Recent Developments in Stochastic Volatility: Statistical Modelling and General Equilibrium Analysis

George Tauchen
Duke University

September 3, 2004 *†

Abstract

The paper reviews findings from recent estimations using various techniques of parametric continuous time models for financial price data, and it highlights some of the identification problems. There is a clear need to make use of the high frequency data in estimation in order to break the identification problems. The evidence from nonparametric volatility and jump component estimation is used to conjecture the characteristics of a better fitting parametric model, though the estimation of that model, which would formidable, is deferred to the future. The paper also explores some recently proposed schemes to simulate from stochastic volatility models built up from Lévy processes, and it considers some general equilibrium aspects of stochastic volatility.

Keywords: stochastic volatility, estimation of diffusions, realized variance, quadratic variation.

JEL classification: G12, C51, C52.

---

*Presented at the joint 59th European Meeting of the Econometric Society and 19th Annual Congress of the European Economics Association, Universidad Carlos III de Madrid, August 2004. Author Contact: George Tauchen, Department of Economics, Duke University, Box 90097, Durham NC 27708, E-mail george.tauchen@duke.edu

†I am grateful to Neil Shephard, who provided substantial guidance at the outset of the project; to Tim Bollerslev, who framed for me many of the current issues in the use of very high frequency data; and to James Davidson and Mikhail Chernov who provided many insightful comments. I also thank Viktor Todorov for development of the very practicable scheme to simulate from a CARMA(p, q) model and for helpful comments on an initial draft.

1 Introduction

Over the past four to five years, the stochastic volatility literature has proceeded along two parallel lines. One line has entailed the specification and estimation of parametric continuous time models applied to daily data. The model (i.e. the maintained DGP) is thought of as evolving in continuous time, and the econometrician observes realizations at discrete intervals. The task of estimating a discretely sampled diffusion is itself a most interesting problem because the likelihood of the discretely sampled data under the model is not available in a convenient closed-form manner. There are more complications associated with the need to introduce additional state variables to capture time varying volatility, among other things. The additional variable(s) are needed to mimic the fact that an observed financial price process, regardless of the sampling interval, is not Markov. The probability distribution of the price from the current value forward depends upon the entire observed history of the price, because that history contains information about the future evolution of the volatility. The loss of information due to the discrete sampling, together with the unobserved state variables, creates formidable econometric challenges that have attracted intense interest in the econometrics literature. There are now simulated-moments techniques, approximate numerical methods, and Markov Chain Monte Carlo (MCMC) methods that work with varying degrees of empirical success.

The other line is much more nonparametric in character and entails using very high frequency price data to produce estimates, typically the realized variance, of the integrated variance of the process over an interval of time. The time interval is almost invariably a day. The realized variance estimates have been used to a modest extent to estimate parametric stochastic volatility models, but largely they are viewed as nonparametric estimates of volatility, which can be modelled in their own right using reduced form techniques. More recently, extensions of these methods have been proposed to separately estimate the integrated variance and cumulative contributions of jumps over the time interval.

This paper reviews the findings from recent estimations using various techniques of parametric continuous time models, and it highlights some of the identification problems. There is a clear need to make use of the high frequency data in estimation in order to break these identification problems. The evidence from nonparametric volatility and jump component estimation is used to conjecture the characteristics of a better fitting parametric model, though the estimation of that model, which would formidable, is deferred to the future. The paper also explores some recently proposed schemes to simulate from stochastic volatility models built up from Lévy processes, and it considers some general equilibrium aspects of stochastic volatility.

2 Volatility Factor Models

Nearly all of the early stochastic volatility literature is set in discrete time. The model describes the movement of the financial process from one discrete point, typically the daily closing price, to the next, without a description the price evolves over the day.

More recently, however, it has become common to think of the price process as discrete, usually equi-spaced, observations on a process that is evolving in continuous time. Let p_t denote the log of a stock price (or stock index) defined for the continuous index $t \in [0, T]$. In what follows, the unit of measurement of time is one day; the sequence $\{p_t\}_{t=0,1,2,\ldots,T}$ denote the discretely sampled log price process, and $y_t = p_t - p_{t-1}$ denote the daily (geometric) return. The timing convention differs from the financial derivatives literature, where time is measured in annual units, with 252 trading days per year, so one day is 1/252 on this scale. The annual time scale is convenient for derivatives modelling, but it leads to poor scaling from an econometric/statistical perspective, so throughout we treat the basic time unit as one trading day.

2.1 Diffusive Models

2.1.1 Affine Factor Models

The affine models of Duffie and Kan (1996) provide a very rich class of convenient models and a natural starting point for a discussion of stochastic volatility. One of the main appeals of the affine class of models is that derivatives prices can be computed by Fourier transform methods (Heston, 1993, Duffie, Pan, and Singleton, 2000). A basic continuous time affine stochastic volatility model takes the form

$$
\begin{aligned}
dp_t &= \mu dt + \sigma_t dw_{pt} \\
dv_t &= (\alpha_0 + \alpha_1 v_t)dt + \sqrt{v_t} dw_{vt} \\
\sigma_t &= \sqrt{\xi_0 + \xi_1 v_t}
\end{aligned}
\tag{1}
$$

where the drift μ reflects the average rate of increase in the stock price, usually constant, dw_{pt} and dw_{vt} are Wiener processes (standard Brownian motions), σ_t is the local price volatility, the parameters α_0 and α_1 control the average level of volatility and the speed of mean reversion of the volatility process, while ξ_0 and ξ_1 control the mapping from the underlying volatility factor to the local price volatility. Common practice is to allow for correlation between the Wiener processes, $\text{Corr}(dw_{pt}, dw_{vt}) = \rho_{pv}$, and estimates of ρ_{pv} are almost always negative to reflect the leverage effect discussed below. The process v_t is a standardized CIR or Feller process, and under suitable regularity conditions the differential equation admits a solution on $(0, \infty)$. The above model is frequently called the Heston (1993) model. Of course, in practice there is a need to impose a normalization to achieve identification (e.g., α_0 and ξ_0 are not separately identified), but the normalization is generally not shown here because the formulas are a bit easier to interpret without it.

As now well understood, the model (1) cannot be estimated directly because the likelihood function $f(y_t|\rho, \{y_{t-j}\}_{j=1,2,\cdots})$, where ρ is a vector containing all of the parameters, is not readily available in convenient closed form. However, the model can be easily simulated using an Euler scheme, or a more complicated scheme (Kloeden and Platen, 1992) and estimated via a simulation-based technique such as the Efficient Method of Moments (EMM) of Gallant and Tauchen (1996, 2002), numerical maximum likelihood (Elerain, Chib, and Shephard 2001, Durham and Gallant, 2002), or MCMC-related methods (Eraker, 2001, Chib, Nardari, Shephard, 2002, Eraker, Jacqier, Polson, 2003). In all cases, the role of simulation is to integrate out the unobserved volatility factor v_t in order to reduce the problem down to one involving functions only of observed variables, namely the y_t process.

The model (1) does not fit the daily stock index returns data (Anderson, Benzoni, and Lund, 2002, Chernov *et al*, 2003, Eraker, Johannes, Polson, 2003). There are two reasons for the poor fit, as anticipated by Meddahi (2002). The observed y_t process is conditionally heteroskedastic with an error distribution that is thick-tailed or leptokurtic relative to the Gaussian distribution. A stochastic volatility model can certainly produce persistent stochastic volatility. Likewise, it can certainly produce non-Gaussian leptokurtic data; indeed, Clark (1973) introduced a random volatility model that was a scale mixture of normals precisely to explain the non-Gaussian character of financial price changes in a situation where one would expect a central limit theorem to apply. The problem with the model (1) is that there is only one volatility process, but it has to accomplish two purposes, which it cannot simultaneously do. In practice, if one estimates the model via a simulation-type technique, then there are often two isolated extreme values for the objective function. One extreme entails a value for α_1 that is negative but very small in magnitude; the other entails a value for α_1 that is negative but very large in magnitude. In the former case, the model is attempting to capture the persistent stochastic volatility and in the latter it is trying to accommodate the thick tails.

A natural empirical strategy is to introduce a second volatility factor, so that there is one factor for each task. Two-factor volatility structures have been developed empirically by Engle and Lee (1999), Gallant, Hsu, and Tauchen (1999), Barndorff-Nielsen and Shephard (2001a, 2002a, and additional papers at www.levyprocess.org), Alizadeh, Brandt, and, Diebold (2002), among others. In the affine case the model is

$$\begin{aligned} dp_t &= \mu dt + \sigma_t dw_{pt} \\ dv_{1t} &= (\alpha_{10} + \alpha_{11} v_{kt})dt + \sqrt{v_{kt}}\, dw_{v_1,t} \\ dv_{2t} &= (\alpha_{20} + \alpha_{21} v_{kt})dt + \sqrt{v_{2t}}\, dw_{v_2,t} \\ \sigma_t &= \sqrt{\xi_0 + \xi_1 v_{1t} + \xi_2 v_{2t}} \end{aligned} \tag{2}$$

As documented in (Chernov *et al*, 2003), this two-factor affine model does much better than the one factor model, but it still does not fit the daily returns data. The difficulty appears to be that the $\sqrt{\ }$ volatility function (2) is concave. Stock prices take sudden large movements. With a concave volatility function, large realizations of the volatility factors have diminished impact on price volatility itself, thus limiting the model's ability to reflect such sudden movements.

2.1.2 Exponential Linear Models

The use of exponential models with a linear Gaussian volatility factor has a long history in the stochastic volatility and financial econometric literatures.

The basic model is

$$
\begin{aligned}
dp_t &= \mu dt + \sigma_t dw_{pt} \\
dv_t &= (\alpha_0 + \alpha_1 v_t)dt + dw_{vt} \\
\sigma_t &= \exp(\xi_0 + \xi_1 v_t)
\end{aligned}
\tag{3}
$$

The model (3), can generate more violent stock price movements via the exponential volatility function, but it also will not fit the daily returns data, and for the same reasons the one-factor affine des not.

Chernov *et al* (2003) generalize the linear exponential setup to two factors with linear drift and diffusion terms. They use linear diffusions for the volatility factors along with an exponential form for the volatility function. Their model is

$$
\begin{aligned}
dp_t &= \mu dt + \sigma_t dw_{pt} \\
dv_{1t} &= (\alpha_{10} + \alpha_{11} v_{kt})dt + (\beta_{10} + \beta_{11} v_{1t})dw_{v_1,t} \\
dv_{2t} &= (\alpha_{20} + \alpha_{21} v_{kt})dt + (\beta_{20} + \beta_{22} v_{2t})dw_{v_2,t} \\
\sigma_t &= \exp(\xi_0 + \xi_1 v_{1t} + \xi_2 v_{2t})
\end{aligned}
\tag{4}
$$

Correlations $\mathrm{Corr}(dw_{pt}, dw_{v_k,t}) = \rho_{pk}$ are allowed to capture leverage effects. This model differs from (3) in that it has two volatility factors and the feedback terms $\beta_{kk} v_{kt}$ in the diffusion terms. Actually, the raw exponential function cannot be used directly, or there are problems with sufficient conditions for existence of a solution and for the Euler simulation scheme; the strategy is to splice the exponential to a near-linear function for extremely large values of the argument. Chernov *et al* (2003) present evidence that (4) is adequate for daily post-WWII aggregate stock market data. As expected, one of the processes is strongly mean reverting and the other very slowly mean reverting. The leverage effects are important and, interestingly, the volatility feedback $\beta_{kk} v_{kt}$ is very important for the strongly mean reverting factor but is unneeded for the very slowly mean reverting factor, which can just be a linear Gaussian process.

