Published version http://dx.doi.org/10.4236/ojapps.2017.78030

http://hdl.handle.net/2440/111642

Scientific
Research
Publishing

Volatility in High-Frequency Intensive Care Mortality Time Series: Application of Univariate and Multivariate GARCH Models

John L. Moran[1*], Patricia J. Solomon[2,3]

[1]Department of Intensive Care Medicine, The Queen Elizabeth Hospital, Woodville, Australia
[2]School of Mathematical Sciences, University of Adelaide, Adelaide, Australia
[3]Australian and New Zealand Intensive Care Society, Carlton, Australia
Email: *john.moran@adelaide.edu.au, patty.solomon@adelaide.edu.au, anzics@anzics.com.au

Abstract

Mortality time series display time-varying volatility. The utility of statistical estimators from the financial time-series paradigm, which account for this characteristic, has not been addressed for high-frequency mortality series. Using daily mean-mortality series of an exemplar intensive care unit (ICU) from the Australian and New Zealand Intensive Care Society adult patient database, joint estimation of a mean and conditional variance (volatility) model for a stationary series was undertaken via univariate autoregressive moving average (ARMA, lags (p, q)), GARCH (Generalised Autoregressive Conditional Heteroscedasticity, lags (p, q)). The temporal dynamics of the conditional variance and correlations of multiple provider series, from rural/regional, metropolitan, tertiary and private ICUs, were estimated utilising multivariate GARCH models. For the stationary first differenced series, an asymmetric power GARCH model (lags (1, 1)) with t distribution (degrees-of-freedom, 11.6) and ARMA (7,0) for the mean-model, was the best-fitting. The four multivariate component series demonstrated varying trend mortality decline and persistent autocorrelation. Within each MGARCH series no model specification dominated. The conditional correlations were surprisingly low (<0.1) between tertiary series and substantial (0.4 - 0.6) between rural-regional and private series. The conditional-variances of both the univariate and multivariate series demonstrated a slow rate of time decline from periods of early volatility and volatility spikes.

1. Introduction

Mortality time series analyses in the biomedical literature traditionally utilise monthly or yearly aggregates [1], albeit log-linear (Poisson) approaches to the assessment of the effects of air-borne pollution report daily mortality [2]. The recent application of statistical process control (SPC) to monitor provider (for example intensive care unit, ICU) mortality has seen the use of EWMA (exponentially weighted moving average) charts to plot sequential patient admissions and progressively updated aggregate (mean) mortalities [3] [4]. The data generating process (DGP) of mortality series at this degree of temporal aggregation has not been appropriately characterised and would have implications for performance monitoring strategies such as residual-EWMA control charts, which we have previously advocated [5]. The latter study investigated the DGP of monthly ICU mortality time-series, which displayed autocorrelation, seasonality and (G)ARCH ((Generalised) Autoregressive Conditional Heteroscedasticity) effects. That is, the conditional variance of the time series random component (ϵ_t, or white noise) followed an autoregressive process with time varying volatility. In the financial time series literature, "volatility" is conventionally equated with (conditional) standard deviation [6] or (conditional) variance [7], albeit such focus has been subjected to critique [8] [9]. We now extend the previous perspective to daily mortality time series, which, for the current purpose, we will term "high-frequency" and draw inspiration from the paradigm of economic and financial time series [10] [11]. As opposed to financial time series, we do not consider intra-day events [12] [13] [14] on the basis that deaths within a "day" are relatively few in number and occur at irregular time intervals, precluding conventional time series analysis [15]. This being said, the stylised facts of financial "returns" $\left(\log\left(p_t/p_{t-1}\right)\right)$, where p_t is the asset price at time t, [16] have similarities with mortality time series [7].

We first undertake an analysis of the daily (mean) mortality of an exemplar ICU continuously contributing data (1996-2010) to the ANZICS (Australian and New Zealand Intensive Care Society) adult patient database [17]. In particular: characterisation of the raw series in terms of moments, auto-correlation and ARCH effects; specification of a mean equation and model to remove any linear dependence (for example, ARMA, autoregressive moving average); identification of residual ARCH effects and formulation of a volatility model (in this case, a (G) ARCH model [18]), and joint estimation of the mean and volatility equations [19]. Secondly, and more ambitiously, we undertake the joint analysis of multiple-provider series on the basis of presumed temporal dependencies [20]. Within a time series paradigm, and inheriting the insights of our first stage analysis, this modelling task presents itself in the domain of multivariate GARCH (MGARCH) models [21] [22], whereby the dynamics of conditional-variance and covariance of multiple provider-series are estimated; specifically the relations across series in the second order moment. In itself, this task is by no means facile, due to the attendant computational burden of the heavily pa-

rameterised MGARCH models [23].

2. Methods and Materials

2.1. Ethics Statement

Access to the data was granted by the ANZICS (Australian and New Zealand Intensive Care Society) Database Management Committee in accordance with standing protocols; local hospital (The Queen Elizabeth Hospital) Ethics of Research Committee waived the need for patient consent to use their data in this study. The data set analysed was anonymised before release to the authors by the ANZICS Centre for Outcome and Resource Evaluation (CORE) of the Australian and New Zealand Intensive Care Society (ANZICS), custodians of the database. The dataset is the property of the ANZICS Data base and contributing ICUs and is not in the public domain. Access to the data by researchers, submitting ICUs, jurisdictional funding bodies and other interested parties is obtained under specific conditions and upon written request ("ANZICS CORE Data Access and Publication Policy.pdf",

http://www.anzics.com.au/Downloads/ANZICS%20CORE%20Data%20Access%20and%20Publication%20Policy%20July%202017.pdf).

As previously described [5] [24], the ANZICS adult patient database [17] was utilised to define an appropriate patient set, 1996-(end)2010. Physiological variables collected in accordance with the requirements of the APACHE (Acute Physiology and Chronic Health Evaluation) III algorithm [25] [26] were the worst in the first 24 hours after ICU admission, and all first ICU admissions to a particular hospital for the period 1995-2009 were selected. Records were used only when all three components of the Glasgow Coma Score were provided, records for which all physiologic variables were missing were excluded, and for the remaining records, missing variables were replaced with the normal range and weighted accordingly. Ventilation status in the data base was recorded with respect to invasive mechanical ventilation on or within the first 24 hours of ICU-admission. The mortality endpoint was at hospital discharge. Exclusions: unknown hospital outcome, patients with an ICU length of stay ≤ 4 hours, and patients aged < 16 years of age.

2.2. Mortality Series

1) Exemplar univariate analysis: a running (mean) sum (window, 1 day) of daily mortality was computed over the period 1st January 1996 to 30th December 2010, with a run-in period of calendar year 1995 to establish an average baseline mortality, for ICU site 14.

2) Multivariate analysis: within a single state of the Commonwealth of Australia, for each of the hospital types (rural/regional, metropolitan, tertiary and private), as defined in the ANZICS CORE data dictionary [25], similar daily mortality series were generated, allowing a minimum 6 month run-in period.

3) The choice of exemplar and multivariate sets was made on the basis of

maximizing series length (including run-in period) with no missing values and on this basis was empirical. We have previously noted the problem of missing values in the ANZICS Adult Patient data base [27].

2.3. Statistical Analysis

Analyses were performed using Stata™ version 14 [28], the G@RCH™ 7 module [29] of OxMetrics™ 7 statistical software [30] and the "forecast" (V 6.1) package [31] of R (V 3.2.0; 2015) statistical software [32]. Continuous variables were reported as mean (SD), except where otherwise indicated, and statistical significance was ascribed at P ≤ 0.05. Summary statistics of the univariate series were characterised in terms of location (mean), scale (SD), skewness and kurtosis (tail-heaviness) by (i) classical estimators based upon (centred) moments of the distribution and (ii) recently described estimators based upon pairwise comparison of observations; in particular the user written Stata command "robjb" [33], which provides a robust Jarque-Bera normality test [34] and a robust measure of asymmetry and tail heaviness ("medcouple"; tail heaviness is compared against a value of 0.2 for the standard normal, for both observations smaller (left) and larger (right) than the median). Seasonality was explored using the "tbats" module of the "forecast" package [31]. This module implements an exponential smoothing state space model with Box-Cox transformation, ARMA errors, and trend and seasonal components [35].

Establishment of daily time-series models at the individual ICU level was based upon classic Box-Jenkins methodology (ARMA models) with investigation of (G)ARCH effects, as previously described [5] [24]. A stationary time series $\{x_t; t = 0, \pm 1, \pm 2, \ldots\}$ has an autoregressive moving average (ARMA(p,q)) structure: $x_t = \phi_1 x_{t-1} + \cdots + \phi_p x_{t-p} + \omega_t + \theta_1 \omega_{t-1} + \cdots + \theta_q \omega_{t-q}$ where $\phi_1, \phi_2, \cdots, \phi_p$ are the "autoregressive" (AR) coefficients relating the value of x at time t to its past p values, and $\theta_1, \theta_2, \cdots, \theta_q$ are the "moving average" (MA) coefficients, relating the current "white-noise", ω_t, to its past q values and $\omega_t \sim N(0, \sigma_\omega^2)$. Initial autoregressive integrated moving average model specification (ARIMA; #p, #d, #q, where "#"denotes the lags [p, q] of autocorrelations and moving averages, respectively and the degree of differencing [d]; and "1/4", say, indicates "1 through 4") was established using the "auto.arima" function of the R statistical package "forecast" [31]. Volatility of the (squared) residuals (ε) of the mean equation (conditional heteroscedasticity [36]) was checked using the PACF (partial autocorrelation function) of the squared residuals and the user-written Stata™ "armadiag" module [37], that is, (G)ARCH effects of the error variance process. The latter module, which may be implemented after the "arima", "arch" or "regress" (ordinary least squares regression, OLS) commands in Stata, plots the residual (standardized residuals with arch) autocorrelations, partial autocorrelations and P-values of the Ljung-Box Q-statistic. For an ARCH model, the variance equation is $\sigma_t^2 = \gamma_0 + \gamma_1 \varepsilon_{t-1}^2 + \gamma_2 \varepsilon_{t-2}^2 + \cdots + \gamma_q \varepsilon_{t-q}^2$, where $\varepsilon_t \sim N(0, \sigma_t^2)$, ε_t^2 are the squared residuals (innovations) and γ_i are the ARCH parameters; the conditional variance is thus modelled as an AR process. A GARCH (p, q)

model includes lagged values of the conditional variance itself $\left(\sigma_t^2 = \gamma_0 + \gamma_1\varepsilon_{t-1}^2 + \gamma_2\varepsilon_{t-2}^2 + \cdots + \gamma_q\varepsilon_{t-q}^2 + \delta_1\sigma_{t-1}^2 + \delta_2\sigma_{t-2}^2 + \cdots + \delta_p\sigma_{t-p}^2\right)$, where δ_i are the GARCH parameters (an ARMA process) [5] [24] [38].

2.4. Univariate Series

Various univariate GARCH models were considered and implemented in Stata™. As originally proposed by Engle [39], in the ARCH model, the variance of a regression model was modelled as a linear function of the lagged values of the squared regression disturbances. The conditional mean of the series (y_t) was given by $y_t = x_t\beta + \varepsilon_t$ (where $x_t\beta$ is a linear combination of lagged endogenous and exogenous variables and the (unknown) regression parameters, and ε_t are the residuals or "innovations") and the (conditional) variance $\left(\sigma_t^2\right)$ was variously specified and both normal and t (degrees of freedom (df) estimated from the data) distributions were utilised. A (1, 1) lag formulation was utilised for each variant [40]. Other than the vanilla GARCH model [41], the models assessed were those that formally deal with the stylized facts of financial data such as persistence (the conditional volatility process is not mean reverting), asymmetry (positive and negative shocks have different volatility impacts) and leverage (volatility is increased by negative shocks and decreased by positive) [42] [43]. In particular: the GARCH (p, q) model, as formulated by Bollerslev [41]; the exponential GARCH (p, q) model of Nelson (EGARCH [44]); the GJR-(Glosten, Jagaannathan and Runkle [45])-GARCH model; and the asymmetric power GARCH (APGARCH (p, q)), as described by Ding *et al* [46]. Full technical details are provided in **Appendix 1**.

