

A comment on Christoffersen, Jacobs and Ornathanalai (2012),
“Dynamic jump intensities and risk premiums: Evidence from
S&P500 returns and options”

Garland Durham¹

John Geweke²

Pulak Ghosh³

February 25, 2014

¹garland.durham@gmail.com; Quantos Analytics, LLC and California Polytechnic State University.

²John.Geweke@uts.edu.au; University of Technology Sydney (Australia), Erasmus University Rotterdam (Netherlands) and Colorado State University (USA). Geweke acknowledges partial financial support from Australian Research Council grants DP110104732 and DP130103356.

³pulak.ghosh@iimb.ernet.in; Indian Institute of Management, Bangalore.

Abstract

Christoffersen, Jacobs, and Ornathanalai (2012) propose an interesting and useful class of GARCH-like models with dynamic jump intensity, and find evidence that the models not only fit returns data better than some commonly used benchmarks but also provide substantial improvements in option pricing performance. While such models pose difficulties for estimation and analysis, CJO propose an innovative approach to filtering intended to address them. However, there are some statistical issues that their approach leaves unresolved, with implications for the option pricing results. This note proposes a solution based on using the filter and estimator proposed by CJO, but interpreted in the context of an alternative model. With respect to this model, the estimator is consistent and likelihood-based model comparisons and hypothesis tests are valid.

JEL classification: G13

Keywords: Compound Poisson jumps; Risk premiums; Analytical filtering; Dynamic jump intensity

1. Introduction

There has been much interest in models for asset returns that include dynamic jump intensity, going back to seminal work by Chan and Maheu (2002) and Maheu and McCurdy (2004). In more recent work, Rangel (2011) examines the effects of news events on jump intensity, and Christoffersen, Jacobs, and Ornathanalai (2012) and Santa-Clara and Yan (2010) find that dynamic jump intensity plays an important role in option pricing. In a somewhat different vein, Wright and Zhou (2009) find that evidence extracted from high-frequency stock index returns supports the premise of time-variation in jump mean, variance and intensity, and that jump variance (but not intensity) has strong predictability for excess bond returns. Ait-Sahalia, Cacho-Diaz, and Laeven (2013) suggest an innovative modeling framework with origins in epidemiology to explain the presence of time-varying jump intensity based on mutually exciting jump processes (Hawkes processes).

Christoffersen, Jacobs, and Ornathanalai (2012) (CJO hereafter) propose an interesting and useful class of GARCH-like models with dynamic jump intensity, and find evidence that the models not only fit returns data better than some commonly used benchmarks but also provide substantial improvements in option pricing performance. In the model of primary interest in that paper, the returns process is driven by two dynamic state variables: one of these is closely related to standard GARCH volatility factors, while the other generates time-varying jump intensities. We shall refer to this as the GARCH dynamic jump intensity (GARCH-DJI) model.

While such models pose difficulties for estimation and analysis, CJO propose an innovative approach to filtering that addresses these issues. However, the filter they propose lacks a key property: the filtered states are not equal to their expected values conditional on the relevant information set. Furthermore, while CJO refer to the estimator they propose as a maximum likelihood estimator (MLE), it is not the MLE for the GARCH-DJI model they study.

It is, however, possible to construct a model that represents a well-defined data generating process based on their filtering algorithm. We refer to this as the FILTER-DJI model. For this model, the states are trivially identified; and it is this model for which the estimator proposed by CJO is in fact the MLE.

In this note, we investigate some characteristics of GARCH-DJI and FILTER-DJI models. We find that for a given parameter vector the two classes of models represent similar data generating processes, but that there are some clear differences. When applied to simulated data generated from a GARCH-DJI model, we show that the filter proposed by CJO is biased. The estimator they propose is also biased, and hypothesis tests have incorrect

size. In a Monte Carlo study, we find that the coverage ratio of a nominal 95% confidence region for the model parameter vector is only 39%. Furthermore, the log likelihood values reported by CJO and used for model comparisons are not valid.

With respect to the FILTER-DJI model, in contrast, the estimator maintains the usual attractive properties of maximum likelihood estimation. In particular, when interpreted in the context of this model (rather the GARCH-DJI model actually proposed by CJO), parameter estimates (including those reported by CJO) are consistent, hypothesis tests are correctly sized, and model comparison results are valid.