2.1.3 Semiparametric Factor Models

The models just discussed have a common structure where multiple underlying volatility factors follow relatively simple diffusion models that affect local price volatility through a volatility index function as in (2) and (4). An intriguing alternative would be to preserve the simplicity of the dynamics of the volatility factors and be nonparametric about the functional form of the volatility index function. For example, from Dai and Singleton (2000) we know that all affine diffusions can be decomposed into two blocks. One block consists of Gaussian-type processes with local volatilities drive by a set of nearly autonomous square root type process. Such a model for a vector of volatility factors $v_t = (v'_{gt} v'_{rt})'$ is

$$\begin{aligned} dv_{gt} &= (\alpha_{g0} + \alpha_{g1} v_{gt})dt + \sqrt{\beta_{g0} + \beta_{g1} v_{at}}\, dw_{gt} \\ dv_{at} &= (\alpha_{a0} + \alpha_{a1} v_{at})dt + \sqrt{v'_{at}}\, dw_{gt} \end{aligned} \tag{5}$$

where the α's and β's are suitably conformable parameter vectors or matrices. The square root factors comprising v_{at} are completely autonomous if the matrix α_{a1} is diagonal, though that is not required in the general affine setup. The volatility index function is

$$\sigma_t = \nu(\xi_0 + \xi'_g v_{gt} + \xi'_a v_{at}) \tag{6}$$

and the price dynamics are

$$dp_t = \mu dt + \sigma_t dw_{pt}. \tag{7}$$

In (6) the functional form is viewed as unknown, which gives the models its semiparametric character. Natural constraints to place on the volatility index function are that it is positive and monotone. Adaptation of the techniques in Berestenau (2003) for imposing such constraints on a cubic spline might be very worthwhile.

A very interesting and related method is the eigenfunction approach of Meddahi (2001). In this setup there is one (or a very few) few basic factors following stochastic diffusions, and then the volatility function is constructed as linear combinations of the eigenfunctions of the infinitesimal generator. This way, fairly complicated volatility dynamics can be built up from only one or a few few factors, which adds considerable richness and flexibility. The approach is especially powerful for getting analytical expressions for important functions that arise in the theory of quadratic variation (Andersen, Bollerslev, and Maddahi, 2003).

2.2 Jump Diffusion Models

An alternative to introducing additional volatility factors to accommodate the thick tails of the conditional density of the returns process is simply to allow for discrete jumps in the price process itself. Jump diffusions have a long and rich history in financial economics starting with Merton (1976). Recently Andersen, Benzoni, and Lund (2002) consider affine jump models of the form

$$
\begin{aligned}
dp_t &= \mu dt + \sigma_t dw_{pt} + \kappa_t q_t \\
dv_t &= (\alpha_0 + \alpha_1 v_t)dt + dw_{vt} \\
\sigma_t &= \sqrt{\xi_0 + \xi_1 v_t}
\end{aligned}
\tag{8}
$$

where q_t is a Poisson counting process with intensity λ and κ_t is the jump size; they also examine the exponential version

$$
\begin{aligned}
dp_t &= \mu dt + \sigma_t dw_{pt} + \kappa_t q_t \\
dv_t &= (\alpha_0 + \alpha_1 v_t)dt + dw_{vt} \\
\sigma_t &= \exp(\xi_0 + \xi_1 v_t)
\end{aligned}
\tag{9}
$$

as do Chernov *et al* (2003). Eraker, Johannes, and Polson (2003) go further and explore the role of jumps in volatility as well jumps in returns with a model close in form to:

$$
\begin{aligned}
dp_t &= \mu dt + \sigma_t dw_{pt} + \kappa_{1t} q_{1t} \\
dv_t &= (\alpha_0 + \alpha_1 v_t)dt + (\beta_0 + \beta_1 v_t)dw_{vt} + \kappa_{2t} q_{2t} \\
\sigma_t &= \sqrt{\xi_0 + \xi_1 v_t}
\end{aligned}
\tag{10}
$$

where q_{1t}, q_{2t} are Poisson counting processes and the κ's are the jump sizes.

2.3 Summary on Fitting the Daily Returns Data

A simple message emerges from the Andersen, Benzoni, and Lund (2002), Chernov *et al* (2003), and Eraker, Johannes, and Polson (2003) papers. If estimated on long daily data on broad stock market indexes, the suitably enhanced two-factor model (4), with continuous sample paths, and the jump diffusion models, (8) and (10), all appear to fit the data about equally well. The daily returns data alone are just not sufficiently informative. Data sets on daily (log) returns $y_t = p_t - p_{t-1}$ alone are not strong enough to discriminate among competing models for the price process, and in particular between a model with multiple volatility factors versus a jump diffusion. Despite the radically different internal structure of these models, their implications for daily price movements are about the same.

3 Models Built from Lévy Processes

To discriminate better across models, we have to use the higher frequency data. Options data might seem appealing for dealing with this identification problem, but we lack a good options pricing model for both the objective and risk neutral distributions. Also, there is evidence that the options prices themselves exhibit economic problems such as abnormally large risk premiums to writing options (Eraker, 2004). So it makes sense to confine attention to the stock price data alone, although processes with only Brownian increments and possibly rare jumps (jump diffusions) will not suffice for the very high frequency. Consequently, we turn to some more general models built up from Lévy processes. A cádlág stochastic process is a Lévy process if the increments are strictly stationary, independent, and the process is continuous in probability. There are many good references on Lévy processes such as Bertoin (1996) and Cont and Tankov (2004); see also Barndorff-Nielsen and Shephard (2001b).

Although Lévy processes are valuable components for building models for stock price dynamics, it needs to be kept in mind that the stock price series itself is never a Lévy process. The stock price increments over short intervals of time are essentially uncorrelated, but the increments are not independent because of the persistent time varying volatility. Indeed, once the discretely sampled process is verified to display GARCH — as essentially all financial price series do — then the underlying continuous process is non-Markovian and perforce not a Lévy process.

3.1 Lévy-Driven Volatility (Subordinators)

In Section 2 the underlying driving process(s) are locally Gaussian and positivity of the volatility process ensured by the functional form of the volatility index function (6). Barndorff-Nielsen and Shephard (2001a, 2001b) suggest an alternative approach, where the driving process is a Lévy process with nonnegative increments; simple parametric sign restrictions ensure positivity. In their basic model, the price and volatility dynamics are governed by

$$\begin{aligned} dp_t &= \mu dt + \sigma_t dw_{pt} \\ dv_t &= \alpha v_t dt + d\mathcal{L}_t \\ \sigma_t &= \sqrt{\beta_v v_t} \end{aligned} \tag{11}$$

where μ is the price drift, $\alpha < 0$, $\beta_v > 0$ and $\mathcal{L}_t$ is a Lévy process with nonnegative increments. Of course the constant β_v could be absorbed into the Lévy process, but it proves convenient to keep it separate in order to scale simulations. The expression for v_t is to be interpreted as

$$v_t = e^{\alpha t} v_0 + \int_{s=0}^{t} e^{\alpha(t-s)} d\mathcal{L}_s, \tag{12}$$

and $v_t > 0$, so long as $v_0 > 0$, because the increments of $\mathcal{L}_t$ are nonnegative. The model for v_t is a nonnegative Ornstein-Uhlenbeck (OU) process. As emphasized by Barndorff-Nielsen and Shephard (2001a, 2002a), a two-factor structure for volatility can easily be handled in this setup by simply using two nonnegative OU processes in a fashion entirely analogous to (2) and (4) above,

$$\begin{aligned} dp_t &= \mu dt + \sigma_t dw_{pt} \\ dv_{1t} &= \alpha_1 v_t dt + d\mathcal{L}_{1t} \\ dv_{2t} &= \alpha_2 v_t dt + d\mathcal{L}_{2t} \\ \sigma_t &= \sqrt{\beta_{1v} v_{1t} + \beta_{2v} v_{2t}}. \end{aligned} \tag{13}$$

Here the α_k are negative, the β_k are positive, and the $\mathcal{L}_{kt}$ are Lévy processes with nonnegative increments. The extension to the superposition of many such processes is straightforward,

$$\sigma_t = \sqrt{\sum_1^K \beta_{kv} v_{kt}} \tag{14}$$

where each of the v_t is a nonnegative OU processes.

Brockwell and Davis (2001) and Brockwell (2003) propose an extension, termed CARMA(p, q), that has a convenient notation and offers potential efficiency gains for generating simulations. A Lévy-driven continuous time CARMA(p, q) process is defined by the formal equation

$$a(D) v_t = b(D) D \mathcal{L}_t, \quad t \geq 0, \tag{15}$$

in which D denotes differentiation with respect to t, $\mathcal{L}_t$ is a Lévy process with $\mathrm{E}(\mathcal{L}_1^2) < \infty$, and a and b are polynomials

$$\begin{aligned} a(z) &= a_0 + a_1 z^1 + \cdots + a_p z^p \\ b(z) &= b_0 + b_1 z^1 + \cdots + b_q z^q, \quad q < p. \end{aligned} \tag{16}$$

The solution of (15) can be represented as

$$v_t = \kappa(t) v_0 + \int_{s=0}^{t} \kappa(t-s) d\mathcal{L}_s, \quad v_0 > 0, \tag{17}$$

where κ is a deterministic function, called the *kernel* function, and integral is defined as the L^2 limit of Riemann-Stieltjes sums. If the roots of the polynomial $a(z)$ all have negative real parts, then $\lim_{t\to\infty} \kappa(t) = 0$ and the solution admits a stationary representation

$$v_t = \int_{s=-\infty}^{t} \kappa(t-s) d\mathcal{L}_s^* \tag{18}$$

where $\mathcal{L}^*$ is an extended Lévy process (Brockwell, 2003). The kernel function κ is easily obtained from the coefficients of the a and b polynomials. So long as the kernel function is nonnegative, then the variance process (17) is ensured to be positive.

The CARMA(p, q) stochastic volatility model is

$$\begin{aligned} dp_t &= \mu dt + \sigma_t dw_{pt} \\ a(D)v_t &= b(D)D\mathcal{L}_t \\ \sigma_t &= \sqrt{\beta_v v_t} \end{aligned} \tag{19}$$

The Lévy process driving volatility is a pure jump (no diffusive component) non-decreasing process with characteristic function

$$\varphi_t(\omega) = \mathrm{E}\exp(i\omega v_t) = \exp[\ t \int_{0^+}^{\infty} (e^{i\omega y} - 1) d\nu(y)\,], \tag{20}$$

where ν is the Lévy-measure on $(0, \infty)$. The process may be finitely active or infinitely active, $\nu[(0, \infty)] \leq \infty$, and since

$$\int_{0^+}^{1} y d\nu(y) < \infty \tag{21}$$

the volatility process is of finite variation.

It turns out to be quite feasible to simulate directly trajectories v_t from a CARMA(p, q) volatility model. We always work with the case where the roots of $a(z)$ have negative real parts. The unconditional distribution of v is generally unknown, so the strategy is to start the process from a sensible value for v_0 in (17), let the process run for a burn-in period, and take values after that. Thus, the basic task is to simulate

$$\int_{s=0}^{t} \kappa(t-s) d\mathcal{L}_s. \tag{22}$$

It turns out to be quite simple to simulate directly the integrated variance over the interval $[t - a, t]$ as

$$\begin{aligned} iv_t^a &= \textstyle\int_{u=t-a}^{t} v_u du \\ &= \textstyle\int_{u=t-a}^{t} \int_{s=0}^{u} \kappa(u-s) d\mathcal{L}_s\ du \quad t \geq a > 0, \end{aligned} \tag{23}$$

and applying the stochastic Fubini Theorem (everything is square integrable) gives

$$iv_t^a = \int_{s=0}^{t} \int_{u=\max(s,t-a)}^{t} \kappa(u-s)du\ d\mathcal{L}_s \tag{24}$$

or

$$iv_t^a = \int_{s=0}^{t} \kappa^*(t-s)\, d\mathcal{L}_s, \tag{25}$$

where the kernel κ^* is readily obtained from the basic kernel κ. The integrated volatility process (25) takes the same form as (22) except with a different kernel function. Since the weighted sums of AR models is an ARMA model, the CARMA(p,q) can be expected to capture the basic features of a multi-factor additive model such as (14), but it only entails simulating a single Lévy process instead of multiple Lévy processes. Also, as noted by Brockwell (2003), the CARMA(p,q) model entails fewer monotonicity constraints on the autocorrelation function and the kernel function.

