Bootstrap Prediction Intervals for ARCH Models

by

Jonathan J. Reeves *
Department of Economics
Queen's University
Kingston, Canada

email: reevesj@qed.econ.queensu.ca

April 2000

Abstract

In this paper we construct prediction intervals for ARCH models using the bootstrap. We use both a parametric and non-parametric bootstrap, which take account of parameter uncertainty. We compare our prediction intervals to traditional asymptotic prediction intervals, and find that the bootstrap leads to improved accuracy. The accuracy of the bootstrap is empirically demonstrated with the Yen/$US exchange rate.

KEY WORDS: ARCH models; Bootstrap; Prediction intervals; Foreign exchange rates.

*I would like to thank Allan Gregory for his insight and guidance. I also thank James MacKinnon for helpful comments.

1 Introduction

The Autoregresive Conditional Heteroskedasticity (ARCH) class of models was originally introduced by Engle (1982) and has become a core part of empirical finance. The issue of forecasting with these models has been discussed by Koenker and Zhao (1996), Baillie and Bollerslev (1992), Granger, White, and Kamstra (1989), Geweke (1989), Diebold (1988), and Engle and Kraft (1983). However, the procedures developed are restricted to a limited class of ARCH models, often do not take account of parameter uncertainty, and often have questionable finite sample properties. Baillie and Bollerslev (1992) Gaussian asymptotic prediction intervals were developed for the class of ARMA-GARCH models. However, for other ARCH models, such as ARCH-in-mean, Engle, Lilien, and Robins (1987), analytical closed form formulas for prediction intervals do not currently exist.

Recently, financial practitioners have become increasingly interested in the forecasting of the distribution of future asset prices. Modern financial instruments have enabled practitioners access to the pay-offs from any part of the distribution of future values of an asset, thus accurate measures of this entire distribution is very important. Value at Risk (VaR) calculations, which involve the calculation of an estimate of a quantile of the distribution of future values of an asset(s), are widespread. For an introduction to VaR, see Schachter (1998).

This paper proposes bootstrap procedures to estimate the future distribution of an asset which follows an ARCH process. This procedure is applicable to the general class of ARCH models, which includes the ARCH-in-mean models. Our main measure of the future distribution is through the construction of prediction intervals. We compare the accuracy of our prediction intervals against the Baillie and Bollerslev (BB) Gaussian asymptotic in a Monte Carlo experiment. The results strongly favour the bootstrap.

There are two distinct types of bootstrap prediction interval: Parametric and the non-parametric. The parametric bootstrap prediction interval was first discussed by Beran (1990) and Thombs and Schucany's (1990) developed the non-parametric bootstrap prediction interval for AR processes. We extend the application of these bootstrap techniques to ARCH processes.

An advantage of the bootstrap prediction intervals is they can account for parameter uncertainty. We show that asymptotically, the procedure when applied to ARCH processes has the correct coverage. Moreover, our Monte Carlo study shows it also does well in finite samples. We find the performance of the BB asymptotic prediction intervals to often be disappointing.

The non-parametric bootstrap prediction interval is especially useful to the applied researchers for in situations in which there is no clear-cut parametric model. We show in a Monte Carlo study that the non-parametric bootstrap prediction interval for non-normal distributions is much more accurate than the Gaussian BB prediction interval.

Finally, we look at the empirical performance of the bootstrap relative to the Gaussian BB. This is done by fitting AR(1)-GARCH(1,1) models to the Yen/$US exchange rate and assessing the one-step ahead predictions. We find the non-parametric bootstrap to do extremely well, whereas the Gaussian BB often does not do well.

The paper is organised as follows. Section 2 describes the general ARCH model. Section 3 describes the method for prediction interval construction, and Section 4 gives the asymptotic justification for the procedure. Section 5 provides Monte Carlo evidence using sample sizes consistent with empirical finance applications. Section 6 examines the empirical performance of these procedures with the Yen/$US and Section 7 concludes.

2 The ARCH Model

Let Y_t , $t = 1, \ldots, \infty$ be a sequence of scalar random variables, $\{y_t : t = T-n+1, \ldots, T\}$ be a realisation, and $\boldsymbol{x}_t = (y_{t-1}, y_{t-2}, \ldots, y_{T-n+1})$ denote the predetermined variables.

The conditional mean and variance functions are jointly parameterized by a finite dimensional vector $\boldsymbol{\theta}$:

$$\{\mu_t(\boldsymbol{x}_t, \boldsymbol{\theta}) : \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$$

$$\{\Omega_t(\boldsymbol{x}_t, \boldsymbol{\theta}) : \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$$

where $\boldsymbol{\Theta}$ is a compact subset of $\Re^p$ that has nonempty interior, and μ_t and Ω_t are known continuous functions of $\boldsymbol{x}_t$ and $\boldsymbol{\theta}$ that are twice continuously differentiable on the interior of $\boldsymbol{\Theta}$ for all $\boldsymbol{x}_t$.