2.5. Multivariate Series

Multivariate GARCH models [21] [22] [47] allow the conditional covariance matrix of the dependent variables to follow a flexible dynamic structure and the conditional mean to follow a vector-autoregressive (VAR) structure [24]. Thus, if {x_t} is a vector stochastic process of dimension N x 1, and conditioning on past information, then $x_t = \mu_t(\theta) + \epsilon_t$, where θ is a finite vector of parameters, $\mu_t(\theta)$ is the conditional mean vector and $\epsilon_t = H_t^{1/2}(\theta)z_t$; $H_t^{1/2}$ is the Cholesky factorisation of the time varying conditional covariance matrix H_t and z_t is a random innovations vector. Both H_t and μ_t depend on the unknown parameter vector θ (which can be split into two parts, one for μ_t and one for H_t [47]). MGARCH models differ in specification of H_t: direct generalisations of the univariate GARCH model of Bollerslev [41], for instance, the BEKK models [48]; linear combinations of univariate GARCH models, such as the orthogonal and G(eneralised)O-GARCH [49] [50] models; and conditional correlation models [29]. As noted by van der Weide: "The 'holy grail' in multivariate GARCH modeling is without any doubt a parameterization of the covariance matrix that is feasible in terms of estimation at a minimum loss of generality" [51]. For our purposes, the conditional mean μ_t of these models was of lesser importance and we follow Laurent *et al* [23] and Tsay [49] and impose a

constant conditional mean and consider the conditional covariance matrix $\left(H_t \right)$ as the primary objective of investigation [52]. The particular problems of forecasting squared innovations (and determining appropriate loss functions) from MGARCH models, first addressed by Andersen and Bollerslev [53], reiterated by Laurent *et al* [23], and resolved in the concept of realized variance [54], persuaded us not to undertake multivariate forecasting, which is more appropriate for construction of hedging ratios and portfolio weights [55] [56] [57] and lacks import for mortality series. We therefore considered the conditional correlations between ICUs over time and the ICU conditional variance over time, and contrast the following MGARCH models, using the G@RCH™ 7 module of Oxmetrics™ 7: GO-GARCH [50]; and the conditional correlation models: constant conditional correlation (CCC) [58], and dynamic conditional correlation (DCC) [59]. Full technical details are provided in Appendix 2. The program allows specification of different univariate GARCH models within the overall MGARCH process: GARCH, EGARCH, APGARCH and GRJ-GARCH (see above, "Univariate series") may be selected [47].

Model selection was guided by a combination of penalized information criteria, assessment of model diagnostics and comparison of model predictive performance. We utilised the Akaike (AIC) and Bayesian (BIC) information criteria (the latter for non-nested comparisons, [60]; smaller values are advantageous for AIC, with, in general, a difference of >5 indicating potential model discrimination). Model diagnostics: the use of auto- (ACF) and partial-autocorrelation (PACF) function displays, testing for time series stationarity via the KPSS (Kiawtowski-Phillips-Schmidt-Shin) test (null hypothesis of stationarity [61]) and residual white-noise (Bartlett's periodogram-based- and Portmanteau (Q)-test) were undertaken after Shumway & Stoffer [62] and as previously described [5] [24]. Lag length for various tests used Schwert's criterion (a function of sample size) where applicable [63].

Model performance (univariate series) was assessed by (i) graphical comparison from one-step ahead predictions and dynamic forecasts, the latter (from 1st July 2010 to 31st December 2010) utilising the Kalman filter (see "Dynamic forecasting" in [64]) and (ii) various loss criteria, using the "accuracy" function of the R-software "forecast" package: Mean Error (ME, $\frac{1}{T}\sum_{t=1}^{T} e_{t+h,t}$, where $e_{t+h,t} = y_{t+h} - y_{t+h,t}$); Root Mean Square Error (RMSE, $\sqrt{\frac{1}{T}\sum_{t=1}^{T} e_{t+h,t}^2}$), Mean Absolute Error (MAE, $\frac{1}{T}\sum_{t=1}^{T} |e_{t+h,t}|$); and Mean Absolute Percentage Error (MAPE, $\frac{1}{T}\sum_{t=1}^{T} |p_{t+h,t}|$, where $p_{t+h,t} = \left(\left(y_{t+h} - y_{t+h,t} \right) \big/ y_{t+h} \right)$ [65]. Forecast comparison between competing models was assessed by the Stata™ user-written module "dmariano" which computes the Diebold-Mariano comparison of predictive accuracy (for loss criteria MSE, MAE and MAPE [66]), albeit we note the caution that "… different [GARCH] models can lead to almost equivalent pre-

dictive formulas" [7].

3. Results

The initial data set, 1995-2010, contained 674,193 patient records from 157 ICUs. For the exemplar univariate analysis (ICU site 14), there were 5479 observations over the calendar years 1996-2010, with no missing values. The mean series mortality was 0.17 (0.01) and summary statistics for the raw and first differenced series are seen in Table 1, where tail-heaviness for the first differenced series is noted. Not surprisingly the raw series demonstrated a high degree of autocorrelation to the 100th lag (and beyond, data not shown). The raw series (Figure 1, top panel) displayed a downward mortality trend and rejected the null of stationarity (KPSS test) at all lags (n = 10, p < 0.01). The first differenced series (Figure 1, bottom panel) displayed stationarity at all lag lengths (n = 10, p > 0.1) and *the latter series was used for model development*, the marked kurtosis being a feature (Table 1). Residuals from OLS regression of the raw series against time displayed autocorrelation and ARCH effects with p-values of the Ljung-Box Q-statistic approximating zero. We were unable to establish season-
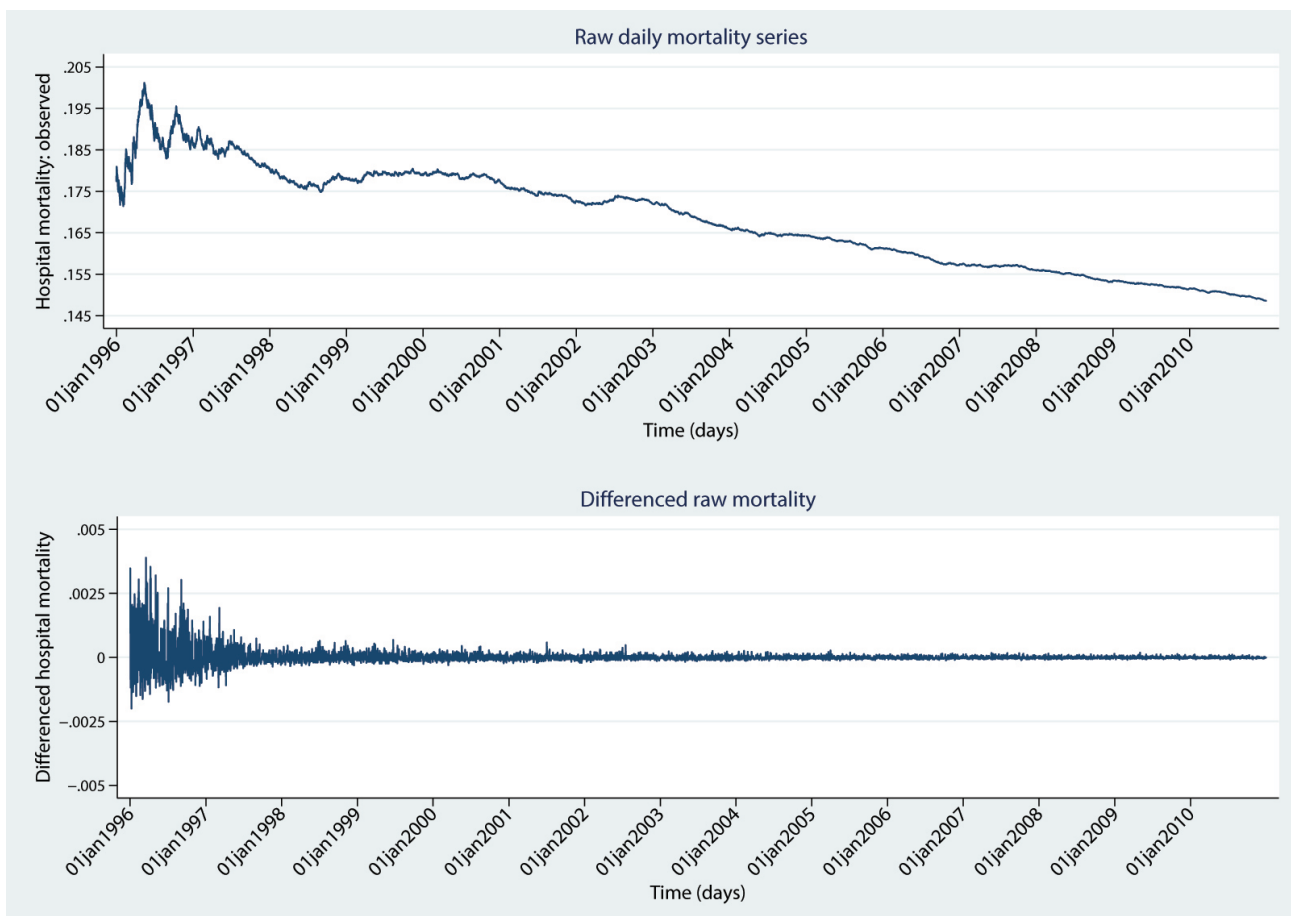
Upper panel: Y-axis, observed daily hospital mortality for ICU site 14. X-axis, time in days. Lower panel: Y-axis, first differenced daily hospital mortality for ICU site 14. X-axis, time in days.

Figure 1. Raw and differenced daily mortality series.

Table 1. Summary statistics and autocorrelations for raw and first differenced mortality series for ICU site 14.

Series	N	Mean	SD	Skewness	Kurtosis	Minimum	Maximum	S-Wilk z	S-Wilk p	robjb-s	robjb-k	medcouple-L	medcouple-R
Raw	5479	0.16835	0.01178	0.104	2.000	0.149	0.201	12.718	0.000	0.000	0.000	−0.190	0.006
First differenced	5478	−0.00001	0.00029	4.215	48.663	−0.002	0.004	18.954	0.000	0.000	0.000	0.513	0.282

Autocorrelations								
Lag	1	2	3	4	5	10	20	40	100
Raw	0.9994	0.9987	0.9981	0.9974	0.9968	0.9938	0.9881	0.9759	0.9358
First differenced	0.0053	0.0597	−0.0103	−0.0479	0.0477	0.0150	0.0180	0.0395	−0.0003

SD, standard deviation. S-Wilk, Shapiro-Wilk normality test. z, z-statistic. p, p-value. robjb-s, robust Jarque-Bera normality test (skewness). robjb-k, robust Jarque-Bera normality test (kurtosis). medcouple-L, left medcouple (observations less than median). medcouple-R, right medcouple (observations greater than median). Tail-heaviness (medcouple) is compared with a value of 0.2 for the standard normal.

ality of the (raw) series at the monthly or yearly level using the "tbats" module of the R-software "forecast" package.