Todorov (2004) adapts the work of Rosinski (2001) to develop a very practicable scheme to simulate from a CARMA(p,q) model. An illustration for a CARMA$(2,1)$ follows. A convenient parameterization is to set

$$\begin{aligned} a(z) &= (z-\rho_1)(z-\rho_2), \quad \rho_1 < 0,\ \rho_2 < 0 \\ b(z) &= 1 + b_1 z. \end{aligned} \tag{26}$$

The implied kernel function κ is easily determined from Brockwell (2003), but care must be taken to choose b_1 small enough to ensure that the kernel function is nonnegative over the domain of integration. After generating a simulated trajectory for v_t, then use the parameter β_v to scale $\sigma_t = \sqrt{\beta_v v_t}$ to a typical value for volatility. Here we scale so that the average value of σ_t is unity, because the mean average volatility of the S&P Index is one percent per day (or 16 percent per annum). The driving Lévy process is a gamma process with Lévy density given by

$$\frac{d\nu(y)}{dy} = c\frac{e^{-\lambda y}}{y}, \quad y > 0, \tag{27}$$

where c is a positive parameter that governs the intensity of the jumps and λ is a positive parameter that governs the magnitude of the jumps. The process is infinitely active with paths of finite variation.

Given the kernel function κ and the gamma specification (27) for the Lévy density, then the simulations based on Todorov's method proceed as follows: Let $[0,T]$ denote the interval over which the simulation is defined, where T is in units of days. For a given level of precision τ, generate independent T_i from the standard exponential distribution (intensity parameter equal to unity) and set $\Gamma_i = \sum_{j=1}^{i} T_j/T$; generate independent V_i from the standard exponential distribution and U_i from the uniform $[0, T]$. Continue the generation until $\Gamma_i > \tau$. Set

$$\hat{v}_t = (\frac{1}{\lambda}) \sum_i e^{-\frac{\Gamma_i}{c}} V_i \, \kappa(t - U_i) 1_{U_i \leq t}, \quad t \in [0, T], \tag{28}$$

and $\hat{v}_t \to v_t$ uniformly almost surely as $\tau \to \infty$, T fixed.

Figures 1–4 display very long simulated trajectories of CARMA$(2, 1)$ volatility process for the different parameter settings shown in Table 1. The simulations presume eight trading hours per day with the Lévy process and spot variance, σ_t^2, sampled at 30-minute intervals, or $M = 16$ observations per day. The simulation runs for 3,000 trading days, which corresponds twelve years worth of data, after a burn-in period of 1,000 days. Also shown are the thirty minute return, the kernel function κ, and the driving Lévy process. The return simulations abstract from the leverage effect and are achieved by multiplying the square root of the integrated variance over the thirty minute inteveral by a standard $\mathrm{N}(0, 1)$ Gaussian random variable. Both the spot variance and the integrated variance are direct simulations from the underlying CARMA$(2, 1)$; there is no Euler-type discretization anywhere. The accuracy parameter $\tau = 10$ and the appearances of the figures are quite insensitive to values above that. The simulations take about 30 minutes in Fortran 90 run under Linux on new vintage (circa 2004) PC; there is substantial room for improving the simulation code and generating simulations far more quickly.

The figures, especially Figures 1 and 2, show how the CARMA$(2, 1)$ model can capture a two-factor type structure, with one very strongly mean reverting factor and another more slowly mean reverting. The figures also show the great range, from very rugged to very smooth, types of volatility trajectories that a CARMA model can produce.

Common practice in estimating stochastic volatility models is to use much longer data sets, commonly 40 years or about 10,000 daily data points. Thus it is interesting to see if the CARMA approach holds promise for that length. Figure 5 shows a simulation run for 10,000 days, eight trading hours per day, with a sampling rate of 30-minute intervals. The parameter settings are in Table 1, with the main feature being that one root corresponds to extreme persistence and the other to quick mean reversion. Figure 5 shows the simulated daily integrated variance and the implied simulated daily return process. The daily integrated variance is simulated directly; it is not an aggregation from an Euler-type discretization. The implied daily log return process appears quite similar to that typically observed. Figure 6 shows 100 $\times$ the daily (log) return process for the Dow Jones Industrial Average from 1953 to mid-2004. The appearance of the simulated daily process in the bottom panel of Figure 5 is very much like that of the observed process in Figure 6, except perhaps that it lacks the very most extreme fluctuations seen in the observed data.

Taken together, the simulations shown in Figures 1–5 suggest that a CARMA(p,q) with a nondecreasing driving Lévy process appears to be a very promising volatility model. The setup can generate a wide ranging set of volatility dynamics and, with only rough calibration (a little tinkering and some intuition), can produce patterns of price movements similar to what is actually observed. Thus, when used in formal estimation, a topic discussed in more detail below, one can expect a CARMA(p,q) to perform quite well.

An intriguing, and very promising alternative to the CARMA(p,q) with a nondecreasing driving Lévy process is the Wishart Autoregressive (WAR) model of Gourieroux-Jasiak-Sufana (2004). That comparison is deferred to future work.

3.2 The Nonparametric Evidence on Price Level Jumps Variances

The purpose of this subsection is to review the recent evidence generated from very high frequency data on the existence of jumps in price process itself. With the exception of the jump diffusion model of Subsection 2.2, all of the previously discussed models, including the CARMA(p,q), produce price processes with continuous sample paths. The major objective of the new nonparametric tests for jumps is to determine whether the additional complication of jumps in the price process itself is needed. It is quite different from the analytical approaches of Ait-Sahalia (2002, 2003).

The starting point is the model for the log price p_t evolving in continuous time as

$$dp_t = \mu\, dt + \sigma_t dw_{pt} + \beta_p\, d\mathcal{L}_{pt} \tag{29}$$

where μ is the drift, assumed constant for simplicity, σ_t is the local or spot volatility such that σ_t is cádlág and $0 < \int_0^t \sigma_t^2 ds < \infty$, w_{pt} is a standardized

Brownian motion, and $\mathcal{L}_{pt}$ is a compound Poisson process, equivalently a Lévy process with piecewise constant trajectories so that its Lévy measure ν_p satisfies $\nu_p[(-\infty,\infty)] < \infty$. The parameter β_p is kept separate from the driving Lévy $\mathcal{L}_{pt}$ process only to facilitate scaling.

Again, time is measured in daily units and for integer t define the within-day geometric returns

$$r_{t,j} = p_{t-1+\frac{j}{M}} - p_{t-1+\frac{j-1}{M}}, \quad j = 1, 2, \ldots, M, \tag{30}$$

where M is the sampling frequency. The objective is to provide evidence on whether the jump component is needed in (29), but without making strong parametric assumptions as done for the models discussed in Section 2.

Barndorff-Nielsen and Shephard (2004a, 2004b) study general measures of realized within-day price variance, and two natural measures emerge from their work. The first is the familiar **Realized Variance**

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2, \tag{31}$$

and the other is the realized **Bipower Variation**

$$BV_t = \mu_1^{-2} \sum_{j=2}^{M} |r_{t,j}||r_{t,j-1}| = \frac{\pi}{2} \sum_{j=2}^{M} |r_{t,j}||r_{t,j-1}|, \tag{32}$$

where

$$\mu_a = \mathrm{E}(|Z|^a), \quad Z \sim N(0,1),\ a > 0. \tag{33}$$

The notation here absorbs μ_1^{-2} into the definition of the Bipower variation and thereby makes it directly comparable to the Realized Variance. As is well known the Realized Variance satisfies

$$\lim_{M\to\infty} RV_t = \int_{t-1}^{t} \sigma^2(s)ds \ + \sum_{t-1\le s\le t} (\Delta p_s)^2, \tag{34}$$

where the sum is over the nonzero jumps $\Delta p_s = p_s - p_{s-}$, of which there are only finitely many over any interval. Thus, the RV_t is an estimator of the integrated variance $\int_{t-1}^{t} \sigma^2(s)ds$ plus the jump contribution.

The key result of Barndorff-Nielsen and Shephard (2004a, 2004b) is that under mild regularity conditions

$$\lim_{M\to\infty} BV_t = \int_{t-1}^{t} \sigma^2(s)ds. \tag{35}$$

The statistic BV_t provides an estimate of the integrated variance unaffected by jumps. Evidently, the difference

$$DV_t = RV_t - BV_t \tag{36}$$

is an estimate of the pure jump contribution and can clearly provide a test for jumps. Another measure is the **Relative Jump** measure

$$RJ_t = \frac{RV_t - BV_t}{RV_t}, \tag{37}$$

which is an indicator of the contribution (if any) of jumps to the total within-day variance of the process. Huang and Tauchen (2004) show that RJ_t can be interpreted as a Hausman-type test for jumps. Under the null of no jumps, RV_t is the most efficient estimate of $\int_{t-1}^{t} \sigma^2(s)ds$ while BV_t is a less efficient estimator but more robust (with respect to jumps) and so the difference DV_t can be expected to be asymptotically independent of RV_t, which turns out to be the case. An equivalent statistic, -RJ, called the ratio statistic is studied by Barndorff-Nielsen and Shephard (2004a, 2004b), who provide the requisite asymptotic theory needed to studentize the measures DV_t and RJ_t. The statistics are

$$z_{Dt} = \frac{DJ_t}{\sqrt{\text{avar}(DJ_t)}} \xrightarrow{\mathcal{D}} N(0,1) \tag{38}$$

and the other off the ratios

$$z_{Rt} = \frac{RJ_t}{\sqrt{\text{avar}(RJ_t)}} \xrightarrow{\mathcal{D}} N(0,1) \tag{39}$$

where the limit is taken as $M \to \infty$ for the fixed interval $[t-1, t]$. The interval does not have to be a day, but it usually is because that is the largest length to work without having to deal with the overnight movement in prices for markets not operating 24 hours per day. The number of within-day returns M is chosen so that the sampling interval is no finer than five minutes, as anything more frequent leads to statistics corrupted by within-day micro-structure noise.

Huang and Tauchen (2004) conduct a Monte Carlo analysis of wide class of tests based on comparison of the BV_t to the RV_t. They find that tests based on the relative jump as in (39) always have better size properties than tests based on the difference as in (38), which tend to over reject somewhat when the null of no jumps is true. They also find that the ratio tests have very good power properties, and that the choice of the method to estimate the asymptotic variance does not matter a lot, so long as it is done in a jump-robust manner.

Andersen, Bollerslev and Diebold's (2003) use evidence from z-statistics much like (39) and generate evidence suggesting that there are too many large within-day movements in equity, fixed income, and foreign exchange prices to be consistent with the standard continuous time stochastic volatility model with continuous sample paths. They also produce time series reflecting the jump component within the day by defining

$$J_t = (RV_t - BV_t)I[z_{Rt} > z_{1-\alpha}] \tag{40}$$

where $I(\cdot)$ is the 0-1 indicator function and α is a small number such as 0.005 or 0.001. The idea is to pull out the jump contribution for those days when the z-statistic signals strong evidence for a jump.

Huang and Tauchen (2004) provide complementary empirical evidence on jumps, which is seen in Figure 7. The five panels show time series of z-statistics computed from S&P futures returns sampled at five minute intervals and computed on a daily basis for each of the 5186 days in the sample. The top three panels show daily z-statistics based on the difference (38), with slightly different methods used to estimate the asymptotic variance; the bottom two panels show plots of statistics based on the ratio (39) with two alternative methods used to estimate the asymptotic variance. The statistics in the bottom two panels are known to have the best finite samples properties. The figure also shows in each panel the upper 0.99 and 0.999 values of the standard normal distribution. Evidently, there are far many more observed z-statistics above the critical points than could be expected under the null of no jumps; this evidence strongly discredits continuous diffusion models for stock prices.

