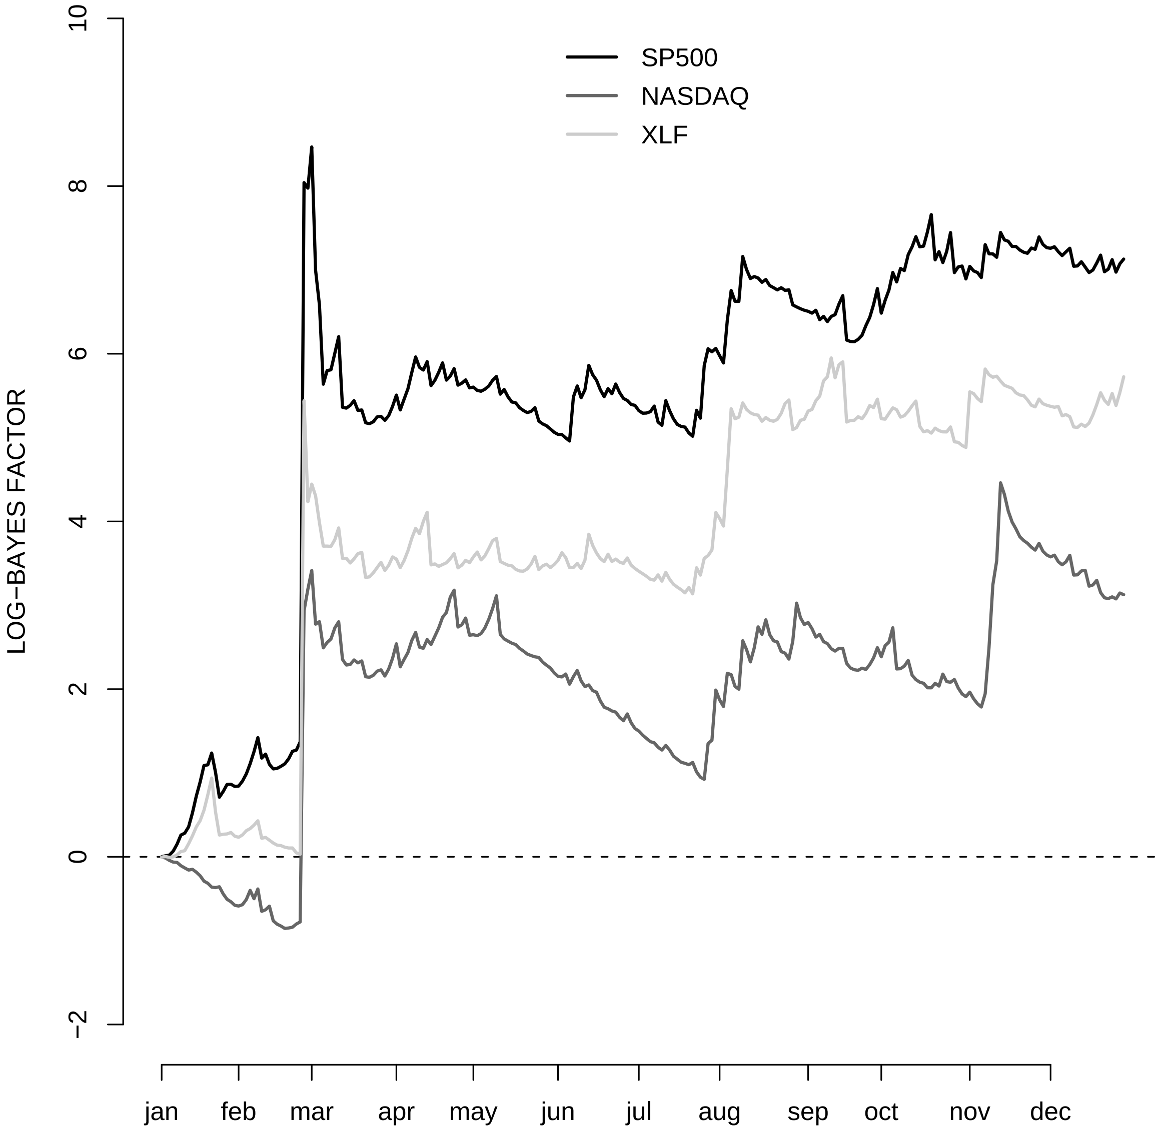

Figure 6: *Year 2007:* Sequential (log) Bayes factor, $BF(M_1, M_0)$. $M_1 \equiv$SVJ model $M_0 \equiv$SV model.