The model formulated by the "auto-arima" module of the R-software "forecast" package [31] was ARIMA (1/4, 1, 1/3). Alternate specifications up to ARIMA (1/9, 1, 1/3) were considered, but such extensive parameterisation of the mean dynamic was considered to lack interpretation and a simpler mean model, ARIMA (7, 1, 0) was chosen to reflect the daily series (additive seasonality), albeit the latter model and all other ARIMA variants demonstrated substantial ARCH effects. Of the 8 GARCH models initially considered, an asymmetric power (G)ARCH model (APGARCH, [46] [67]) with t-distribution (df, 11.63) and ARMA (7, 0) for the mean-model, was the most parsimonious and, not surprisingly, had substantial information criterion advantage over the ARIMA mean model (ARIMA (7, 1, 0)); BIC −86,324 versus −73,873. Information criteria (AIC and BIC) with model and estimated t df for all univariate GARCH models are detailed in Table 2. For each of the GARCH variants, a t-distribution had information criterion advantage, but between-model differences based upon BIC were rather modest. Graphical display of ARCH specification tests for each of the t-distribution variants is seen in Figure 2. The APGARCH-t model (lower right panel) appeared the most parsimonious, especially with regard to the lack of residual (ε_t^2) serial correlation, as indicated by the lag p-values ($\gg$0.05) of the (Ljung-Box) Q-statistic [68]. APGARCH parameter estimates are shown in Table 3 (cross-referenced to model formula in Statistical analysis, Univariate series (iii), above), the scalar sum of α and β (=1.002) suggesting persistence of conditional volatility. Conditional variance plots (Figure 3) from the APGARCH-t model exhibited extreme volatility during the period 1996-1998 (upper panel) and substantial, but declining volatility, from 1998-2010 (lower panel). This being said, a "level" shift (coded 1/0) at 1st January 1998 lacked significance (p = 0.23). The predictive performance of the APGARCH-t model was also considered. One-step-ahead predictions are seen in Figure 4 (upper panel), demonstrating, not surprisingly, virtual identity of the raw series signal and the one-step predictions (Table 4). The dynamic 6-month forecast is shown in Figure 4

(lower panel), with some divergence between the raw series and forecast, reflected in the increment of the MAPE loss function for the out-of sample forecast, 2nd June 2010 to 31st December 2010 (Table 4(a), third column). With respect to comparative predictive accuracy between GARCH variants (in particu-

Table 2. GARCH model comparison.

Model	Model-df	AIC	BIC	t-df
GARCH	3	−86133.62	−86113.80	
GARCH-t	4	−86332.11	−86305.68	11.35
EGARCH-t	6	−86359.50	−86432.02	11.61
APGARCH	5	−86169.02	−86319.85	
APGARCH-t	6	−86363.79	−86324.14	11.63
GJR-GARCH	4	−86171.55	−86145.12	
GJR-GARCH-t	5	−86351.47	−86332.43	11.63

df, degrees of freedom. −t, t distribution. t-df, estimated t degrees of freedom, EGARCH model, no convergence.
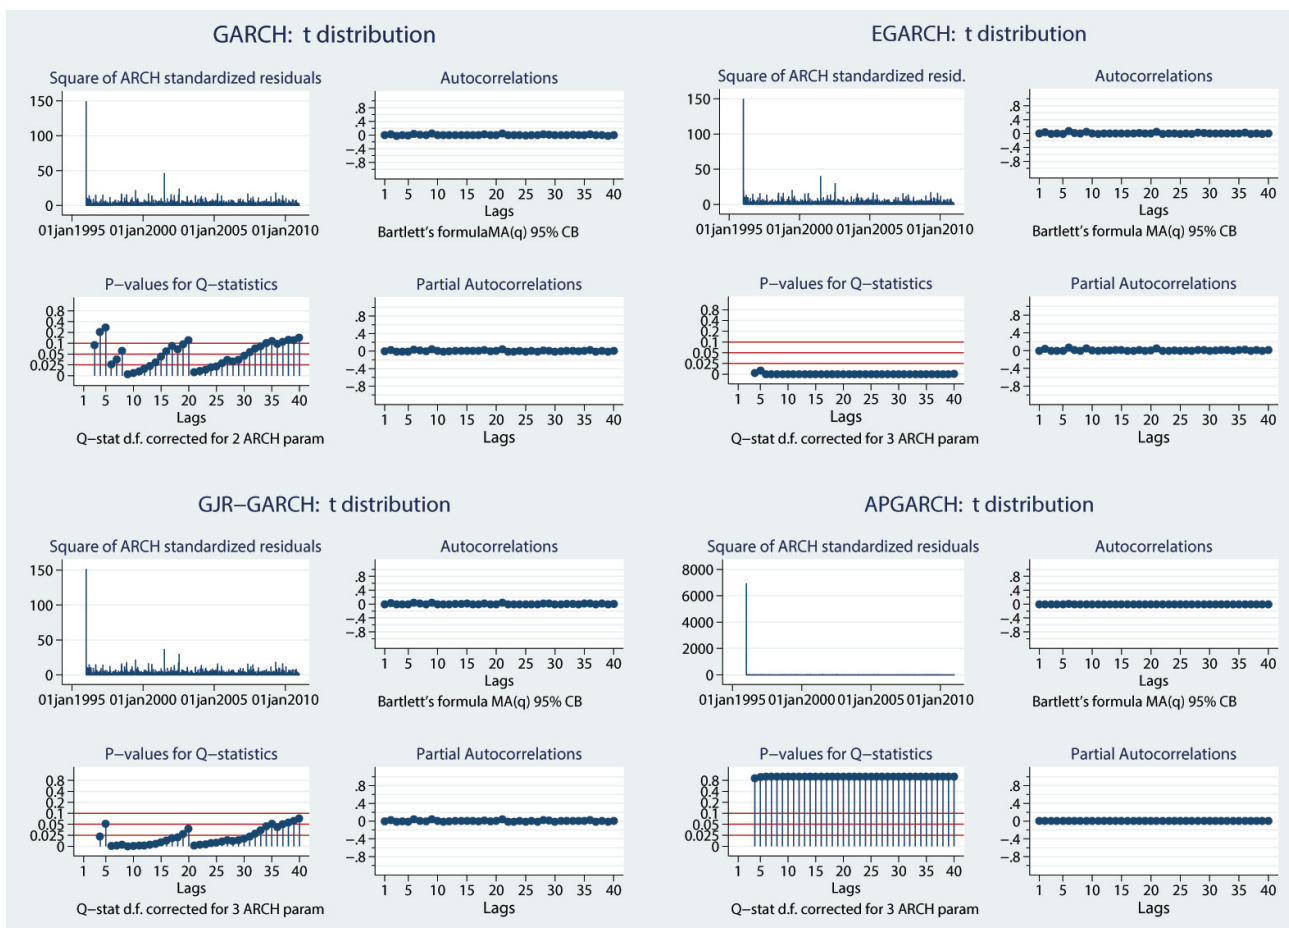
Upper left panel: GARCH (1, 1) model, t-distribution. Upper right panel: EGARCH (1, 1) model, t-distribution. Lower left panel: GJR-GARCH (1, 1) model t-distribution. Lower right panel: APGARCH (1, 1), t-distribution. 95% CB, 95% confidence bands. Q-stat, Ljung-Box Q-statistic [80]. ARCH param, ARCH parameters.

Figure 2. ARCH specification tests for GARCH models.

Table 3. Parameter estimates of the APGARCH-t model.

Equation	Parameter	Estimate	p-value	Lower 95% CI	Upper 95% CI
ARMA	L7.AR	0.054	0.000	0.028	0.079
ARCH	L1.aparch (α)	0.032	0.000	0.020	0.044
ARCH	L1.aparch_e (γ)	−0.330	0.000	−0.450	−0.210
ARCH	L1.pgarch (β)	0.969	0.000	0.962	0.977
ARCH	Constant	0.000	0.741	0.000	0.000
POWER	power (δ)	1.791	0.000	1.210	2.372
t-df		11.626		9.944	13.665

L, lag. t-df, estimated t degrees of freedom.
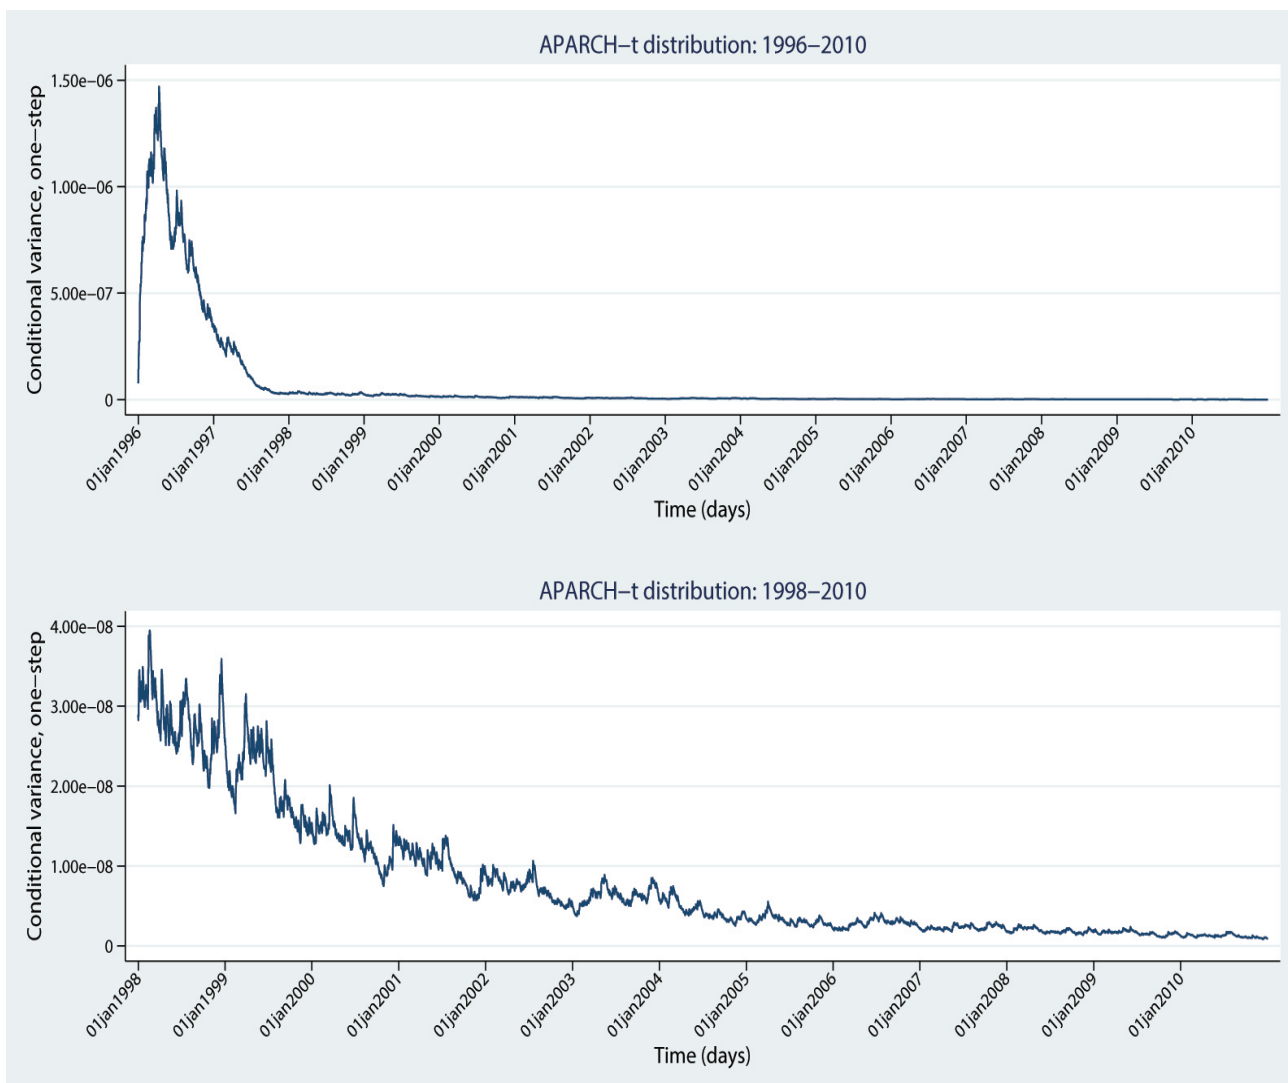
Upper panel: 1996-2010, Y-axis, one-step conditional variance. X-axis, time (days); Lower panel: 1998-2010, Y-axis, one-step conditional variance. X-axis, time (days).

Figure 3. APGARCH-t conditional variance plots.

Table 4. (a) Forecast evaluation of APGARCH-t model. (b) Comparative forecasts.

(a)

	01 Jan. 1966-31 Dec. 2010	01 Jan. 1966-01 June 2010	02 June 2010-31 Dec. 2010
Estimation sample N	5478	5265	213
Mean Error (ME)	−0.00004	−0.00005	−0.0373
Root Mean Square Error (RMSE)	0.0004	0.0003	0.0382
Mean Absolute Error (MAE)	0.0002	0.0001	0.0373
Mean Absolute Percentage Error (MAPE)	0.0707	0.0707	24.906

Dec., December. Jan., January.