4 Statistical Models and Methods

To summarize the evidence from the previous two sections, an appropriate continuous time model for the stock price p_t would be of the form

$$\begin{aligned} dp_t &= \mu dt + \sigma_t dw_{pt} + \beta_p d\mathcal{L}_{pt} \\ a(D)v_t &= b(D)D\mathcal{L}_{vt} \\ \sigma_t &= \sqrt{\beta_v v_t} \end{aligned} \tag{41}$$

where the drift is presumed constant and $a(z)$ and $b(z)$ are polynomials of degree $p > q$, respectively. In the above, $\mathcal{L}_{pt}$ is a Lévy pure jump process included to accommodate the nonparametric evidence on jumps in the log stock price process itself; the volatility factor is described by a CARMA(p, q) model driven by a nondecreasing Lévy process $\mathcal{L}_{vt}$. The CARMA(p, q) is used to capture the known multi-factor structure of volatility. Finally, the increments dw_{pt} and $d\mathcal{L}_{pt}$ must be correlated with the increments of $\mathcal{L}_{vt}$ in order to capture the leverage effect.

Daily returns data will not suffice to identify the model (41). As noted in Section 2, there are simpler, special cases, that fit the daily returns data over 1953–2001 equally well. Two recent efforts to estimate specific versions of a model like (41) but without the CARMA(p, q) volatility structure, on daily data are Carr, Gemen, Madan, and Yor (2002, 2003) and Li, Wells, and Yu (2004). These papers utilize more complicated Lévy processes than those discussed in Section 2, but the data sets are shorter, and the simpler two-factor volatility models like (2) and (4) are not considered, so the identification issue is not confronted directly.

The estimation of a parametric specification of the general model (41) on a high frequency data presents very formidable challenges, both from the data perspective and an estimation perspective.

From the data perspective, consider the sequence of high frequency returns $r_\tau = p_{\tau\delta} - p_{(\tau-1)\delta}$, for $\tau = 1, 2, \cdots, T$ where δ is the sampling interval. The sampling interval δ cannot be too small or the returns will be dominated by microstructure noise (Bai, Russell, and Tiao, 2004, and references therein). Complications such as bid-ask bounce overwhelm the information in the data. It is generally regarded that, for heavily traded stocks or indices, data sampled at any more frequently than five minutes are too noisy to be of value for inferring price dynamics. Methods for handling the microsctructure noise (Andreou and Ghysels, 2002) and the optimal sampling interval is a topic of very intensive research (Ait-Sahalia, Mykland, and Zheng, 2004, Bandi and Russell, 2004, and the references therein). Also the return volatility is known to shown a deterministic pattern over the day, with high volatility in the early part of the day, lower in the middle, and higher again towards the later part of the day. In addition, at regularly spaced intervals a return will be the overnight return, which which also as a different variance. The raw high frequency returns will have to be adjusted for these and other diurnal patterns (Andersen and Bollerslev, 1998). The net effect is to produce and adjusted returns series r_τ^{adj} for $\tau = 1, 2, \cdots, T$ that is not actually the return on any traded security and is not sampled as finely as one might think possible.

These data issues could be argued to be minor and easily addressed by simple data transformations, but the estimation issues are truly formidable. The fundamental problem is that the transition density of r_τ^{adj} given $r_{\tau-j}^{adj}, j = 1, 2, \cdots$ is not available in simple closed form. Likelihood-based methods are thus intractable. Simulation will have to play a role somewhere in order to integrate out the unobserved volatility variable(s). Simulated likelihood in the manner of Durham and Gallant (2002) is infeasible because one lacks the joint density of the observables and unobservables, so there is no convenient way to define an importance function to keep the simulated likelihood from becoming hopelessly inefficient. The EMM technique of Gallant and Tauchen (1996, 2002) can potentially be applied to this context. Their method, however, requires a good reduced form statistical model for the r_τ conditional on $r_{\tau-j}, j = 1, 2, \cdots$, and it is not clear if the extant models are adequate for this purpose; also EMM requires very long simulations to obtain accurate numerical integrations, so very efficient simulation schemes would need to be devised for (41). Doubtless MCMC in the style of Eraker, Jacqier, and Polson (2003) and Li, Wells, and Yu (2004) could be adapted to this problem. However, MCMC implementations to date typically assume the sampling interval is the same as one tick on the continuous time clock (Eraker (2001), however, is an exception), which greatly reduces the number of unobserved variables to deal. But then it is not feasible to explore the model's implied behavior of the stock price at intervals finer than the sampling interval, and there is the issue of discretization bias to confront.

Andersen, Bollerslev, and Diebold (2003) argue for an approach different from a direct estimation of parametric versions of (41). Their approach uses the very high frequency data only to extract out the integrated variance and the jump component. That is, the observed vector at the daily frequency would be

$$y_t = \begin{pmatrix} p_t - p_{t-1} \\ BV_t \\ J_t \end{pmatrix}. \tag{42}$$

The three components are the daily return, the daily jump-robust bipower variation, BV_t in (32), and the daily jump component as defined in (40). The argument is that these three summary measures adequately capture the within day activity, and so reduced form modelling of y_t, at the daily frequency, will be adequate for risk management over a multi-day or multi-week horizon. Andersen, Bollerslev, and Meddahi (2003) present analytical evidence that, the for simple continuous diffusions, the realized variance RV_t in (31) comes very close to the optimal estimate under the model of the integrated variance. Whether that same statement is true for BV_t and J_t under dynamics (41) is an open question, however. Nonetheless, it seems clear that small specification errors in a parametric specification of (41) would get greatly magnified if the model were spun out for several days or weeks in an effort to simulated price trajectories.

If we are to persist in the task of estimating parametric versions of the (41) and make use of the high frequency data, there are actually two routes one might take and the best approach is not at all obvious *a priori*. One route, as previously discussed, would be attempt direct estimation using adjusted high frequency returns. Modified likelihood-based methods might adaptable to this task as are simulated method of moments techniques. The other route would be to be work at the daily level and retain from the high frequency only the summary measures as shown in (42). The idea would be to force underlying model (41) to confront the daily dynamics of daily process y_t in (42). It is not obvious how to do this within a likelihood-based context, although with enough cleverness one might be able to get the likelihood of the y_t process in (42) as implied by a parametric specification of (41). Simulated method of moments techniques currently appear more directly applicable to the task, though that can certainly change over time. Either way, the implications of the estimated model for a wide range of frequencies, five-minute, daily, weekly, and perhaps monthly, will have to be scrutinized as diagnostic checks on the overall sensibility of the estimated model.

5 General Equilibrium Analysis: The All-Important Leverage Effect

So far little attention has been paid to the drift of the price process, and often it has been assumed constant for simplicity. But from an economic perspective, the drift is perhaps the primary objective of interest. The extent to which the expected rate of return stands above the risk free rate is the risk premium, the reward for bearing risk. One common specification is to parameterize the drift as

$$\mu_t = a_0 + a_1 \sigma_t^2 \tag{43}$$

to reflect a continuous time version of GARCH-in mean.

Mention has been made above of the leverage effect, which is the well-known observed negative correlation between returns and volatility. Common practice in the stochastic volatility literature is to treat the leverage effect as a statistical regularity that can be accommodated by allowing for correlations among random shocks. The effect is typically thought of as static without much attention paid to the dynamic leverage effect, which is the cross-correlation between returns and volatility at various leads and lags.

Tauchen (2004) examines risk, return, and the dynamic leverage effect from a general equilibrium perspective. The model is quite different than the quasi-reduced form continuous time models set forth in the previous sections. In that setup, all interactions among variables like returns and endowments arise from the characteristics of agents' preferences, and closed form expressions are available are by using a log-linear approximation to the solution of the general equilibrium model. The model operates in discrete time and is better thought of as applying to far more coarsely sampled data such as monthly or quarterly. Nonetheless, the findings from that setup do have some bearing the continuous time modelling discussed above.

One finding pertains to the interpretation of risk premium specifications like (43), which is commonly thought as a risk premium on volatility. That interpretation is incorrect. In a model with expected utility and time varying conditional consumption volatility, a specification somewhat like (43) arises, but it simply reflects a time-varying risk premium on consumption (endowment) risk. The role of stochastic volatility is to create time variation in the factor loading on consumption risk. The stochastic volatility factor itself not priced in this model, and the reason is that the volatility variable does not enter the marginal rate of substitution (or pricing kernel) and thereby does not carry a risk premium. The expected return of an asset with dividend payment proportional to volatility exactly equals the risk free rate. There is no volatility risk premium at all, despite a well-defined GARCH-in-mean like (43).

It turns out from Tauchen (2004) that there are substantive connections among non-expected utility preferences, the volatility risk premium, and the dynamic leverage effect. Only in the case of non-expected utility preferences will there be a volatility risk premium and a leverage effect. Furthermore, the signs of the volatility risk premium and the leverage effect both directly depend upon two key economic parameters, the intertemporal marginal rate of substitution (ψ) and risk aversion (γ). There is general agreement in the economics literature that $\gamma > 1$ and considerable debate about whether $\psi < 1$ or $\psi > 1$. In the case $\psi < 1$, the volatility risk premium is negative and the leverage effect is positive, in contrast to extensive empirical evidence. On the other hand, if $\psi > 1$, then the the volatility risk premium is positive and the leverage effect is negative. Thus it seems that the common empirical finding in the stochastic volatility literature of a negative leverage effect has direct implications for the magnitude of an economic parameter. In addition, calculations presented in Tauchen (2004) show that with $\psi > 1$ and reasonable values for other parameters the general equilibrium model can produce a dynamic leverage effect much like that observed for the S&P 100 Index and the VIX volatility indicator.

6 Conclusion

The specification and estimation of fairly tightly constrained parametric stochastic volatility models has a longstanding history in financial econometrics. Within the past few years, continuous time stochastic volatility models have been estimated on long time series of daily stock market returns. Looking across several studies, it turns out that there are several different, equally plausible, parametric specifications that appear to fit the daily returns data about equally well. Evidently, if the research effort is to continue, then the now-available very high frequency data must be utilized in some sensible manner in order to discriminate better across models and produce sharper parameter estimates.

The high frequency data are known to be noisy, and some might argue that the effort of estimating parametric continuous time on models them might simply be abandoned. The maintained continuous time model is an overly simplistic paradigm that is useful only as a rough guide on how to think about the true DGP. An effort to estimate it directly would produce a model very sensitive to the microstructure noise and to small errors in specification. The model would likely produce unreliable price trajectories when run forward to a week or a month, as the small errors would get greatly magnified. Instead of parametric estimation, we can use the high frequency data to produce simple, and reliable, daily measures of quantities like the integrated variance and the jump contributions. For the purpose of risk management, these measures can modelled in a reduced-form manner, without much attention the underlying continuous time model that might have generated the data.

On the other hand, the effort to continue estimating parametric continuous time stochastic volatility models remains appealing, but there are serious issues to be confronted. If the model is to estimated directly on high frequency returns data, then the data cannot be sampled too frequently and must be adjusted to reflect known intra-day diurnal patterns. Data issues aside, there are serious problems associated with the estimation. It remains almost certain, just as with estimation on lower frequency returns, that simulation has to be used in someway to integrate out unobserved variables, in either a likelihood or moments-based approach. A complicating factor is that at the very high frequency, much more general Lévy processes are almost certainly more appropriate than the familiar diffusion and jump diffusion models. Efficient simulation schemes will clearly be needed, and there evidence in Section 3 that some recently developed schemes are extremely promising. An alternative to applying the model directly to the high frequency data is to apply it to the daily summary measures. That presents a serious challenge to likelihood based methods but less so to simulated moments methods.