There are implications for option pricing. The approach proposed by CJO involves simulating option prices conditional on an estimated parameter vector and filtered states. But while the estimator is consistent for the parameters of the FILTER-DJI model and the filter is unbiased (conditional on a parameter vector), neither of these properties hold for GARCH-DJI, the model under which option prices are simulated. Option prices implied by the two models are close at short times to maturity, but diverge as time to maturity increases.

There is a more fundamental issue underlying the one involved here that is a subtle but important one for economists working with state space models. Even if a filter that provides unbiased estimates for the states is available (a condition not satisfied by the GARCH-DJI model), treating the extracted states as known and using them for maximum likelihood estimation will not in general yield a consistent estimator or correct values for the log likelihood. For valid statistical analysis, either the state uncertainty must be integrated out or the model reformulated in such a way as to eliminate it (in special cases, such as the linear Gaussian state space model, the integral may be computable analytically). This point has been noted previously in a related setting by Fleming and Kirby (2003).

Since the class of models proposed by CJO, and others like it, have properties that are of considerable interest for applied work, it is important that subsequent work have a solid theoretical foundation to build upon. The FILTER-DJI model proposed here provides one possible workaround for the estimation issues left unresolved by that paper. While the estimation results and model comparisons reported by CJO are not valid in the context of the GARCH-DJI model, they are valid if interpreted in the context of the FILTER-DJI model instead. This note thus provides a constructive solution that reaffirms the usefulness of CJO's empirical findings and helps open the way for further research building on their work.

2. Two dynamic jump intensity models

The GARCH-DJI model proposed by CJO is given by

$$R_t = \mu_t + z_t + y_t \quad (1)$$

$$\mu_t = r_t + (\lambda_z - \frac{1}{2})h_{zt} + (\lambda_y - \xi)h_{yt} \quad (2)$$

where R_t is a log return, λ_z and λ_y are risk premia, r_t is the risk-free rate, $z_t \sim N(0, h_{zt})$, and y_t is a Poisson jump process with intensity h_{yt} , mean μ_J and variance σ_J^2 . The dynamics of variance (h_{zt}) and jump intensity (h_{yt}) are given by

$$h_{z,t+1} = w_z + b_z h_{zt} + \frac{a_z}{h_{zt}} (z_t - c_z h_{zt})^2 + d_z (y_t - e_z)^2 \quad (3)$$

$$h_{y,t+1} = w_y + b_y h_{yt} + \frac{a_y}{h_{zt}} (z_t - c_y h_{zt})^2 + d_y (y_t - e_y)^2 \quad (4)$$

with initial conditions h_{z0} and h_{y0} . The terms $h_{zt}/2$ and $\xi h_{yt} = (e^{\mu_J + \sigma_J^2/2} - 1)h_{yt}$ in (2) are convexity adjustments. The full model has parameter vector $\theta = (\lambda_z, \lambda_y, \mu_J, \sigma_J, a_z, b_z, c_z, d_z, e_z, w_z, a_y, b_y, c_y, d_y, e_y, w_y)$.

Let n_t denote the number of jumps at time t . Then $(R_t | r_t, h_{zt}, h_{yt}, n_t) \sim N(\mu_t + n_t \mu_J, h_{zt} + n_t \sigma_J^2)$ and $(n_t | h_{yt}) \sim \text{Poisson}(h_{yt})$. Integrating across n_t , $R_t | r_t, h_{zt}, h_{yt}$ is a mixture of normals with density

$$p(R_t | r_t, h_{zt}, h_{yt}) = \sum_{j=0}^{\infty} p(j | h_{yt}) \phi(R_t | \mu_t + j \mu_J, h_{zt} + j \sigma_J^2), \quad (5)$$

where $p(j | h_{yt})$ is the Poisson(h_{yt}) density and ϕ is the Gaussian density. For future reference, note that

$$p(n_t | R_t, r_t, h_{zt}, h_{yt}) \propto p(R_t | r_t, h_{zt}, h_{yt}, n_t) p(n_t | h_{yt}) \quad (6)$$

by Bayes rule.