(b)

Model: compared with APARCH-t	MSE	MAE	MAPE
	p	p	p
GARCH-t	0.817	0.052	0.032
EGARCH-t	0.852	0.040	0.024
GJR-GARCH-t	0.319	0.677	0.687

MSE, means squared error. MAE, mean absolute error. MAPE, mean absolute percentage error; Significant p-values indicate superior APARCH-t forecast performance.
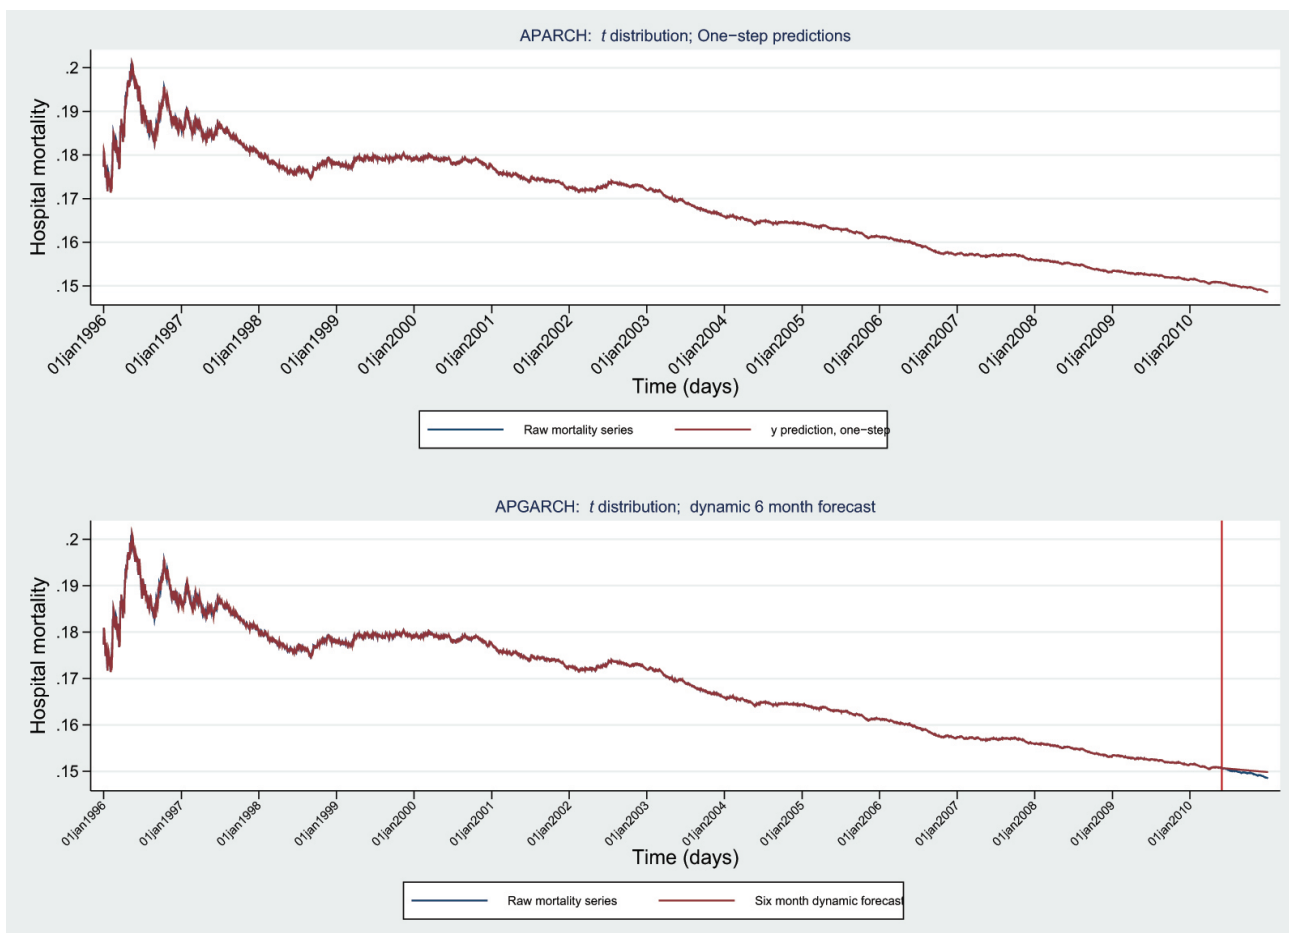
Upper panel: One-step ahead predictions for the APGARCH-t (mean) model; Y-axis, mortality; X-axis, time (days); Lower panel: Six month dynamic forecast for APGARCH-t (mean) model; Y-axis, mortality; X-axis, time (days).

Figure 4. APGARCH-t model: One step prediction and dynamic forecast.
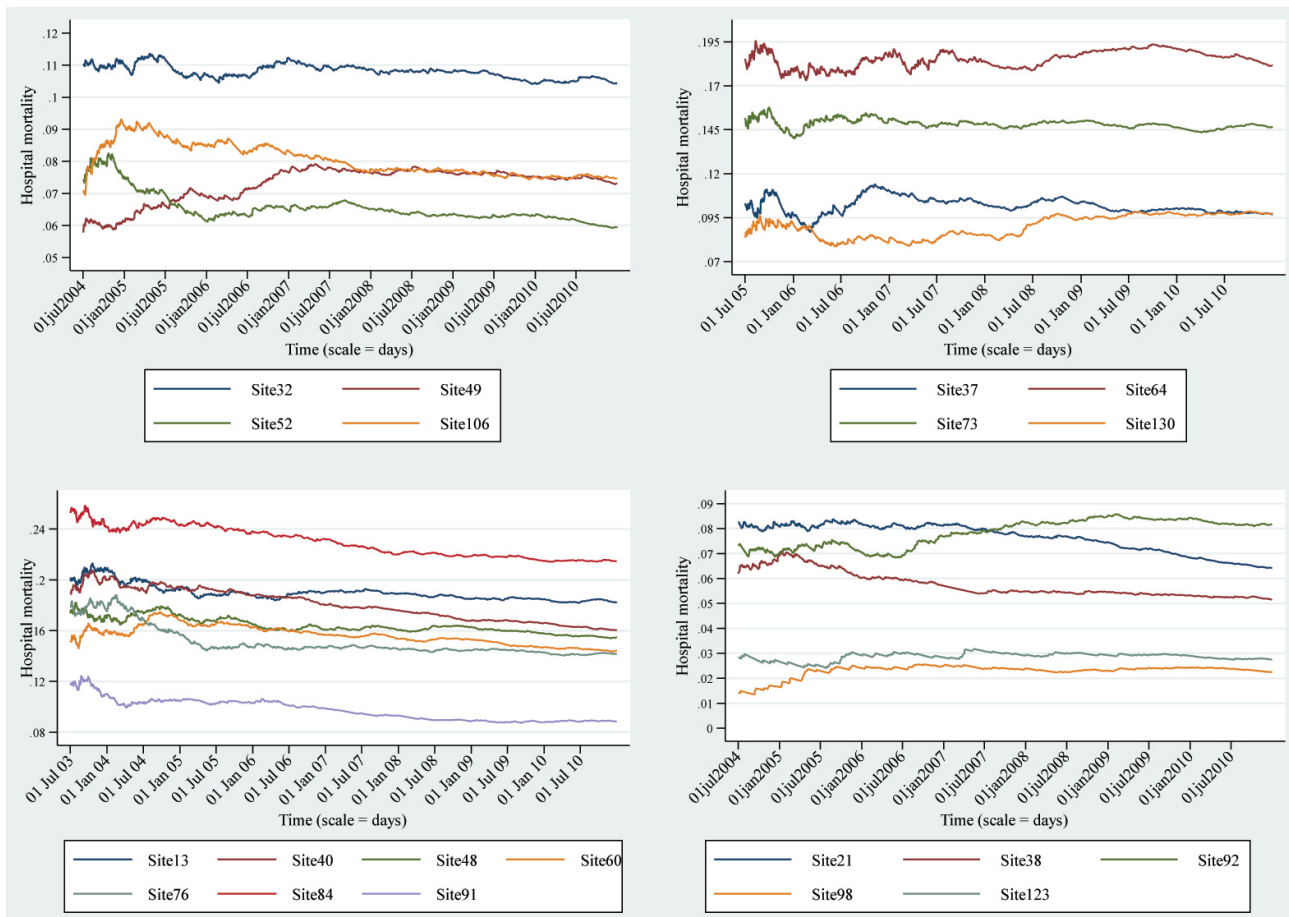
Hospital raw mortality (Y-axis) plotted against time (days) for various sites in the multivariate series: upper left, rural/regional; upper right, metropolitan; lower left, tertiary; lower right, private.

Figure 5. Multivariate rural/regional, metropolitan, tertiary and private series.

lar, using the t distribution): for a maximum lag of 32 (chosen by Schwert criterion [63], with a uniform kernel to calculate long-run variance) a variable superiority of the APGARCH-t forecasts (compared with GARCH-t, EGARCH-t and GJR-GARCH-t) was demonstrated, as indicated in Table 4(b), test significance being dependent upon the particular loss criterion [69]. Compared with the conventional ARIMA (7, 1, 0) model, the APGARCH-t demonstrated a superior forecast (MAPE, $p = 0.015$; MAE, $p = 0.026$).

The four multivariate component raw series (rural/regional, metropolitan, tertiary and private) are seen in **Figure 5**, demonstrating varied levels of, and trend decline in, mortality. The series are further characterised in terms of summary statistics and autocorrelations for raw and differenced series in **Table 5** and **Table 6** respectively. The most notable findings were (i) marked kurtosis and rejection of normality for both the raw and differenced series, and (ii) for each of the multivariate series, the raw series rejected the null of stationarity (KPSS test) at all lags ($n = 10$; $p < 0.01$) and the first differenced series displayed stationarity at all lag lengths ($n = 10$; $p > 0.1$). The 3 MGARCH model variants (GO-GARCH (normal distribution only), CCC (normal and t-distribution) and

Table 5. Summary statistics for rural/regional, metropolitan, tertiary and private multivariate series.