7 Tables and Figures

Table 1: Parameter Settings for Gamma-Driven CARMA$(2, 1)$

	Figure 1	Figure 2	Figure 3	Figure 4	Figure 5
Parameter					
Half life of ρ_1 (days)	10.00	125.0	750.0	125.0	See Figure 2
Half life of ρ_2 (days)	0.10	0.50	.50	0.50	”
b_1	0.1e-2	0.1e-02	0.1e-02	0.1e-02	”
c	0.10	0.5e-02	0.5e-03	0.5e-01	”
λ	1.0	1.0	1.0	1.0	”

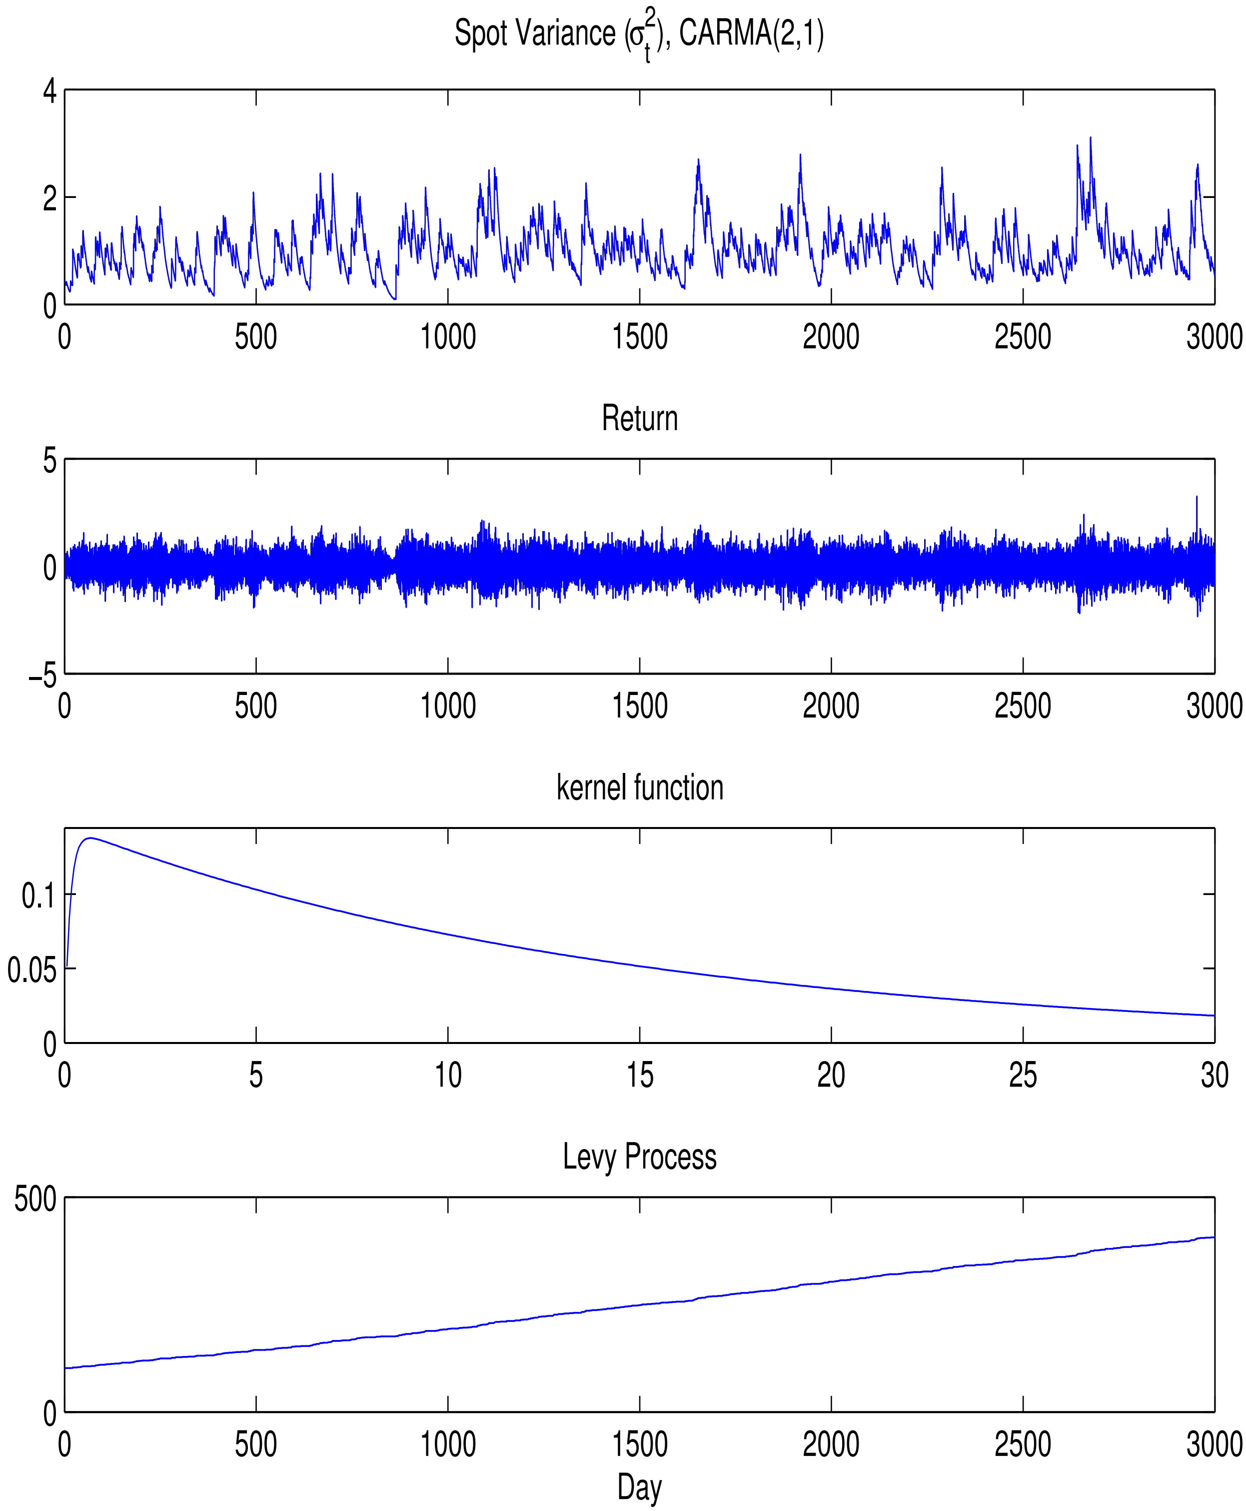

Figure 1: Simulated Realization of a CARMA(2,1) stochastic volatility model. The top panel shows the spot volatility at thirty minute intervals, eight hours per day; the second shows the return over the thirty minute intervals; the third shows the kernel function, and fourth the Lévy process at thirty minute intervals. Parameter settings shown in Table 1.

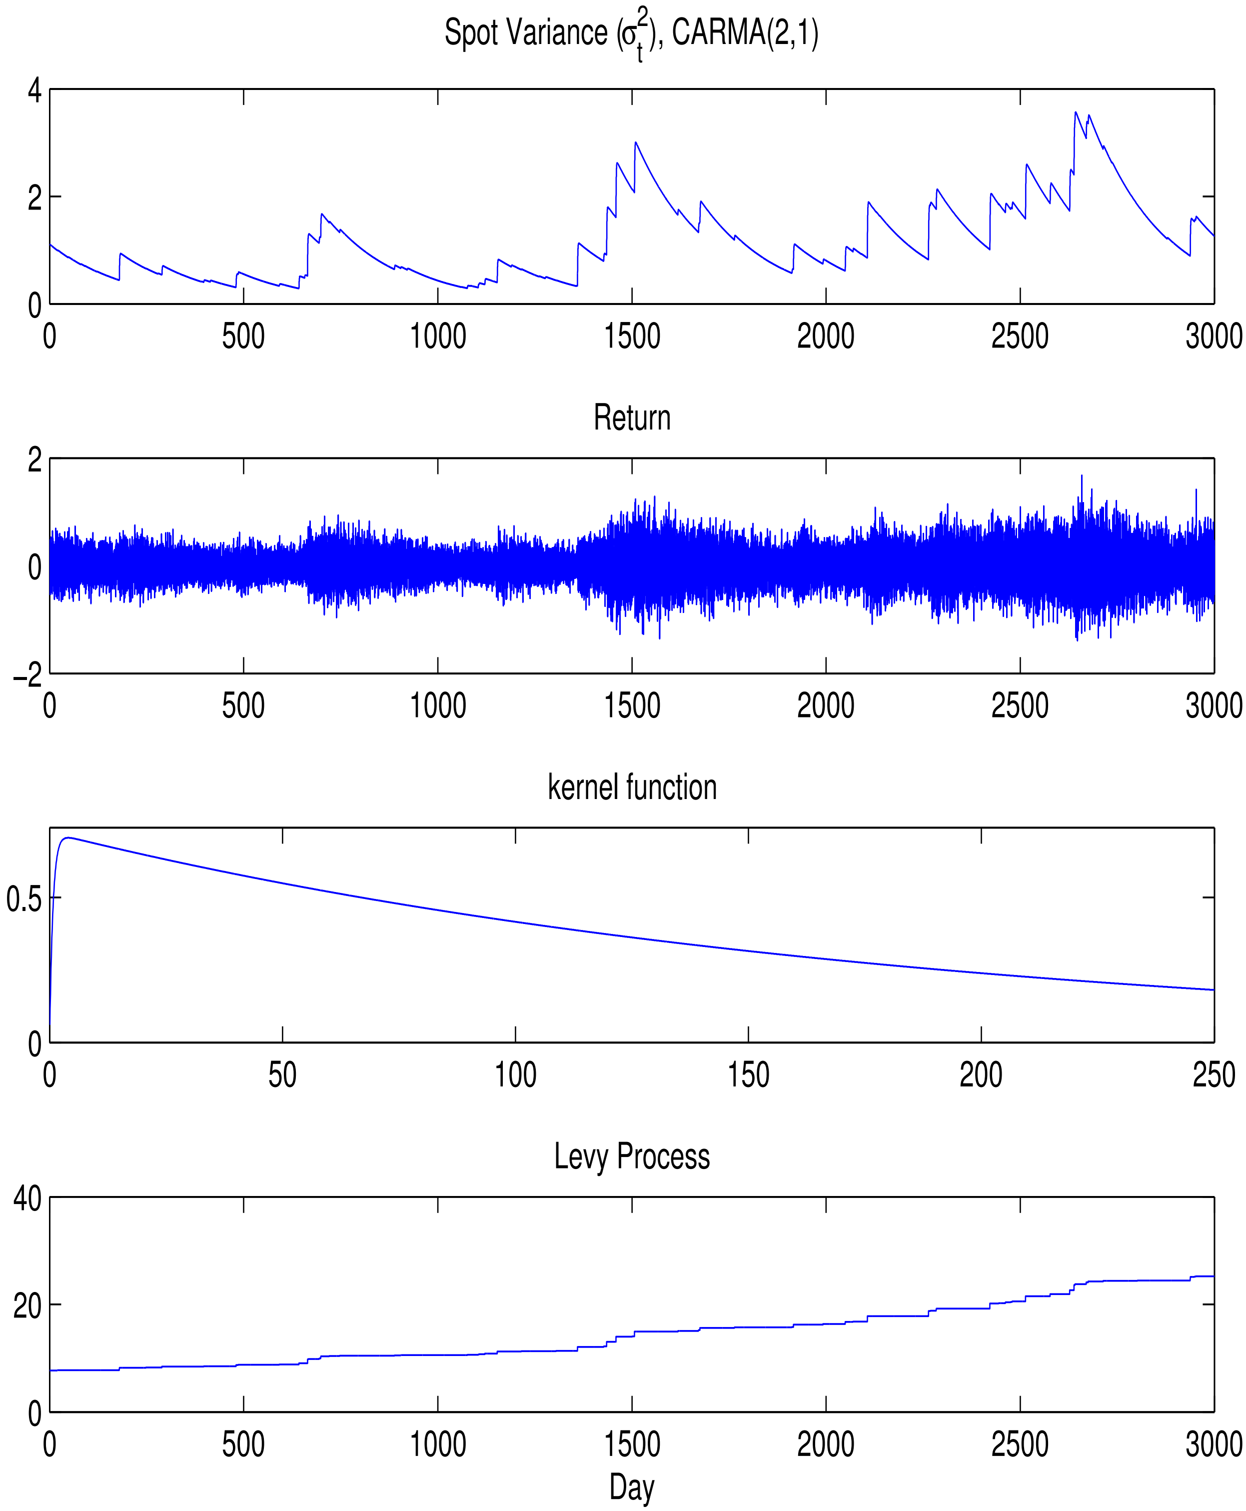

Figure 2: Simulated Realization of a CARMA(2,1) stochastic volatility model. The top panel shows the spot volatility at thirty minute intervals, eight hours per day; the second shows the return over the thirty minute intervals; the third shows the kernel function, and fourth the Lévy process at thirty minute intervals. Parameter settings shown in Table 1.