In practice, n_t , z_t , y_t , h_{zt} and h_{yt} are not observable. To estimate the model, CJO

propose the “filter”,

$$\tilde{\mu}_t = r_t + (\lambda_z - \frac{1}{2})\tilde{h}_{zt} + (\lambda_y - \xi)\tilde{h}_{yt} \quad (7)$$

$$\tilde{z}_t = \sum_{j=0}^{\infty} \frac{\tilde{h}_{zt}}{\tilde{h}_{zt} + j\sigma_j^2} (R_t - \tilde{\mu}_t - j\mu_J) p(n_t = j | R_t, r_t, \tilde{h}_{zt}, \tilde{h}_{yt}) \quad (8)$$

$$\tilde{y}_t = R_t - \tilde{\mu}_t - \tilde{z}_t \quad (9)$$

$$\tilde{h}_{z,t+1} = w_z + b_z \tilde{h}_{zt} + \frac{a_z}{\tilde{h}_{zt}} (\tilde{z}_t - c_z h_{zt})^2 + d_z (\tilde{y}_t - e_z)^2 \quad (10)$$

$$\tilde{h}_{y,t+1} = w_y + b_y \tilde{h}_{yt} + \frac{a_y}{\tilde{h}_{zt}} (\tilde{z}_t - c_y h_{zt})^2 + d_y (\tilde{y}_t - e_y)^2, \quad (11)$$

with initial conditions \tilde{h}_{z0} and \tilde{h}_{y0} . CJO show that (1)–(2) with $h_{zt} = \tilde{h}_{zt}$ and $h_{yt} = \tilde{h}_{yt}$ implies that $\tilde{z}_t = E(z_t | R_t, r_t, \tilde{h}_{zt}, \tilde{h}_{yt})$ and $\tilde{y}_t = E(y_t | R_t, r_t, \tilde{h}_{zt}, \tilde{h}_{yt})$.

Using this filter, it is straightforward to back out implied values of \tilde{h}_{zt} and \tilde{h}_{yt} conditional on data $\{R_t, r_t\}_{t=1}^{t-1}$, parameter vector θ and initial conditions \tilde{h}_{z0} and \tilde{h}_{y0} . The parameter vector is then estimated by optimization,

$$\hat{\theta} = \operatorname{argmax}_{\theta \in \Theta} \sum_{t=1}^T p(R_t | r_t, \tilde{h}_{zt}, \tilde{h}_{yt}, \theta), \quad (12)$$

where the summands are given by (5) but with \tilde{h}_{zt} and \tilde{h}_{yt} in place of h_{zt} and h_{yt} . CJO refer to this as a maximum likelihood estimator (MLE).

In fact, $\hat{\theta}$ is the maximum likelihood estimator not for the GARCH-DJI model (1)–(4), but instead for the model defined by (7)–(11) in conjunction with

$$R_t = \tilde{\mu}_t + \epsilon_t \quad (13)$$

where ϵ_t is a mixture of normals with density $\sum_{j=0}^{\infty} p(j | \tilde{h}_{yt}) \phi(\cdot, j\tilde{\mu}_J, \tilde{h}_{zt} + j\tilde{\sigma}_j^2)$ and $p(j | \tilde{h}_{yt})$ is the Poisson(\tilde{h}_{yt}) density. We shall refer to this as the FILTER-DJI model. It is motivated by GARCH-DJI but not equivalent to it.

While CJO show that $\tilde{z}_t = E(z_t | R_t, r_t, \tilde{h}_{zt}, \tilde{h}_{yt})$ and $\tilde{y}_t = E(y_t | R_t, r_t, \tilde{h}_{zt}, \tilde{h}_{yt})$, it does not follow that $\tilde{h}_{zt} = E(h_{zt} | \mathcal{F}_{t-1})$ or that $\tilde{h}_{yt} = E(h_{yt} | \mathcal{F}_{t-1})$, where \mathcal{F}_t is the σ -algebra generated by $\{R_\tau, r_\tau\}_{\tau=1}^t$. So (7)–(10) does not possess a key feature typically desired of a filter. And while (12) is the MLE for FILTER-DJI, it is not the MLE for GARCH-DJI. Furthermore, even if the state filter were unbiased, (12) would not be the MLE for GARCH-DJI. Computing the likelihood for this model requires integrating across state uncertainty; a plug-in estimate of the states, even an unbiased one, is not sufficient. The properties of

this estimator with respect to the parameters of the GARCH-DJI model are unknown.

3. Findings

3.1. Comparison of true and filtered states

This section investigates the extent to which the filtered states \tilde{h}_{zt} and \tilde{h}_{yt} are informative about the true states h_{zt} and h_{yt} . The experiment performed here uses simulated data generated using the GARCH-DJI model with the parameter vector reported by CJO in Table 1 of that paper. We simulate one million observations with a burn-in period of 1000 observations to minimize the effects of initial conditions. The filter (7)–(11) is then applied to the simulated returns to extract filtered states \tilde{h}_{zt} and \tilde{h}_{yt} and innovations \tilde{z}_t and \tilde{y}_t . Since the data are simulated, the true states h_{zt} and h_{yt} and innovations z_t and y_t , which are latent in empirical applications, can be observed. Thus comparison of the true and filtered states and innovations is possible.