Locationz	ICU site	Series	N	Mean	SD	Skewness	Kurtosis	Minimum	Maximum	S-Wilk	S-Wilk p
Rural/Regional 1st-July-2004 to 31st-December-2010	32	Raw	2374	0.10800	0.0021	0.2275	2.4469	0.1040	0.1137	8.105	0.0000
		Differenced	2373	0.00000	0.0002	2.7344	12.4898	−0.0005	0.0015	15.767	0.0000
	49	Raw	2374	0.07285	0.0053	−1.1625	3.1737	0.0578	0.0792	14.041	0.0000
		Differenced	2373	0.00001	0.0002	3.4719	16.9665	−0.0003	0.0013	16.688	0.0000
	52	Raw	2374	0.06555	0.0047	1.8186	5.9761	0.0592	0.0826	14.561	0.0000
		Differenced	2373	−0.00001	0.0002	4.8580	43.4592	−0.0006	0.0029	16.616	0.0000
	106	Raw	2374	0.08039	0.0050	0.5939	2.2198	0.0694	0.0932	12.393	0.0000
		Differenced	2373	0.00000	0.0002	4.5056	27.158	−0.0004	0.0025	17.030	0.0000
Metropolitan 1st-July-2005 to 31st-December-2010	37	Raw	2010	0.10175	0.0046	0.1195	3.2859	0.0867	0.1141	8.471	0.0000
		Differenced	2009	0.00000	0.0004	5.2374	47.9392	−0.0018	0.0061	16.135	0.0000
	64	Raw	2010	0.18514	0.0048	−0.2384	2.0787	0.1728	0.1957	9.002	0.0000
		Differenced	2009	0.00000	0.0006	2.3950	12.9483	−0.0037	0.0035	14.721	0.0000
	73	Raw	2010	0.14827	0.0026	0.1365	3.9181	0.1398	0.1580	7.408	0.0000
		Differenced	2009	0.00000	0.0004	2.6539	24.7187	−0.0024	0.0039	14.634	0.0000
	130	Raw	2010	0.09023	0.0064	−0.2206	1.4799	0.0784	0.0988	12.531	0.0000
		Differenced	2009	0.00001	0.0004	4.7265	37.2795	−0.0013	0.0048	15.982	0.0000
Tertiary 1st-July-2003 to 31st-December-2010	13	Raw	2741	0.18985	0.0062	1.4893	4.9537	0.1818	0.2133	14.030	0.0000
		Differenced	2740	−0.00001	0.0004	2.2167	17.7470	−0.0023	0.0037	15.237	0.0000
	40	Raw	2741	0.18045	0.0124	0.0787	1.7626	0.1604	0.2077	11.391	0.0000
		Differenced	2740	−0.00001	0.0003	2.7742	32.2987	−0.0019	0.0038	15.470	0.0000
	48	Raw	2741	0.16403	0.0060	0.6958	2.7709	0.1540	0.1825	11.664	0.0000
		Differenced	2740	−0.00001	0.0004	2.8927	20.2068	−0.0015	0.0041	15.605	0.0000
	60	Raw	2741	0.15641	0.0073	0.2747	2.4559	0.1435	0.1749	9.506	0.0000
		Differenced	2740	0.00000	0.0003	3.0224	27.7510	−0.0017	0.0039	15.460	0.0000
	76	Raw	2741	0.15100	0.0118	1.7084	4.5631	0.1403	0.1884	15.883	0.0000
		Differenced	2740	−0.00001	0.0004	2.3249	18.1360	−0.0024	0.0033	15.334	0.0000
	84	Raw	2741	0.22975	0.0119	0.2812	1.7031	0.2140	0.2589	12.688	0.0000
		Differenced	2740	−0.00001	0.0004	2.4361	22.5103	−0.0024	0.0045	15.313	0.0000
	91	Raw	2741	0.09748	0.0084	0.6651	2.8833	0.0873	0.1247	13.223	0.0000
		Differenced	2740	−0.00001	0.0003	4.5859	56.5747	−0.0015	0.0052	16.536	0.0000
Private 1st-July-2004 to 31st-December-2010	21	Raw	2374	0.07638	0.0057	−0.7457	2.2253	0.0641	0.0839	13.098	0.0000
		Differenced	2373	−0.00001	0.0002	2.9992	17.5986	−0.0005	0.0015	15.581	0.0000
	38	Raw	2374	0.05749	0.0051	0.9777	2.6550	0.0516	0.0706	13.807	0.0000
		Differenced	2373	0.00000	0.0001	5.9317	61.3591	−0.0004	0.0022	16.887	0.0000
	92	Raw	2374	0.07823	0.0054	−0.4022	1.6272	0.0683	0.0858	12.957	0.0000
		Differenced	2373	0.00000	0.0002	3.8774	24.9860	−0.0005	0.0021	16.503	0.0000
	98	Raw	2374	0.02300	0.0025	−2.2758	7.3654	0.0135	0.0257	15.725	0.0000
		Differenced	2373	0.00000	0.0001	6.8764	65.5402	−0.0001	0.0017	17.678	0.0000
	123	Raw	2374	0.02871	0.0015	−0.9272	3.4962	0.0241	0.0318	11.843	0.0000
		Differenced	2373	0.00000	0.0001	7.5953	75.4797	−0.0002	0.0019	17.805	0.0000

SD, standard deviation. S-Wilk, Shapiro-Wilk normality test. z, z-statistic. p, p-value.

Table 6. Autocorrelations of raw and first differenced mortality series for rural/regional, metropolitan, tertiary and private multivariate series.

Location	ICU number	N	Lag	1	2	3	4	5	10	20	40	100
Rural/Regional 1st-July-2004 to 31st-December-2010	32	2374	Raw	0.9946	0.9895	0.9847	0.9799	0.9751	0.9518	0.9014	0.8154	0.6208
		2373	Differenced	−0.0274	−0.0299	0.0061	0.0089	−0.0275	0.0126	−0.0282	−0.0031	0.0330
	49	2374	Raw	0.9979	0.9959	0.0042	0.9926	0.9910	0.9826	0.9690	0.9403	0.8263
		2373	Differenced	0.0483	0.0070	−0.0198	0.0406	−0.0509	0.0002	0.0232	0.0449	−0.0391
	52	2374	Raw	0.9981	0.9959	0.9939	0.9918	0.9897	0.9796	0.9555	0.8957	0.6809
		2373	Differenced	0.1227	−0.0310	0.0147	−0.0085	−0.0683	−0.0223	−0.0421	0.0203	0.0405
	106	2374	Raw	0.9979	0.9957	0.9935	0.9911	0.9886	0.9748	0.9573	0.9228	0.8240
		2373	Differenced	−0.0087	0.0426	0.0297	0.0599	0.0184	−0.0355	0.0295	0.0246	−0.0160
Metropolitan 1st-July-2005 to 31st-December-2010	37	2010	Raw	0.9952	0.9902	0.9855	0.9810	0.9769	0.9538	0.8970	0.7653	0.4058
		2009	First differenced	0.0236	−0.0315	−0.0125	−0.0493	0.0210	−0.0138	0.0259	−0.0336	0.0313
	64	2010	Raw	0.9925	0.9847	0.9778	0.9711	0.9637	0.9336	0.8807	0.7573	0.5021
		2009	First differenced	0.0231	−0.0596	−0.0128	0.0459	−0.0535	−0.0247	−0.0042	−0.0026	−0.0059
	73	2010	Raw	0.9864	0.9735	0.9605	0.9495	0.9418	0.8972	0.8215	0.6175	0.0536
		2009	First differenced	−0.0184	0.0070	−0.0537	−0.1208	−0.0808	−0.1250	0.1100	0.0428	−0.0090
	130	2010	Raw	0.9974	0.9948	0.9925	0.9904	0.9881	0.9781	0.9624	0.9315	0.8112
		2009	First differenced	0.0060	−0.0752	0.0037	0.0354	−0.0651	−0.0541	0.0412	0.0486	0.0535
Tertiary 1st-July-2003 to 31st-December-2010	13	2741	Raw	0.9971	0.9944	0.9919	0.9894	0.9871	0.9746	0.9520	0.9111	0.7533
		2740	Differenced	−0.0777	−0.0374	0.0142	0.0023	−0.0006	−0.0106	−0.0894	0.0445	−0.0232
	40	2741	Raw	0.9991	0.9982	0.9973	0.9964	0.9954	0.9901	0.9779	0.9577	0.8925
		2740	Differenced	−0.0760	0.0779	0.0403	0.0342	0.0778	0.0276	−0.0590	0.0317	0.0364
	48	2741	Raw	0.9971	0.9943	0.9916	0.9890	0.9862	0.9727	0.9501	0.8988	0.7471
		2740	Differenced	−0.0601	−0.0285	−0.0169	−0.0042	0.0426	−0.0533	0.0634	0.0438	−0.0078
	60	2741	Raw	0.9985	0.9969	0.9953	0.9938	0.9923	0.9852	0.9686	0.9352	0.8285
		2740	Differenced	0.0553	0.0121	−0.0501	0.0311	−0.0066	0.0438	−0.0293	0.0770	0.0005
	76	2741	Raw	0.9984	0.9969	0.9953	0.9935	0.9915	0.9819	0.9662	0.9421	0.8495
		2740	Differenced	−0.0092	0.0685	0.0262	0.9190	−0.0339	0.0037	−0.0039	0.0513	−0.0237
	84	2741	Raw	0.9984	0.9968	0.9952	0.9936	0.9920	0.9837	0.9680	0.9441	0.8678
		2740	Differenced	−0.0662	−0.0074	−0.0168	−0.0019	−0.0177	−0.0360	0.0512	−0.0323	−0.0077
	91	2741	Raw	0.9980	0.9960	0.9940	0.9922	0.9905	0.9806	0.9623	0.9308	0.7933
		2740	Differenced	0.0348	−0.0244	0.0759	−0.0313	0.0088	−0.1068	−0.0329	−0.0409	−0.0347
Private 2nd-July-2004 to 31st-December-2010	21	2374	Raw	0.9984	0.9967	0.9951	0.9935	0.9919	0.9843	0.9709	0.9422	0.8631
		2373	Differenced	0.0336	−0.0198	−0.0094	0.0087	−0.0007	−0.0169	−0.0424	0.0026	0.0056
	38	2374	Raw	0.9991	0.9982	0.9973	0.9964	0.9955	0.9900	0.9795	0.9590	0.8877
		2373	Differenced	−0.0077	−0.0204	−0.0176	0.0727	−0.0273	−0.0134	−0.0441	0.0162	−0.0050
	92	2374	Raw	0.9992	0.9984	0.9977	0.9969	0.9961	0.9922	0.9832	0.0962	0.8991
		2373	Differenced	−0.0409	−0.0085	−0.0089	0.0422	0.0229	0.0260	0.0140	−0.0051	0.0107
	98	2374	Raw	0.9966	0.9931	0.9895	0.9859	0.9822	0.9632	0.9283	0.8571	0.6486
		2373	Differenced	0.0788	0.0396	−0.0115	0.0572	0.0883	0.0083	−0.0246	0.0299	−0.0222
	123	2374	Raw	0.9969	0.9937	0.9907	0.9878	0.9849	0.9684	0.9317	0.8574	0.6119
		2373	Differenced	0.0094	−0.0259	−0.0182	−0.0062	0.0549	0.0032	0.0082	0.0167	−0.0002

Table 7. Model specifications for the MGARCH series.

Model	Rural/Regional Series = 4		Metropolitan Series = 4		Tertiary Series = 7		Private Series = 5	
	AIC	BIC	AIC	BIC	AIC	BIC	AIC	BIC
GO-GARCH-normal								
GARCH univariate	−59.18	−59.14	−54.16	−54.11	−99.65	−99.56	−77.22	−77.16
APARCH univariate	No−convergence		No−convergence		No−convergence		No-convergence	
CCC-normal								
GARCH univariate	−59.14	−59.09	−54.11	−54.06	−99.56	−99.47	−77.08	−77.02
APARCH univariate	−59.38	−59.32	−54.25	−54.18	No−convergence		No-convergence	
CCT-t								
GARCH univariate	−61.85	−61.81	−54.81	−54.76	−99.97	−99.88	−81.95	−81.89
APARCH univariate	−61.85	−61.78	−54.86	−54.79	No−convergence		−82.11	−82.02
DCC-normal								
GARCH univariate	−59.14	−59.09	−54.11	−54.05	−99.56	−99.87	−77.08	−77.01
APARCH univariate	−59.38	−59.31	−54.25	−54.17	No−convergence		−77.42	−77.43
DCC-t								
GARCH univariate	−61.86	−61.81	−54.81	−54.75	−99.97	−99.88	−81.95	−81.89
APARCH univariate	−61.85	−61.78	−54.86	−54.78	No−convergence		−82.11	−82.02

GO-GARCH: generalised orthogonal GARCH. CCC: constant conditional correlation. DCC: dynamic conditional correlation. -normal: normal distribution. −*t*: *t* distribution (estimated from the data)