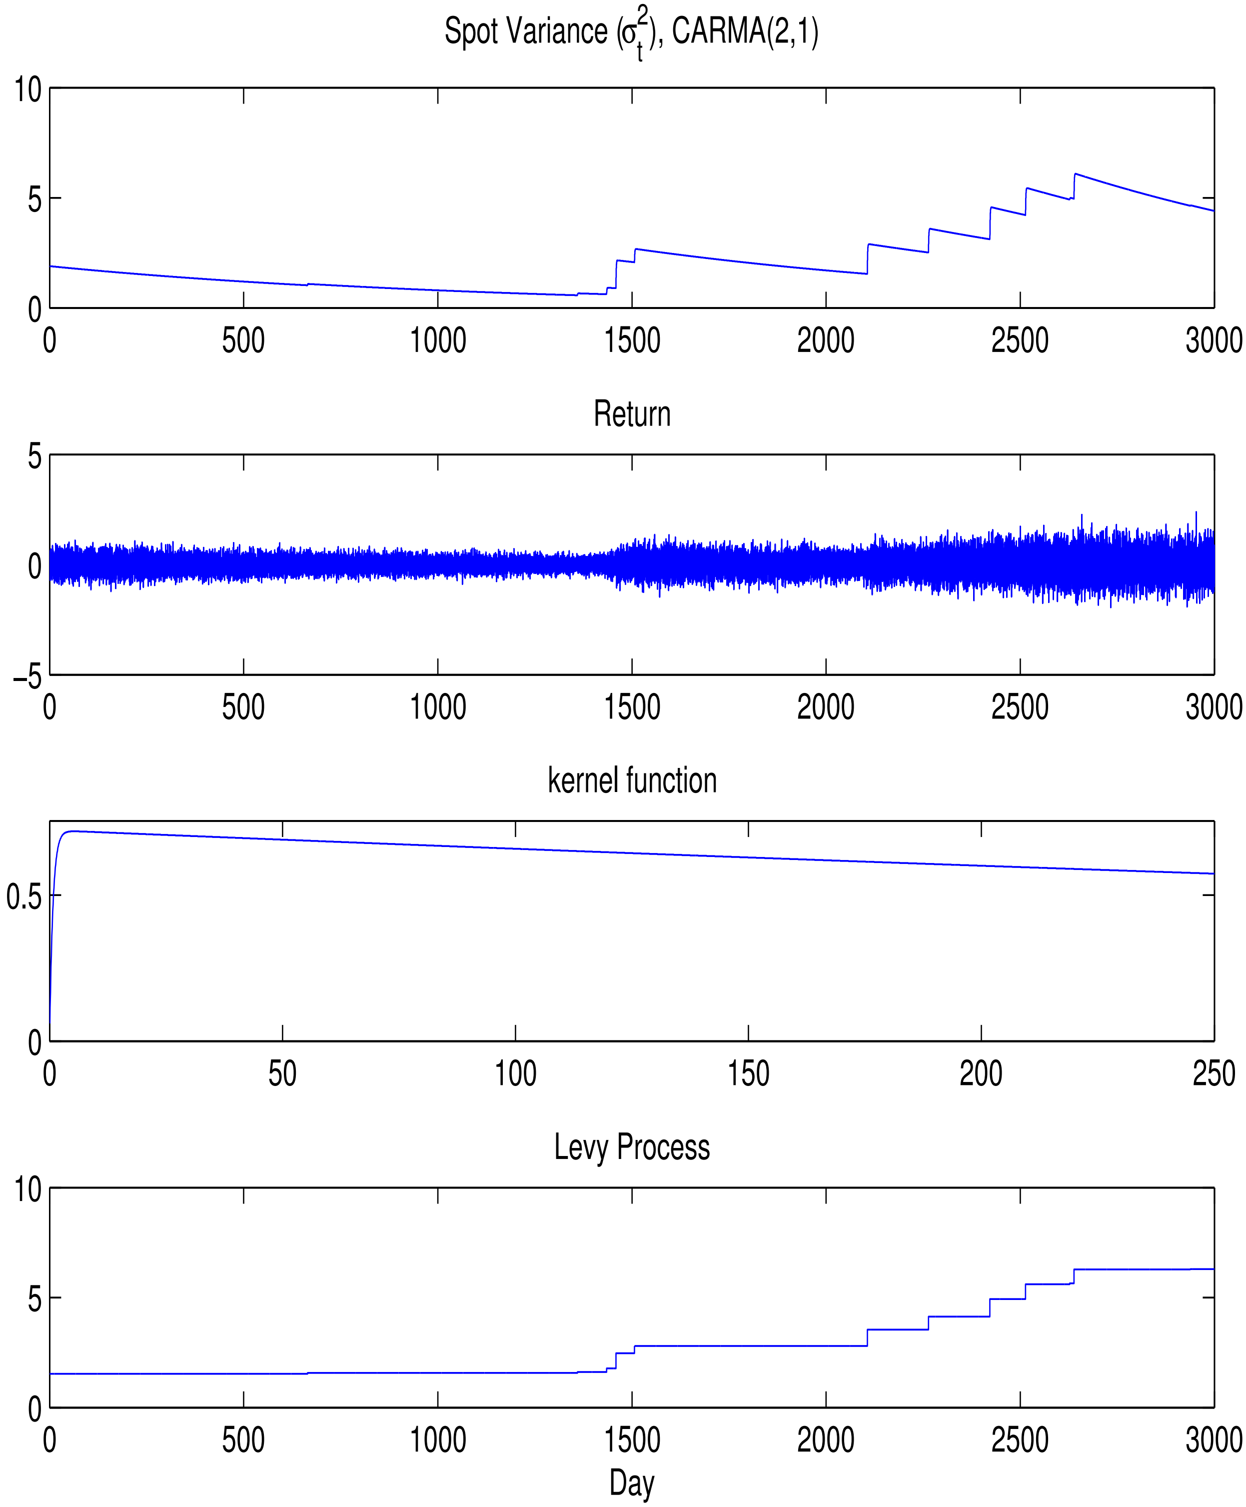

Figure 3: Simulated Realization of a CARMA(2,1) stochastic volatility model. The top panel shows the spot volatility at thirty minute intervals, eight hours per day; the second shows the return over the thirty minute intervals; the third shows the kernel function, and fourth the Lévy process at thirty minute intervals. Parameter settings shown in Table 1.