Table 1 reports summary statistics for the true and filtered states and innovations. Relative to the true states, the filtered states \tilde{h}_{zt} and \tilde{h}_{yt} are systematically biased downwards, have smaller standard deviation, and are less skewed and less leptokurtic.

3.2. Comparison of true and estimated parameters

This section reports the results of a Monte Carlo study investigating the issue of potential bias in the estimator (12) with respect to the parameters of the GARCH-DJI model. Each replication in the study involves generating $N = 11,979$ observations of simulated data (equal to the sample size used in the application provided by CJO) using the GARCH-DJI model with the parameters from Table 1 of CJO (as described above), and then estimating the model using (12) (as proposed by CJO). We perform 1000 replications and report the bias and RMSE of the resulting parameter estimates. For comparison, we then repeat this procedure using data generated from the FILTER-DJI model with the same parameter vector. In this case, (12) represents the true MLE.

In the full model, some of the parameters are difficult to pin down accurately, with correlations implied by estimates of the asymptotic covariance matrix as high as 0.995. This is analogous to the situation of near-multicollinearity in linear regression, and results in very large standard errors. In addition to being very flat in some dimensions, the likelihood surface has multiple local maxima, making global optimization problematic. For the results reported here, we treat the parameters, d_z , e_z , d_y and e_y —for which the likelihood function is particularly uninformative—as given and focus on the remaining parameters, which largely

alleviates these issues.

Table 2 reports bias, standard deviation and RMSE of parameter estimates across the 1000 replications. When applied to data generated from the GARCH-DJI model, the estimator shows evidence of bias with a magnitude equal to about one standard error for several parameters, notably a_z , c_z , a_y , μ_J , and σ_J . No evidence of bias is apparent when the estimator is applied to data generated from the FILTER-DJI model.

The columns of the table labeled “CR” report the coverage rate of the nominal 95% confidence interval for each parameter. Confidence intervals are slightly conservative for the FILTER-DJI model (that is, for most parameters the 95% confidence interval contains the true value just over 95% of the time). For data generated from GARCH-DJI, in contrast, coverage rates are substantially lower than 95% for several parameters (notably, 71.2% for μ_J). The results are approximately as expected: recall that for a normally distributed estimator with known variance, a bias of one standard error implies that the 95% confidence interval fails to include the true parameter value about 16% of the time.

A multivariate confidence region can be constructed by inverting the likelihood ratio statistic (analogous to the construction of confidence regions in multivariate linear regression). Let $U(\mathbf{X}) = \{\theta : 2L(\hat{\theta}; \mathbf{X}) - 2L(\theta; \mathbf{X}) < F_{\text{crit}}\}$, where L is the likelihood function, $\hat{\theta}$ is an estimator, \mathbf{X} denotes data, and F_{crit} is the appropriate critical value for a chi-squared distribution with degrees of freedom equal to the number of estimated parameters. For example, in this model there are 12 free parameters, so for a 95% confidence region $F_{\text{crit}} = P^{-1}(0.95; 12) = 21.03$, where $P(\cdot; k)$ indicates the CDF of a chi-squared random variable with k degrees of freedom. For a maximum likelihood estimator (assuming regularity conditions), U should contain the true parameter vector 95% of the time asymptotically. In the exercise undertaken here, the coverage rate was 94.9% for data generated from the FILTER-DJI model and 39.0% for data generated using the GARCH-DJI model. For FILTER-DJI, the results are consistent with theory; in contrast, operating under the (false) supposition (maintained by CJO) that (12) is the MLE for GARCH-DJI entails a substantial risk of Type I errors.

4. Conclusions

CJO provide evidence that including dynamic jump intensity provides a significant improvement in fitting stock returns relative to some models without that feature. They also provide evidence that the models they propose perform better than some alternatives in explaining option prices.

But although the GARCH-DJI model is the primary object of interest in that paper, the

estimator they propose is the MLE for a different model, which we refer to as FILTER-DJI. The estimator is not consistent with respect to the parameters of the GARCH-DJI model, the associated filter is biased, and implied log likelihood values (and corresponding model comparisons) are invalid. These issues are resolved by the FILTER-DJI model proposed here.