DCC (normal and *t*-distribution)), displayed varying degrees of convergence difficulties, especially with the 7 component univariate series of the tertiary multivariate set. Table 7 shows model information criteria (AIC and BIC) for (i) two univariate specifications, GARCH (1, 1) and APGARCH (1, 1) and (ii) the 4 MGARCH series (rural/regional, metropolitan, tertiary and private). Within each MGARCH series no model specification dominated, although there was some advantage for the *t*-distribution in the private series. The tertiary series allowed only a univariate GARCH specification and the GO-GARCH model was unable to converge with the univariate APGARCH. Graphical analysis is presented of the conditional correlations and variances from the DCC-*t* model (univariate APGARCH (1, 1)), except for the tertiary series (univariate GARCH (1, 1)). Figure 6 & Figure 7 show the conditional correlations between the univariate series of the rural/regional and metropolitan series. The correlations between the component univariate series were quite variable, and demonstrated reversion to a constant level (see Statistical analysis, Multivariate series, ii(b), above), the private series being noted for relatively high positive correlation (0.27 - 0.85). The variances of the component series of each of the multivariate series demonstrated a variable rate of time decline from periods of early volatility and volatility spikes in the rural/regional and metropolitan and private series during the calendar years mid 2006-2007 to mid 2007-2008.
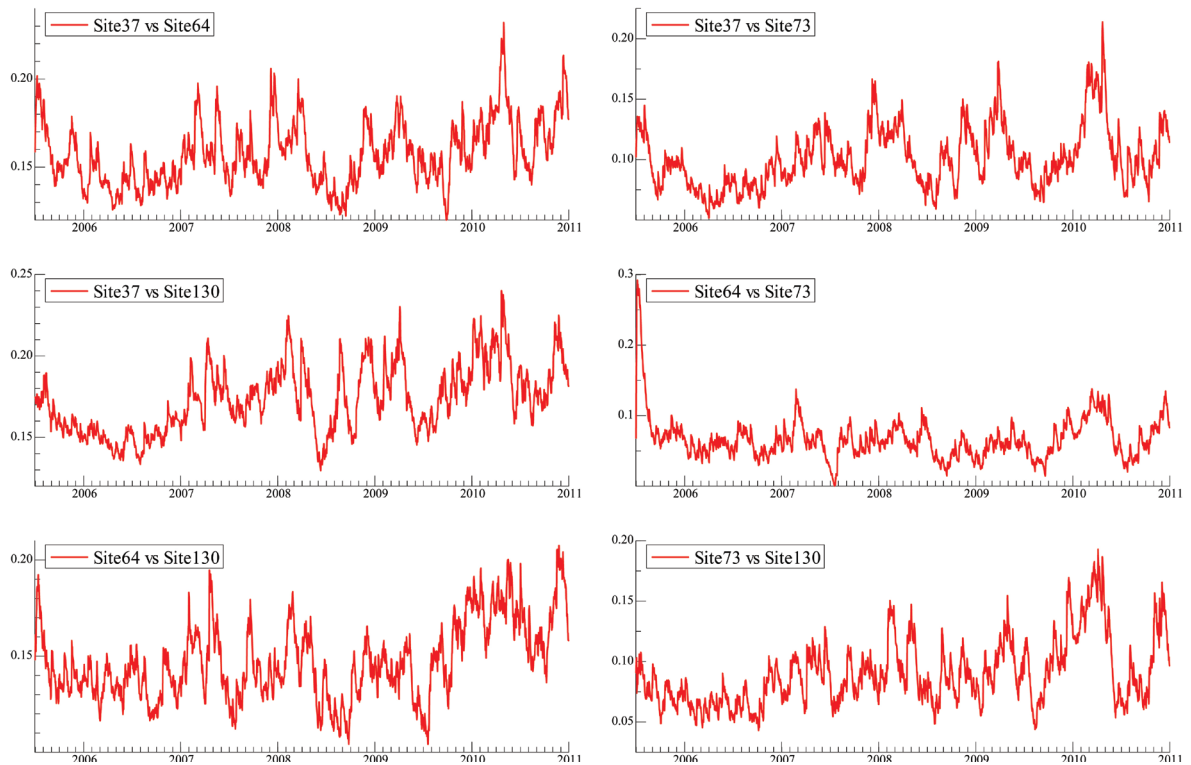
4. Discussion

The current study has demonstrated that high frequency (daily) mortality series

Conditional correlations (Y-axis) over time (X-axis) for various combinations of ICU sites.

Figure 6. Conditional correlations between the ICU sites (32, 49, 52 and 106) of the rural/regional multivariate series.

Conditional correlations (Y-axis) over time (X-axis) for various combinations of ICU sites.

Figure 7. Conditional correlations between the ICU sites (37, 64, 73 and 130) of the metropolitan multivariate series (obtained from a variant of the DCC model: Asymmetric Corrected Dynamic Correlation Model (Aielli) [59], with univariate APGARCH (1, 1).

exhibit (G)ARCH effects, consistent with our two previous studies of the same data-base [5] [24], albeit the specific models differ, most likely reflecting different temporal data aggregation, daily versus monthly [70]. Thus, unlike some financial data, for example exchange rates, the ability to discern ARCH effects did not decrease with increasing sampling interval [71].

In financial time series, conditional asymmetry is a stylized fact [67]. That is, there is a negative correlation between the squared current innovations $\left(\varepsilon_t^2\right)$ and the past innovations, or empirically, the volatility due to, say, a price decrease is greater than that of a comparable price increase. For classical GARCH, the conditional variance is a function of the modulus of the past ε_t, positive $\left(\varepsilon_t^+\right)$ and negative $\left(\varepsilon_t^-\right)$ innovations having the same effect on current volatility: $\mathrm{Cov}\left(\varepsilon_t^+, \varepsilon_{t-h}\right) = \mathrm{Cov}\left(\varepsilon_t^-, \varepsilon_{t-h}\right) = 0, h > 0,$ where Cov = covariance. Under conditions of second-order stationarity, ε_t may be decomposed as $\varepsilon_t = \sigma_t n_t$ where n_t is an iid sequence and σ_t is a measurable positive function of the past of ε_t. The circumstance $\mathrm{Cov}\left(\sigma_t, \varepsilon_{t-h}\right) < 0$ is that of a *leverage* effect. For the APGARCH model (see Statistical analysis, Univariate series, (iii)), if $|\gamma_i| > 0$ then "… negative innovations have more impact on current volatility than positive ones of the same modulus" [67]; that is, there exists a leverage effect. This condition was satisfied in the current APGARCH model ($|\gamma|$ = 0.330, γ was significantly different from 0, P = 0.000, Table 3).

However, a degree of caution is required in considering the application of asymmetric volatility models [43] to non-economic/financial data. Such models *require* "…a specification that can accommodate a leverage effect" [40], such specification being described as "crucial" by the authors of the APGARCH model [46]. We make two points with regard to leverage in the current series: (i) a snapshot of the series, 1996-1998, raw versus the differenced mortality (Figure 8), is suggestive of volatility clustering during sharp falls in mortality, similar to that described by Engle for financial data [72] and (ii) the early volatility maybe at least in part due to reporting artefacts or processes at that time, in addition to the increased variability associated with smaller numbers. The overall trend then is for declining mortality, which would be confounded with improvements in reporting processes, data completeness and increasing numbers. This would suggest that we are not seeing a leverage effect as such, but rather a confluence of trends.

The implications of a volatility model perspective [43] in the context of mortality rates are best considered against the background of (i) the above perspective of the financial paradigm where the trade-off between risk and expected return is a fundamental concern and the measurement and forecasting of volatility a core pursuit [73] and (ii) recent actuarial and demographic literature, where detailed comparisons between (G)ARCH–based stochastic mortality models and the orthodox Lee-Carter model [74], have favoured the former. Of interest, the literature review (section 2.3) of Andersen *et al.*, "Volatility forecasting in fields outside finance" in 2006 [73] made no specific mention of mortality series. The

Upper panel: raw series 1996-1998; Y-axis, raw mortality; X-axis, time (days); Lower panel: differenced series 1996-1998; Y-axis, differenced mortality; X-axis, time (days).

Figure 8. Raw mortality and differenced mortality series: 1996 to 1998.

original Lee-Carter model (modelling the logarithm of the central death rate $m_{x,t}$ for age x at time t) used ARIMA functions to undertake mortality forecasts, but assumed homoscedasticity and constant volatility, which assumptions are belied by the structure of long-term (yearly) mortality series which demonstrate non-stationarity, conditional heteroscedasticity and non-normality [75] [76]. Thus, apposite analysis of mortality time-series mandates the demonstration and appropriately modelling of volatility. As opposed to the population perspective of demography, we model the mortality of the critically-ill, where the interplay of an ensemble of patient factors (severity of illness, patient type) and provider characteristics (ICU occupancy, structure and staffing), not all of which are in principle identifiable, are determinate in the conditional heteroscedasticity of the mortality series.

The parsimony of univariate GARCH models has been shown in actuarial and demographic studies, but the analysis of, say, cross-(nation)-state mortality correlations [75] within the same framework has been "…a largely unchartered territory" [77]. To this end, the recent study of Gao and Hu [78] reports 8 separate GARCH (1, 1) models in a sub-section "An application to multi-country study",

rather than a multivariate approach. Although not undertaking MGARCH forecast assessment in the current series (see Statistical analysis, Multivariate series; above), recent evidence has suggested primacy of the DCC model, at least in financial series [23] [56] and further sophisticated variants of the DCC model have been presented [79], although cautions about the DCC representation have been expressed [80]. In a wide ranging study of time series from finance, physiology and genomics, Podobnik *et al.*, using time-lag random matrix theory, demonstrated that "… cross-correlations are ubiquitously present in many systems … [and] … studying these cross-correlations is a necessary prerequisite for understanding them …" [81]. Similar studies have been presented from the social sciences [82] [83]. As the selection of univariate series was based upon defined provider categories of hospital type and locality, there was an expectation of substantial but variable levels of correlation between these (multivariate) series, more so given the particular structure of critical care practice in Australia and New Zealand (uniform training scheme and closed ICUs [84]). The conditional correlations were surprisingly low (<0.1) between tertiary series and substantial (0.4 - 0.6) between rural-regional and private series. An explanation for this finding could be the similarity / uniformity of patient-mix and treatments in tertiary ICUs. Thus, condition correlations would look relatively independent, as opposed to the cross-correlations, which would be expected to be related. Such explanation would also suggest that other sets of hospitals were more heterogeneous, which seems plausible, although less so for the metropolitan centres. Conditional variance volatility demonstrated, not surprisingly, persistence and the degree of volatility was most marked in non-tertiary series where annual patient admission numbers were lower [85]. Within the statistical process control (SPC) paradigm, where we have demonstrated the facility of univariate GARCH modelling [5], the extension to "… monitor [ing] outcomes at more than one unit simultaneously" has been advocated [86] and, on the basis of the cross-correlations revealed in the current analysis, would also appear to have a plausible empirical basis.

5. Conclusion

High frequency ICU mortality time series display autocorrelation, persistence of conditional variance and volatility which are appropriately modelled using estimators which explicitly account for these attributes. Similarly, multivariate mortality series exhibit these stylised facts and temporal dependencies, reflected in varying degrees of conditional correlations which belie the use of (repeated) univariate approaches to the understanding of the performance of sets of ICUs.

References

[1] Zeger, S.L., Irizarry, R. and Peng, R.D. (2006) On Time Series Analysis of Public Health and Biomedical Data. *Annual Review of Public Health*, **27**, 57-79. https://doi.org/10.1146/annurev.publhealth.26.021304.144517

[2] Bhaskaran, K., Gasparrini, A., Hajat, S., Smeeth, L. and Armstrong, B. (2013) Time Series Regression Studies in Environmental Epidemiology. *International Journal of Epidemiology*, **42**, 1187-1195. https://doi.org/10.1093/ije/dyt092

[3] Cook, D.A., Coory, M. and Webster, R.A. (2011) Exponentially Weighted Moving Average Charts to Compare Observed and Expected Values for Monitoring Risk-Adjusted Hospital Indicators. *BMJ Quality & Safety*, **20**, 469-474. https://doi.org/10.1136/bmjqs.2008.031831

[4] Pilcher, D.V., Hoffman, T., Thomas, C., Ernest, D. and Hart, G.K. (2010) Risk-Adjusted Continuous Outcome Monitoring with a EWMA Chart: Could It Have Detected Excess Mortality among Intensive Care Patients at Bundaberg Base Hospital? *Critical Care and Resuscitation*, **12**, 36-41.

[5] Moran, J., Solomon, P. and ANZICS Centre for Outcome and Resource Evaluation (2013) Statistical Process Control of Mortality Series in the Australian and New Zealand Intensive Care Society (ANZICS) Adult Patient Database: Implications of the Data Generating Process. *BMC Medical Research Methodology*, **13**, 66. https://doi.org/10.1186/1471-2288-13-66

[6] Miller, M.B. (2014) Mathematics & Statistics for Financial Risk Management. John Wiley & Sons Ltd., Hoboken.