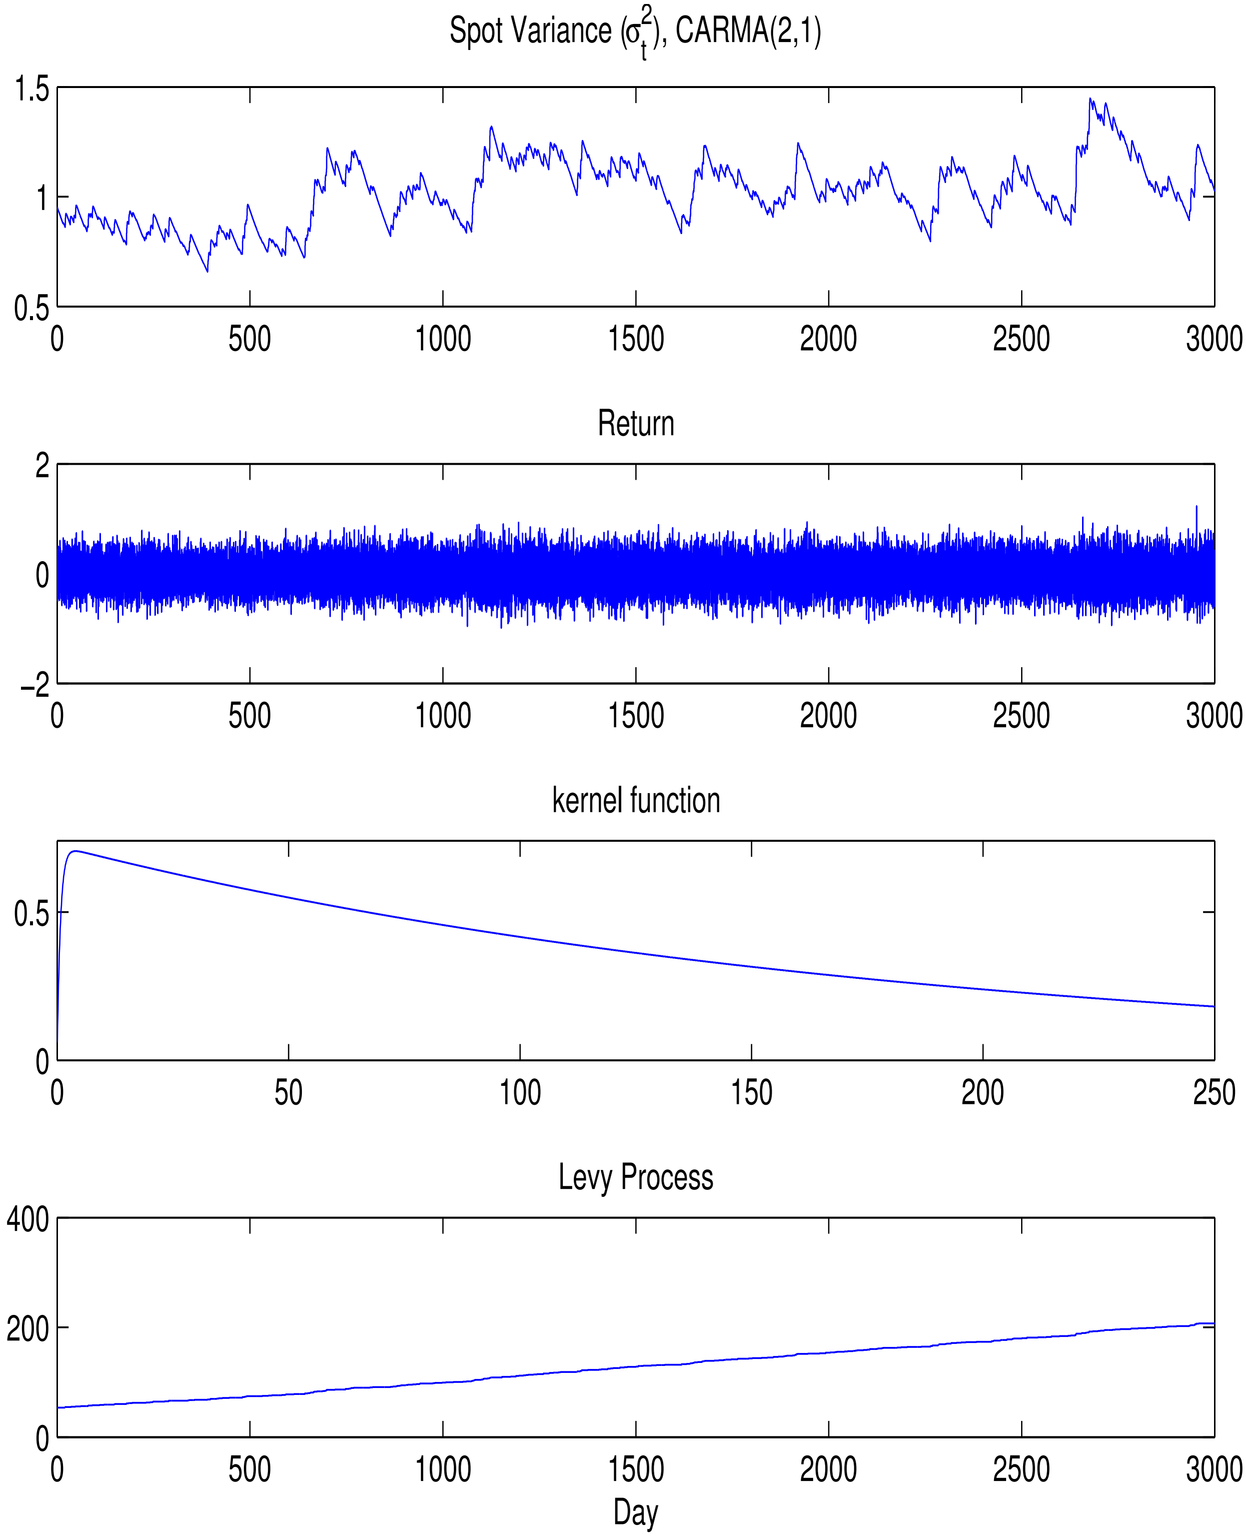

Figure 4: Simulated Realization of a CARMA(2,1) stochastic volatility model. The top panel shows the spot volatility at thirty minute intervals, eight hours per day; the second shows the return over the thirty minute intervals; the third shows the kernel function, and fourth the Lévy process at thirty minute intervals. Parameter settings shown in Table 1.

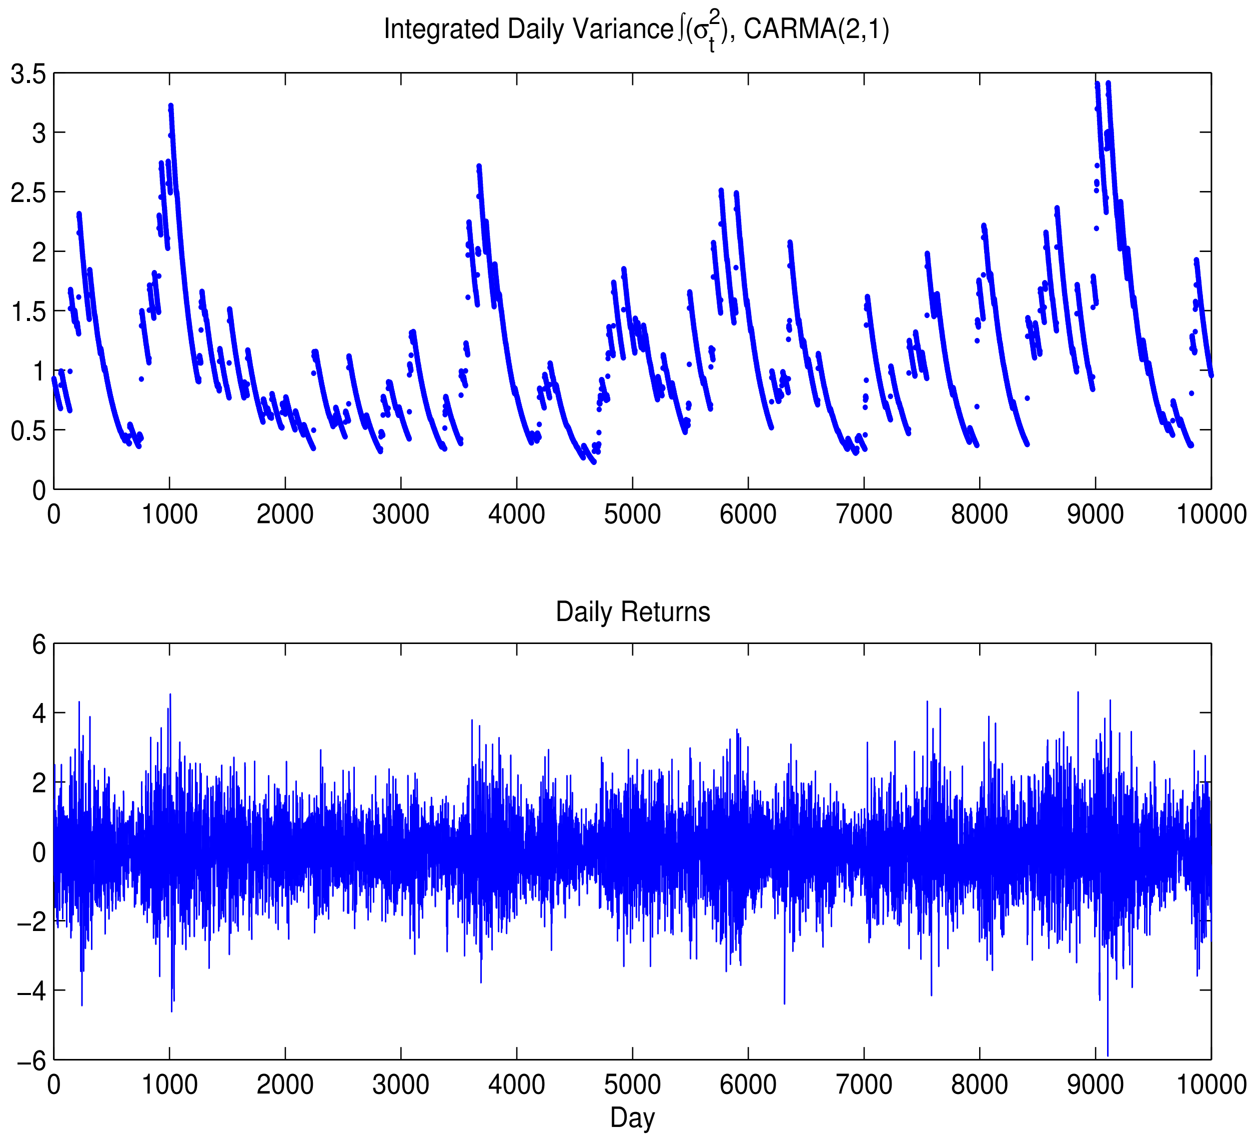

Figure 5: Simulated Daily Realizations Implied by a CARMA(2,1) model. The top panel shows the daily integrated volatility; the bottom shows the daily price change. Parameter settings as per Figure 5.

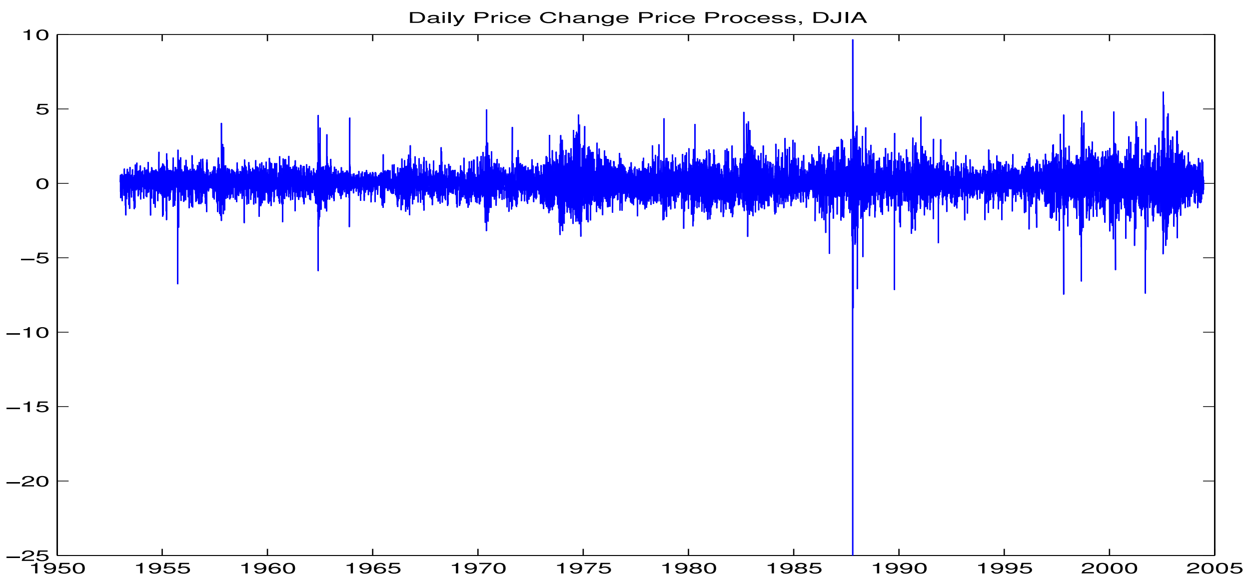

Figure 6: Observed Daily Realizations on the Dow Jones Return, $100 * (\log(p_t) - \log(p_{t-1})$, January 2, 1953 – June 22, 2004.