References

- Aït-Sahalia, Y., Cacho-Diaz, J., Laeven, R. J. A., 2013. Modeling financial contagion using mutually exciting jump processes. Unpublished working paper.
- Chan, W. H., Maheu, J. M., 2002. Conditional jump dynamics in stock market returns. *Journal of Business and Economic Statistics* 20, 377–389.
- Christoffersen, P., Jacobs, K., Ornathanalai, C., 2012. Dynamic jump intensities and risk premiums: Evidence from S&P500 returns and options. *Journal of Financial Economics* 106, 447–472.
- Fleming, J., Kirby, C., 2003. A closer look at the relation between GARCH and stochastic autoregressive volatility. *Journal of Financial Econometrics* 1, 365–419.
- Maheu, J. M., McCurdy, T. H., 2004. News arrival, jump dynamics, and volatility components for individual stock returns. *Journal of Finance* 59, 755–793.
- Rangel, J. C., 2011. Macroeconomic news, announcements, and stock market jump intensity dynamics. *Journal of Banking and Finance* 35, 1263–1276.
- Santa-Clara, P., Yan, S., 2010. Crashes, volatility, and the equity premium: Lessons from S&P500 options. *Review of Economics and Statistics* 92, 435–451.
- Wright, J., Zhou, H., 2009. Bond risk premia and realized jump risk. *Journal of Banking and Finance* 33, 2333–2345.

Table 1

Summary statistics, true and filtered states and innovations using data simulated from the GARCH-DJI model with the parameter vector reported by CJO.

	True	Filtered
GARCH state (h_z)		
Mean \times 1000	0.0719	0.0678
SD \times 1000	0.0605	0.0504
Skewness	2.7662	2.2499
Kurtosis	15.1248	10.7819
Jump intensity (h_y)		
Mean \times 1000	23.3164	21.5456
SD \times 1000	29.3464	23.3113
Skewness	2.9781	2.3117
Kurtosis	16.3256	10.4829
GARCH innovation (z)		
Mean \times 1000	-0.0019	0.0356
SD \times 1000	8.4780	8.1350
Skewness	-0.0008	0.3308
Kurtosis	5.1227	4.4093
Jump innovation (y)		
Mean \times 1000	-0.4083	-0.4047
SD \times 1000	3.1034	2.1649
Skewness	-8.7148	-10.8684
Kurtosis	89.7829	148.9068

Table 2

Monte Carlo study for bias and RMSE of estimator. In each case, data was simulated using the indicated model with parameter vector θ_0 , and estimates were obtained using (12). Each experiment consisted of 1000 replications using simulated datasets of 11,979 observations each. The columns labeled “CR” report the coverage ratio of the nominal 95% confidence interval for each individual parameter.

	θ_0	GARCH-DJI				FILTER-DJI			
		Bias	SD	RMSE	CR(%)	Bias	SD	RMSE	CR(%)
$\lambda_z \times 0.1$	0.0506	-0.0363	0.2430	0.2456	97.1	-0.0088	0.2611	0.2611	97.4
$a_z \times 10^6$	2.2300	0.2006	0.2050	0.2867	83.9	-0.0069	0.1718	0.1719	96.6
b_z	0.9260	0.0032	0.0065	0.0072	88.3	-0.0010	0.0066	0.0067	95.8
$c_z \times 0.01$	1.3200	-0.1057	0.1229	0.1621	81.6	0.0116	0.1243	0.1248	95.7
$w_z \times 10^4$	-1.1600	-0.0014	0.0017	0.0022	85.6	0.0003	0.0016	0.0016	95.7
$\lambda_y \times 100$	0.5880	0.2820	0.7883	0.8369	96.5	0.0708	0.8489	0.8514	96.7
$a_y \times 1000$	0.8110	0.2106	0.2633	0.3371	97.6	0.0434	0.2255	0.2295	98.2
b_y	0.9730	-0.0035	0.0051	0.0062	94.8	-0.0021	0.0058	0.0062	97.8
$c_y \times 0.01$	0.0000	-0.0843	0.2641	0.2771	94.1	-0.0025	0.3077	0.3075	95.3
$w_y \times 10$	-0.5290	-0.0024	0.0028	0.0037	96.6	-0.0006	0.0022	0.0023	97.9
$\mu_J \times 100$	-1.7500	-0.3336	0.2617	0.4239	71.2	-0.0103	0.2816	0.2816	91.4
$\sigma_J \times 100$	0.9780	-0.2235	0.2479	0.3337	86.7	-0.0463	0.2208	0.2255	94.3

Coverage rate for θ_0
95% confidence region: 39.0%

Coverage rate for θ_0
95% confidence region: 94.9%