[7] Francq, C. and Zakoian, J.M. (2010) Classical Time Series Models and Financial Series. GARCH Models: Structure, Statistical Inference and Financial Applications. John Wiley & Sons Ltd., Chichester, 1-15. https://doi.org/10.1002/9780470670057

[8] Taleb, N.N. (2009) Errors, Robustness, and the Fourth Quadrant. *International Journal of Forecasting*, **25**, 744-759. https://doi.org/10.1016/j.ijforecast.2009.05.027

[9] Goldstein, D.G. and Taleb, N.N. (2007) We Don't Quite Know What We Are Talking about. *Journal of Portfolio Management*, **33**, 84-86. https://doi.org/10.3905/jpm.2007.690609

[10] Tsay, R.S. (2013) An Introduction to Analysis of Financial Data with R. John Wiley & Sons Ltd., Hoboken.

[11] Tsay, R.S. (2014) Multivariate Time Series Analysis: With R and Financial Applications. John Wiley & Sons Ltd., Hoboken.

[12] Aldridge, I. (2013) High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems. John Wiley & Sons Ltd., Hoboken.

[13] Gabrys, R., Horman, S. and Kokoszka, P. (2013) Monitoring the Intraday Volatility Pattern. *Journal of Time Series Econometrics*, **5**, 87-116. https://doi.org/10.1515/jtse-2012-0006

[14] Hautsch, N. (2012) Modelling Hogh-Frequency Volatility. Econometrics of Financial High-Frequency Data. Springer-Verlag, Berlin, 195-223. https://doi.org/10.1007/978-3-642-21925-2_8

[15] Wang, Z. (2013) Cts: An R Package for Continuous Time Autoregressive Models via Kalman Filter. *Journal of Statistical Software*, **53**, 1-19. https://doi.org/10.18637/jss.v053.i05

[16] Mandelbrot, B. (1963) The Variation of Certain Speculative Prices. *Journal of Business*, **36**, 394-419. https://doi.org/10.1086/294632

[17] Stow, P.J., Hart, G.K., Higlett, T., George, C., Herkes, R., McWilliam, D. and Bellomo, R. (2006) Development and Implementation of a High-Quality Clinical Database: The Australian and New Zealand Intensive Care Society Adult Patient Database. *Journal of Critical Care*, **21**, 133-141. https://doi.org/10.1016/j.jcrc.2005.11.010

[18] Francq, C. and Zakoian, J.M. (2010) GARCH Models: Structure, Statistical Infe-

rence and Financial Applications. John Wiley & Sons Ltd., Chichester.
https://doi.org/10.1002/9780470670057

[19] Tsay, R.S. (2010) Conditional Heteroscedastic Models. Analysis of Financial Time Series. John Wiley & Sons Ltd., Chichester, 109-173.
https://doi.org/10.1002/9780470644560.ch3

[20] Francq, C. and Zakoian, J.M. (2010) Multivariate GARCH Processes. GARCH Models: Structure, Statistical Inference and Financial Applications. John Wiley & Sons Ltd., Chichester, 273-310. https://doi.org/10.1002/9780470670057.ch11

[21] Bauwens, L., Laurent, S. and Rombouts, J.V. (2006) Multivariate GARCH Models: A Survey. *Journal of Applied Econometrics*, **21**, 79-109.
https://doi.org/10.1002/jae.842

[22] Silvennoimen, A. and Terrasvirta, T. (2009) Multivariate GARCH Models. In: Andersen, T.G., Davies, R.A., Kreib, R.A. and Mikosch, T., Eds., *Handbook of Financial Time Series*, Springer, New York, 201-229.
https://doi.org/10.1007/978-3-540-71297-8_9

[23] Laurent, S., Rombouts, J.V. and Violante, F. (2012) On the Forecasting Accuracy of Multivariate GARCH Models. *Journal of Applied Econometrics*, **27**, 934-955.
https://doi.org/10.1002/jae.1248

[24] Moran, J.L., Solomon, P.J. and Adult Database Management Committee of the Australian and New Zealand Intensive Care Society (2011) Conventional and Advanced Time Series Estimation: Application to the Australian and New Zealand Intensive Care Society (ANZICS) Adult Patient Database, 1993-2006. *Journal of Evaluation in Clinical Practice*, **17**, 45-60. https://doi.org/10.1111/j.1365-2753.2010.01368.x

[25] ANZICS Centre for Outcome and Resource Evaluation (2012) APD Data Dictionary for Software Programmers: Version 5.4 January 2017.
http://www.anzics.com.au/Downloads/ANZICS%20APD%20Dictionary%20Programmers%20V5.4.pdf

[26] Knaus, W.A., Wagner, D.P., Draper, E.A., Zimmerman, J.E., Bergner, M., Bastos, P.G., Sirio, C.A., Murphy, D.J., Lotring, T. and Damiano, A. (1991) The APACHE III Prognostic System. Risk Prediction of Hospital Mortality for Critically ILL Hospitalized Adults. *Chest*, **100**, 1619-1636. https://doi.org/10.1378/chest.100.6.1619

[27] Solomon, P., Kasza, J., Moran, J. and Australian and New Zealand Intensive Care Society (ANZICS) Centre for Outcome and Resource Evaluation (2014) Identifying Unusual Performance in Australian and New Zealand Intensive Care Units from 2000 to 2010. *BMC Medical Research Methodology*, **14**, 53.
https://doi.org/10.1186/1471-2288-14-53

[28] Stata Corporation LP (2015) Stata Statistical Software: Release 14.

[29] Laurent, S. (2013) Estimating and Forecasting ARCH Models Using G@RCH™ 7. Timberlake Consultants Ltd., London.

[30] Doornik, J.A. (2013) OxMetrics™ 7: A Software System for Data Analysis and Forecasting. Timberlake Consultants Ltd., London.

[31] Hyndman, R.J. and Khandakar, Y. (2008) Automatic Time Series Forecasting: The Forecast Package for R. *Journal of Statistical Software*, **27**, 1-22.
https://doi.org/10.18637/jss.v027.i03

[32] R Development Core Team (2009) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna.
http://www.R-project.org

[33] Gelade, W., Veradi, V. and Vermandele, C. (2015) Time-Efficient Algorithms for

Robust Estimators of Location, Scale, Symmetry, and Tail Heaviness. *Stata Journal*, **15**, 77-94.

[34] Brys, G., Hubert, M. and Struyf, A. (2008) Goodness-of-Fit Tests Based on a Robust Measure of Skewness. *Computational Statistics*, **23**, 429-442. https://doi.org/10.1007/s00180-007-0083-7

[35] De Livera, A.M., Hyndman, R.J. and Snyder, R.D. (2011) Forecasting Time Series with Complex Seasonal Patterns Using Exponential Smoothing. *Journal of the American Statistical Association*, **106**, 1513-1527. https://doi.org/10.1198/jasa.2011.tm09771

[36] Tsay, R.S. (2013) Asset Volatility and Volatility Models. An Introduction to Analysis of Financial Data with R. John Wiley & Sons, Hoboken, 176-241.

[37] Karlsson, S. (2009) ARMADIAG: Stata Module to Compute Post-Estimation Residual Diagnostics for Time Series. http://econpapers.repec.org/scripts/search.asp?ft=armadiag

[38] Engle, R. (2001) GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics. *The Journal of Economic Perspectives*, **15**, 157-168. https://doi.org/10.1257/jep.15.4.157

[39] Engle, R.F. (1982) Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United-Kingdom Inflation. *Econometrica*, **50**, 987-1007. https://doi.org/10.2307/1912773

[40] Hansen, P.R. and Lunde, A. (2005) A Forecast Comparison of Volatility Models: Does Anything Beat a GARCH (1, 1)? *Journal of Applied Econometrics*, **20**, 873-889. https://doi.org/10.1002/jae.800

[41] Bollerslev, T. (1986) Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, **31**, 307-327. https://doi.org/10.1016/0304-4076(86)90063-1

[42] Caporin, M. and McAleer, M. (2012) Model Selection and Testing of Conditional and Stochastic Volatility Models. In: Bauwens, L., Hafner, C. and Laurent, S., Eds., *Volatility Models and Their Applications*, John Wiley & Sons, Hoboken, 199-221. https://doi.org/10.1002/9781118272039.ch8

[43] Engle, R.F. and Patton, A.J. (20001) What Good Is a Volatility Model? *Quantitative Finance*, **1**, 237-245. https://doi.org/10.1088/1469-7688/1/2/305

[44] Nelson, D.B. (1991) Conditional Heteroskedasticity in Asset Returns—A New Approach. *Econometrica*, **59**, 347-370. https://doi.org/10.2307/2938260

[45] Glosten, L.R., Jagannathan, R. and Runkle, D.E. (1993) On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *The Journal of Finance*, **48**, 1779-1801. https://doi.org/10.1111/j.1540-6261.1993.tb05128.x

[46] Ding, Z., Granger, C.W.J. and Engle, R.F. (1993) A Long Memory Property of Stock Market Returns and a New Model. *Journal of Empirical Finance*, **1**, 83-106. https://doi.org/10.1016/0927-5398(93)90006-D

[47] Laurent, S. (2014) Multivariate GARCH Models. Estimating and Forecasting ARCH Models Using G@RCH^{TM} 7. Timberlake Consultants Ltd., London, 235-302.

[48] Engle, R. and Kroner, K.F. (1995) Multivariate Simultaneous Generalized ARCH. *Econometric Theory*, **11**, 122-150. https://doi.org/10.1017/S0266466600009063

[49] Tsay, R.S. (2014) Multivariate Volatility Models. Multivariate Time Series Analysis: With R and Financial Applications. John Wiley & Sons, Hoboken, 399-464.

[50] Van der Weide, R. (2002) Go-GARCH: A Multivariate Generalized Orthogonal

GARCH Model. *Journal of Applied Econometrics*, **17**, 549-564.
https://doi.org/10.1002/jae.688

[51] Van der Weide, R. (2012) The Time-Variation of Volatility and the Evolution of Expectations. http://dare.uva.nl/record/411121

[52] Nelson, D.B. (1992) Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model. *Journal of Econometrics*, **52**, 61-90. https://doi.org/10.1016/0304-4076(92)90065-Y

[53] Andersen, T.G. and Bollerslev, T. (1998) Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts. *International Economic Review*, **39**, 885-905. https://doi.org/10.2307/2527343

[54] Hansen, P.R. and Lunde, A. (2011) Forecasting Volatility Using High-Frequency Data. In: Clements, M.P. and Hendry, D.F., Eds., *The Oxford Handbook of Economic Forecasting*, Oxford University Press, New York, 525-556.
https://doi.org/10.1093/oxfordhb/9780195398649.013.0020

[55] Efimova, O. and Serletis, A. (2014) Energy Markets Volatility Modelling Using GARCH. *Energy Economics*, **43**, 264-273.
https://doi.org/10.1016/j.eneco.2014.02.018

[56] Sadorsky, P. (2012) Correlations and Volatility Spillovers between Oil Prices and the Stock Prices of Clean Energy and Technology Companies. *Energy Economics*, **34**, 248-255. https://doi.org/10.1016/j.eneco.2011.03.006

[57] Wang, Y.D. and Wu, C.F. (2012) Forecasting Energy Market Volatility Using GARCH Models: Can Multivariate Models Beat Univariate Models? *Energy Economics*, **34**, 2167-2181. https://doi.org/10.1016/j.eneco.2012.03.010

[58] Bollerslev, T. (1990) Modelling the Coherence in Short-Run Nominal Exchange Rates: A Multivariate Generalized ARCH Model. *Review of Economics and Statistics*, **72**, 498-505. https://doi.org/10.2307/2109358

[59] Engle, R. (2002) Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics*, **20**, 339-350.
https://doi.org/10.1198/073500102288618487

[60] Kuha, J. (2004) AIC and BIC: Comparisons of Assumptions and Performance. *Sociological Methods Research*, **33**, 188-229. https://doi.org/10.1177/0049124103262065

[61] Baum, C.F. (2008) Sts15: Tests for Stationarity of a Time Series. *Stata Technical Bulletin*, **10**, 356-362.