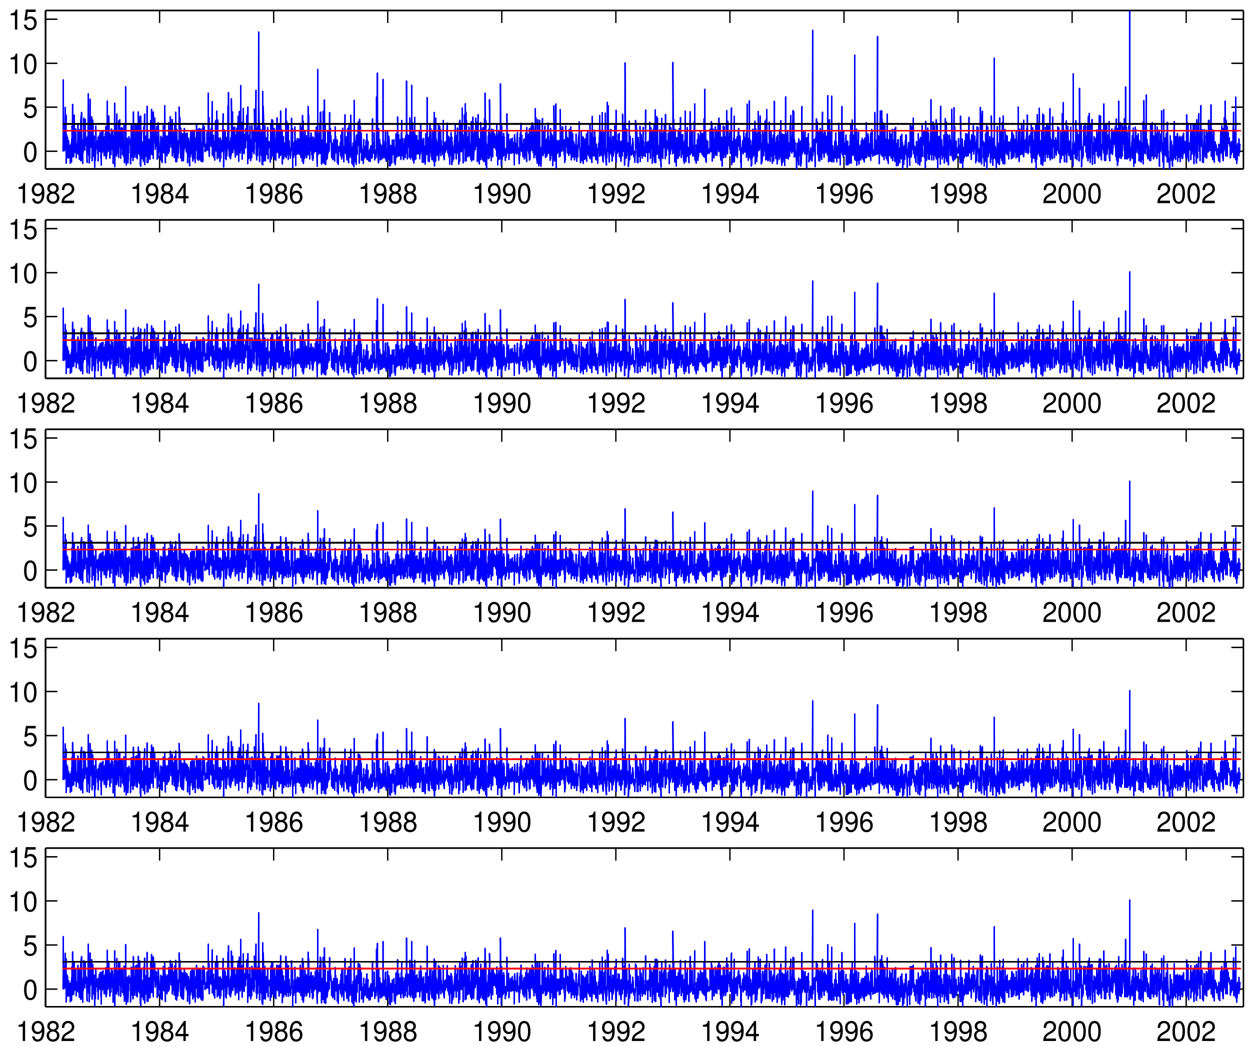

Figure 7: The five panel are the time series plots of the observed values of the daily statistics $z_{TP,t}$, $z_{TP,l,t}$, $z_{TP,lm,t}$, $z_{TP,r,t}$ and $z_{TP,rm,t}$, computed using the five-minute returns on the S&P Index futures, April 21, 1982–December 9, 2002. The horizontal lines are the upper 0.99 and 0.999 critical values of the standard normal distribution.

8 References

Ait-Sahalia, Y. (2003) “Disentangling Diffusion from Jumps,” *Journal of Financial Economics* forthcoming.

Ait-Sahalia, Y., Per Mykland, and Lan Zhang (2004) “How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise,” forthcoming in the *Review of Financial Studies*.

Alizadeh, S., M. Brandt, and F.X. Diebold (2002) “Range-Based Estimation of Stochastic Volatility Models.” *Journal of Finance* 57, 1047-1092.

Andersen, Torben (2004) “Discussion of: Power and Bipower Variation with Stochastic Volatility and Jumps, *Journal of Financial Econometrics* 2, 37–48.

Andersen, T. G., L. Benzoni, and J. Lund (2002) “An Empirical Investigation of Continuous-time Equity Return Models.” *Journal of Finance* 57, 1239-1284.

Andersen, Torben G. and Tim Bollerslev (1998) “Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts,” *International-Economic-Review* 39, 885-905.

Andersen, T. G., T. Bollerslev, and F. X. Diebold (2002) “Parametric and Nonparametric Volatility Measurement.” in L. P. Hansen and Y. Aït-Sahalia, eds., *Handbook of Financial Econometrics*, North Holland, Amsterdam, forthcoming.

Andersen, T.G., T. Bollerslev, and F. Diebold (2003) “Some Like it Smooth, and Some Like it Rough: Untangling Continuous and Jump Components in Measuring, Modeling and Forecasting Asset Return Volatility.” Unpublished Paper: Economics Department, Duke University.

Andersen, Torben G., Tim Bollerslev, and Nour Meddahi (2003) “Analytic Evaluation of Volatility Forecasts,” *International Review*, forthcoming.

Andreou, Elena, and Eric Ghysels (2002) “Rolling Sample Volatility Estimators: Some New Theoretical, Simulation, and Empirical Results,” *Journal of Business and Economic Statistics* 20, 363–76.

Bai, Xuezheng, Jeffrey Russell, and George Tiao (2004) “Effects of Non-Normality and Dependence on the Precision of Variance Estimates Using High-Frequency Financial Data,” Working Paper, University of Chicago.

Bandi, Federico M., and Jeffrey R. Russell (2004) "Separating Microstructure Noise from Volatility," working paper, Universit of Chicago.

Bollerslev, Tim, and Hao Zhou (2002) "Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility," *Journal of Econometrics* 109, 33-65.

Barndorff-Nielsen, Ole E., Elisa Nicolato, and Neil Shephard (2002) "Some Recent Developments in Stochastic Volatility Modelling" *Quantitative Finance* 2002, 2, 11–23.

Barndorff-Nielsen, Ole E., and Neil Shephard (2001a) "Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics", *Journal of the Royal Statistical Society, Series B, Methodological* 63, 167-241.

Barndorff-Nielsen, Ole E., and Neil Shephard (2001b) "Modelling by Lévy Processess for Financial Econometrics," in Ole E. Barndorff-Nielsen, Thomas Mikosch and Sid Resnick (eds.), *Lévy Processes — Theory and Application* New York: Birkhauser, 283–318.

Barndorff-Nielsen, Ole E., and Neil Shephard (2002a) "Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society, Series B, Methodological* 64, 253-280.

Barndorff-Nielsen, Ole E., and Neil Shephard (2004a) "Power and Bipower Variation with Stochastic Volatility and Jumps," *Journal of Financial Econometrics* 2, 1-37.

Barndorff-Nielsen, Ole E., and Neil Shephard (2004b) "Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation," manuscript, www.levy
process.org.

Barndorff-Nielsen, Ole E., and Neil Shephard (2004c) "Impact of Jumps on Returns and Realised Variances: Econometric Analysis of Time-Deformed Lévy Processes," manuscript, www.levy.process.org.

Barndorff-Nielsen, Ole E., and Neil Shephard (2004d) "A Feasible Central Limit Theory for Realised Volatility Under Leverage," manuscript, www.levy
process.org.

Barndorff-Nielsen, Ole E., and Neil Shephard (2004e) "Econometric Analysis of Realised Covariation: High Frequency Based Covariance, Regression and Correlation in Financial Economics," *Econometrica* 72, 885–925.

Barndorff-Nielsen, Ole E., and Neil Shephard (2004f) "Estimating Quadratic Variation using Realised Variance," *Journal of Applied Econometrics* 17, 457–477.

Barndorff-Nielsen, Ole E., Bent Nielsen, Carla Ysusi, and Neil Shephard (2004) "Measuring and Forecasting Financial Variability Using Realised Variance With and Without a Model, " in *Identification and Inference for Econometric Models. A Festschrift for Tom Rothenberg* (edited by Don Andrews, Jim Powell, Paul Ruud and Jim Stock), Cambridge University Press.

Beresteanu, Arie (2004) "Nonparametric estimation of regression functions under restrictions on partial derivatives," Working Paper, Duke University.

Bertoin, Jean (1996) "Lévy Processes," Cambridge, UK: Cambridge University Press.

Brockwell, Peter and R. Davis (2001) "Comment on Barndorff-Nielsen, O. E. and N. Shephard," *Journal of the Royal Statistical Society, Series B, Methodological* 64.

Brockwell, Peter (2003) "Representations of Continuous Time ARMA Processes," *Journal*

Carr, Peter, Hlyette Geman, Dilip B. Madan, and Marc Yor (2002), "The Fine Structure of Asset Returns," *Journal of Business* 75, 305–332.

Carr, Peter Geman, Hlyette , Madan, Dilip B. , and Yor, Marc (2003), "Stochastic Volatility for Lévy Processes", *Mathematical Finance,* 13, 345–382.

Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen (2003) "Alternative Models for Stock Price Dynamics," *Journal of Econometrics* 116, 225–258.

Clark, Peter (1973) "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices." *Econometrica* 41, 135–155.

Cont, R. and P. Tankov (2004) *Financial Modelling With Jump Processes*, Boca Raton, Fla, USA: Chapman & Hall.

Dai, Qiang and Kenneth J. Singleton (2000) "Specification Analysis of Affine Term Structure Models," *Journal of Finance* 55, 1943-78.

Duffie, Darrell and Rui Kan (1996) "A Yield-Factor Model of Interest Rates," *Mathematical-Finance* 6, 379-406.

Duffie, D., J. Pan, and K. Singleton (2000) "Transform Analysis and Asset Pricing for Affine Jump-Diffusions." *Econometrica* 68, 1343-76

Durham, Garland B. and A. Ronald Gallant (2002) "Numerical Techniques for Maximum Likelihood Estimation of Continuous-Time Diffusion Processes," *Journal-of-Business-and-Economic-Statistics* 20, 297-316.

Elerain, Ola, Siddhartha Chib and Neil Shephard (2001) "Likelihood Inference for Discretely Observed Nonlinear Diffusions," *Econometrica* 2001 69 959-93.

Eraker, Bjorn (2001) "MCMC Analysis of Diffusion Models with Application to Finance," *Journal of Business and Economic Statistics* 19, 177-91.

Eraker, Bjorn (2004) "The Performance of Model Based Option Trading Strategies," Working paper, Duke University.

Eraker, Bjorn, Michael Johannes, and Nicholas Polson (2003) "The Impact of Jumps in Volatility and Returns," O: *Journal of Finance* 58, 1269-1300

Gallant, A. R., C. T. Hsu, and G. Tauchen (1999) "Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance." *The Review of Economics and Statistics* 81(4), 617-631.

Gallant, A. R. and G. Tauchen (2002) "Simulated Score Methods and Indirect Inference for Continuous Time Models." in L. P. Hansen and Y. Aït-Sahalia, eds., *Handbook of Financial Econometrics*, North Holland, Amsterdam, forthcoming.

Gourieroux C., Jasiak, J., and R Sufana (2004) "The Wishart Autoregressive Process of Multivariate Stochastic Volatility," Working Paper, York University.

Hausman, J. A. (1978) "Specification Tests in Econometrics," *Econometrica* 46, 1251-71.

Heston, Steven L (1993) "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies* 6, 327-43.

Huang, Xin and George Tauchen (2004) "The Relative Contribution of Jumps to Total Price Variance," Working Paper, Duke University.

Kloeden, P. E. and E. Platen (1992) *Numerical Solution of Stochastic Differential Equations*, New York: Springer Verlag.

Li, Haitao, Martin Wells, and Long Yu (2004) "A MCMC Analysis of Time-Changed Levy Processes of Stock Return Dynamics," Working Paper, Cornell University.

Merton, R. C. (1976) "Option Pricing When Underlying Stock Returns Are Discontinuous." *Journal of Financial Economics* 3, 125-44.

Meddahi, Nour (2001) "An Eigenfunction Approach for Volatility Modeling," Working Paper, CIRANO.

Meddahi, Nour (2002) "Moments of Continuous Time Stochastic Volatility Models," Working Paper, CIRANO.

Rosinski, Jan (2001) "Series representations of Levy processes from the perspective of point processes," In Levy Processes - Theory and Applications. O.E. Barndorff-Nielsen, T. Mikosch and S.I. Resnick, Eds., Birkhauser, Boston, 401–415.

Shephard, Neil (2004) *Stochastic Volatility: Selected Readings*, Oxford University Press. Expected 2004.

Tauchen, George (2004) "Stochastic Volatility in General Equilibrium," Working Paper, Duke University.

Taylor, S.J. (1982) "Financial Returns Modelled by the Product of Two Stochastic Processes – A Study of Daily Sugar Prices 1961-79." *Time Series Analysis: Theory and Practice 1*, 203-226, Amsterdam: North-Holland.

Todorov, Viktor (2004) "PhD Thesis Chapter in Preparation," Duke University.

Taylor, S.J. (1986) *Modelling Financial Time Series*, Chichester: John Wiley and Sons.