[62] Shumway, R.H. and Stoffer, D.S. (2011) Additional Time Domain Topics. Time Series Analysis and Its Applications with R Examples. Springer Science + Business Media, New York, 267-318. https://doi.org/10.1007/978-1-4419-7865-3_5

[63] Schwert, G.W. (1987) Effects of Model-Specification on Tests for Unit Roots in Macroeconomic Data. *Journal of Monetary Economics*, **20**, 73-103.
https://doi.org/10.1016/0304-3932(87)90059-6

[64] Stata Corporation (2013) Arima—ARIMA, ARMAX, and Other Dynamic Regression Models. http://www.stata.com/manuals13/tsarima.pdf

[65] Hyndman, R.J. and Koehler, A.B. (2006) Another Look at Measures of Forecast Accuracy. *International Journal of Forecasting*, **22**, 679-688.
https://doi.org/10.1016/j.ijforecast.2006.03.001

[66] Diebold, F.X. and Mariano, R.S. (1995) Comparing Predictive Accuracy. *Journal of Business & Economic Statistics*, **13**, 253-263.
https://doi.org/10.1080/07350015.1995.10524599

[67] Francq, C. and Zakoian, J.M. (2010) Asymmetries. GARCH Models: Structure, Statistical Inference and Financial Applications, John Wiley & Sons Ltd., Chichester, 245-270. https://doi.org/10.1002/9780470670057.ch10

[68] Ljung, G.M. and Box, G.E.P. (1978) Measure of Lack of Fit in Time-Series Models. *Biometrika*, **65**, 297-303. https://doi.org/10.1093/biomet/65.2.297

[69] Lopez, J.A. (2001) Evaluating the Predictive Accuracy of Volatility Models. *Journal of Forecasting*, **20**, 87-109.
https://doi.org/10.1002/1099-131X(200103)20:2<87::AID-FOR782>3.0.CO;2-7

[70] Drost, F.C. and Nijman, T.E. (1993) Temporal Aggregation of Garch Processes. *Econometrica*, **61**, 909-927. https://doi.org/10.2307/2951767

[71] Baillie, R.T. and Bollerslev, T. (1989) The Message in Daily Exchange-Rates—A Conditional-Variance Tale. *Journal of Business & Economic Statistics*, **7**, 297-305. https://doi.org/10.1080/07350015.1989.10509739

[72] Engle, R. (2004) Risk and Volatility: Econometric Models and Financial Practice. *American Economic Review*, **94**, 405-420.
https://doi.org/10.1257/0002828041464597

[73] Andersen, T.G., Bollerslev, T., Christoffersen, P.F. and Diebold, F.X. (2006) Volatility and Correlation Forecasting. In: Elliott, J., Granger, C.W.J. and Timmermann, A., Eds., *Handbook of Economic Forecasting*: *Volume* 1, Elsevier, Amsterdam, 777-878. https://doi.org/10.1016/S1574-0706(05)01015-3

[74] Lee, R. (2000) The Lee-Carter Method for Forecasting Mortality, with Various Extensions and Applications. *North American Actuarial Journal*, **4**, 80-91.
https://doi.org/10.1080/10920277.2000.10595882

[75] Alders, M., Keilman, N. and Cruijsen, H. (2007) Assumptions for Long-Term Stochastic Population Forecasts in 18 European Countries. *European Journal of Population*, **23**, 33-69. https://doi.org/10.1007/s10680-006-9104-4

[76] Chai, C.M., Siu, T.K. and Zhou, X. (2014) A Double-Exponential GARCH Model for Stochastic Mortality. *European Actuarial Journal*, **3**, 385-406.
https://doi.org/10.1007/s13385-013-0077-5

[77] Alho, J. (2008) Aggregation across Countries in Stochastic Population Forecasts. *International Journal of Forecasting*, **24**, 343-353.
https://doi.org/10.1016/j.ijforecast.2008.05.003

[78] Gao, Q.S. and Hu, C.J. (2009) Dynamic Mortality Factor Model with Conditional Heteroskedasticity. *Insurance Mathematics & Economics*, **45**, 410-423.
https://doi.org/10.1016/j.insmatheco.2009.09.001

[79] Bauwens, L., Hafner, C.M. and Pierret, D. (2013) Multivariate Volatility Modeling of Electricity Futures. *Journal of Applied Econometrics*, **28**, 743-761.
https://doi.org/10.1002/jae.2280

[80] Caporin, M. and McAleer, M. (2013) Ten Things You Should Know about the Dynamic Conditional Correlation Representation.
http://eprints.ucm.es/21803/1/1320.pdf

[81] Podobnik, B., Wang, D., Horvatic, D., Grosse, I. and Stanley, H. (2010) Time-Lag Cross-Correlations in Collective Phenomena. *Europhysics Letters*, **90**, 1632-1652.
https://doi.org/10.1209/0295-5075/90/68001

[82] Kellstedt, P.M., Linn, S. and Hannah, A. (2015) The Polls-Review the Usefulness of Consumer Sentiment: Assessing Construct and Measurement. *Public Opinion Quarterly*, **79**, 181-203. https://doi.org/10.1093/poq/nfu056

[83] Lebo, M.J. and Box-Steffensmeier, J.M. (2008) Dynamic Conditional Correlations in

Political Science. *American Journal of Political Science*, **52**, 688-704. https://doi.org/10.1111/j.1540-5907.2008.00337.x

[84] Moran, J.L., Bristow, P., Solomon, P.J., George, C., Hart, G.K. and Australian and New Zealand Intensive Care Society Database Management Committee (2008) Mortality and Length-of-Stay Outcomes, 1993-2003, in the Binational Australian and New Zealand Intensive Care Adult Patient Database. *Critical Care Medicine*, **36**, 46-61. https://doi.org/10.1097/01.CCM.0000295313.08084.58

[85] Moran, J. and Solomon, P. (2012) Mortality and Intensive Care Volume in Ventilated Patients, 1995-2009, in the Australian and New Zealand Bi-National Adult Patient Intensive Care Database. *Critical Care Medicine*, **40**, 800-812. https://doi.org/10.1097/CCM.0b013e318236f2af

[86] Marshall, E.C., Best, N., Bottle, A. and Aylin, P. (2004) Statistical Issues in the Prospective Monitoring of Health Outcomes across Multiple Units. *Journal of the Royal Statistical Society Series A*, **167**, 541-559. https://doi.org/10.1111/j.1467-985X.2004.apm10.x

[87] Malmsten, H. (2004) Evaluating Exponential GARCH Models. SSE/EFI Working Papers in Economics and Finance. The Economic Research Institute, Stockholm School of Economics, Stockholm, 564.

[88] Laurent, S. (2013) Further Univariate GARCH Models. Estimating and Forecasting ARCH Models Using G@RCH™ 7. Timberlake Consultants Ltd., London, 79-117.

[89] Engle, R.F. and Ng, V.K. (1993) Measuring and Testing the Impact of News on Volatility. *Journal of Finance*, **48**, 1749-1778. https://doi.org/10.1111/j.1540-6261.1993.tb05127.x

[90] Daly, K. (2008) Financial Volatility: Issues and Measuring Techniques. *Physica A—Statistical Mechanics and Its Applications*, **387**, 2377-2393. https://doi.org/10.1016/j.physa.2008.01.009

[91] Higgs, H. (2009) Modelling Price and Volatility Inter-Relationships in the Australian Wholesale Spot Electricity Markets. *Energy Economics*, **31**, 748-756. https://doi.org/10.1016/j.eneco.2009.05.003

[92] Bauwens, L., Hafner, C. and Laurent, S. (2012) Volatility Models. In: Bauwens, L., Hafner, C. and Laurent, S., Eds., *Handbook of Volatility Models and Their Applications*, John Wiley & Sons, Hoboken, 1-45. https://doi.org/10.1002/9781118272039.ch1

Appendix 1

i) The GARCH(p,q) model, as formulated by Bollerslev [41];

$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2$, where ω, α_i and β_j are constants and ε_t^2 are the squared residuals (innovations: $\varepsilon_t = z_t \sigma_t$ and z_t is i.i.d., $E(z_t) = 0$, $var(z_t) = 1$). This model is an obvious comparator and in some financial data (exchange rates) it outperforms more sophisticated models [40].

ii) The exponential GARCH(p,q) model of Nelson (EGARCH [44]).

$\ln(\sigma_t^2) = \omega + \alpha |z_{t-1}| + \gamma z_{t-1} + \beta \ln(\sigma_{t-1}^2)$; where the γ parameter indicates leverage (if $\gamma < 0$ and $\gamma < \alpha < -\gamma$) [42]. Different formulations and software implementations of the EGARCH model exist and we provide a minimal equation where $p = q = 1$ [87].

iii) The GJR-(Glosten, Jagaannathan and Runkle [45])-GARCH model;

$\sigma_t^2 = \omega + \sum_{i=1}^{q} (\alpha_i \varepsilon_{t-i}^2 + \gamma_i S_{t-i}^- \varepsilon_{t-i}^2) + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2$, where S_t^- is a dummy variable of

value 1 when ε_t is negative and 0 otherwise and the model assumes that the sign of ε_t (positive or negative) is determinant of the impact of ε_t^2 on the conditional variance σ_t^2 [88].

iv) The asymmetric power GARCH (APGARCH (p, q)), as described by Ding *et al* [46];

$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i (|\varepsilon_{t-i}| - \gamma_i \varepsilon_{t-i})^\delta + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^\delta$, where $\omega > 0, \delta > 0, \alpha_i \geq 0, \beta_j \geq 0$

and $|\gamma_i| \leq 1$. The δ parameter performs a Box-Cox type transformation of σ_t and γ_i reflects the "leverage" effect [89]; $|\gamma_i| \leq 1$ is a non-restrictive identifiability constraint [67]. If the sum of (scalar) α and β (a persistence coefficient [90] [91]) < 1, the conditional volatility process is mean reverting and shocks are transitory. First published in 1993, this is an encompassing model to the extent that it includes the ARCH, GARCH and GJR-GARCH models as special cases [88].

Appendix 2

i) GO-GARCH [50]. In orthogonal GARCH models, the observed data are assumed to be generated by an orthogonal transform of N (or <N) univariate GARCH processes, which may be considered as factors [47]. In GO-GARCH, the orthogonality condition is relaxed [21]; the original series (contained in r_t) are linked to unobserved, uncorrelated factors (z_t; in GO-GARCH, equal to the series number) through a linear, invertible transformation (W, a non-singular $N \times N$ matrix). The conditional covariance matrix of r_t is expressed as $H_t = W H_t^z W' = \sum_{k=1}^{N} w_{(k)} w_{(k)}' h_{k,t}^z$, where $w_{(k)}$ are the columns of W and $h_{k,t}^z$ are the diagonal elements of H_t^z [22].

ii) Conditional correlation models: use nonlinear combinations of univariate GARCH models to represent the conditional covariances, which are decomposed into conditional variances and correlations. The diagonal elements of H_t (the conditional covariance matrix) are modelled as univariate GARCH

models, whereas the off-diagonal elements are modelled as nonlinear functions of the diagonal terms.

a) constant conditional correlation (CCC) [58]; here the conditional correlations are (time) invariant and the conditional covariances are proportional to the product of corresponding standard deviations. The series (r_t) are modelled as $\mathrm{GARCH}(p,q)$. Although we estimate the CCC model, it is used as a comparator as the "… hypothesis of CCCs is not tenable except for specific cases and short periods" [92].

b) dynamic conditional correlation (DCC) [59]. The CCC model is generalised such that the conditional correlations change over time: $h_{ij,t} = \rho_{ij,t} \sqrt{h_{ii,t} h_{jj,t}}$, where the diagonal elements $h_{ii,t}$ and $\mathrm{h}_{jj,t}$ follow univariate GARCH processes and $\rho_{ij,t}$ follows a time-varying dynamic process. A constraint of this model is that "… all correlations have the same dynamic pattern… [and] … revert to a constant level" [92].