Further, the first two conditional moments are correctly specified. For some $\boldsymbol{\theta}_0 \in$ *int* $\boldsymbol{\Theta}$

$$E(y_t|\boldsymbol{x}_t) = \mu_t(\boldsymbol{x}_t, \boldsymbol{\theta}_0)$$

$$V(y_t|\boldsymbol{x}_t) = \Omega_t(\boldsymbol{x}_t, \boldsymbol{\theta}_0) \ , \qquad t = 1, 2, \ldots$$

with

$$y_t = \mu_t(\boldsymbol{x}_t, \boldsymbol{\theta}_0) + \varepsilon_t^0 \Omega_t^{\frac{1}{2}}(\boldsymbol{x}_t, \boldsymbol{\theta}_0)$$
$$E(\varepsilon_t^0|\boldsymbol{x}_t) = 0$$
$$E(\varepsilon_t^{0^2}|\boldsymbol{x}_t) = 1$$

Finally, the conditional variance satisfies:

$$0 < \Omega_t(\boldsymbol{x}_t, \boldsymbol{\theta}) < \infty \text{ for all } \boldsymbol{\theta} \in \boldsymbol{\Theta}.$$

3 Bootstrap Prediction Intervals for ARCH Models

The bootstrap procedure for prediction interval construction can be broken down into four steps.

Step 1: Fit the ARCH Model

Fit an ARCH model to the data. In the non-parametric setting, a Gaussian likelihood would be used since this quasi-maximum likelihood estimator is generally consistent, and has reasonable finite sample properties, see Bollerslev and Wooldridge (1992). Of course, in applied work the exact form of the model would be determined using a series of diagnostic tests and model selection criteria and would have consequences for any prediction interval calculation. Denote the estimated parameters by $\hat{\boldsymbol{\theta}}$ and the associated residual vector by $\hat{\boldsymbol{\varepsilon}}$.

Step 2: Generate Bootstrap Draws of the Parameters

Use the fitted model to generate bootstrap draws of the parameters. Specifically, we take the fitted model and use the original parameter estimates ($\hat{\boldsymbol{\theta}}$) and any initial conditions, along with random draws from the pool of residuals[2] from the fitted model, $\hat{\boldsymbol{\varepsilon}}$, to generate the dependent variable over the sample period. For the parametric bootstrap case, the random draws are obtained from a specified distribution with mean 0 and variance 1 (such as a standard normal.) Use this generated sample of the dependent variable to re-estimate the original model, and denote $\hat{\boldsymbol{\theta}}^*$ as one bootstrap draw of the set of parameters in the model. Do B draws of the parameters using the bootstrap.

Step 3: Generate Future Realisations

For every bootstrap draw of the parameters ($\hat{\boldsymbol{\theta}}^*$) generate a future realisation vector of the dependent variable. A future realisation vector has first element being a realisation of the dependent variable one period into the future. A future realisation vector has dimension K, where K is the number of forecast periods. To generate a future realisation vector, generate dynamic forecasts with the model using the corresponding bootstrap draw of the parameters ($\hat{\boldsymbol{\theta}}^*$) along with initial conditions, (typically from the last time periods where there are observations), and random draws from $\hat{\boldsymbol{\varepsilon}}$ for the non-parametric case, or a specified distribution with mean 0 and variance 1 for the parametric case.

Step 4: Generate the $100(1-\alpha)\%$ Prediction Intervals

Sort all the future realisations, corresponding to a future time period, $T+k$, from lowest to highest. Trim off the top $100(\frac{\alpha}{2})\%$ of realisations and the bottom $100(\frac{\alpha}{2})\%$ of realisations. Denote the end points of the interval as $L^*_{k,B}$ and $U^*_{k,B}$. This gives the $100(1-\alpha)\%$ bootstrap prediction interval for time period $T+k$.

[2] These empirical residuals should be rescaled to have a mean of 0 and a variance of 1.

4 Asymptotic Justification

Under the conditions of the ARCH model given in Section 2 and the technical assumptions in Appendix A of Bollerslev and Wooldridge (1992) - namely conditions A.1 (iii) - (vi), the quasi-maximum likelihood estimator is generally consistent for $\boldsymbol{\theta}_0$. (See Bollerslev and Wooldridge (1992).) [3] The bootstrap estimator is also consistent (see Bickel and Freedman (1981).) With these results, we can establish the asymptotic coverage for the bootstrap prediction interval.

First we establish $Y^*_{T+k} \xrightarrow{d} Y_{T+k}$. This proof is by induction and is similar to Thombs and Schucany (1990). For $k = 1$, it follows,

$$Y^*_{T+1} = \mu_{T+1}(\boldsymbol{x}_{T+1}, \hat{\boldsymbol{\theta}}^*) + \hat{\varepsilon}^*_{T+1}\Omega^{\frac{1}{2}}_{T+1}(\boldsymbol{x}_{T+1}, \hat{\boldsymbol{\theta}}^*)$$

The lead term converges in conditional probability to $\mu_{T+1}(\boldsymbol{x}_{T+1}, \boldsymbol{\theta}_0)$ by Slutsky's theorem and the consistency of $\hat{\boldsymbol{\theta}}^*$.

The term, $\hat{\varepsilon}^*_{T+1}$, has distribution function $\hat{F}_\varepsilon$ and $|\hat{F}_\varepsilon - F_\varepsilon| \xrightarrow{d} 0$ or $\hat{\varepsilon}^* \xrightarrow{d} \varepsilon$. (Bickel and Freedman (1981).) For similar reasons, the last term, $\Omega^{\frac{1}{2}}_{T+1}(\boldsymbol{x}_{T+1}, \hat{\boldsymbol{\theta}}^*)$, converges in conditional probability to $\Omega^{\frac{1}{2}}_{T+1}(\boldsymbol{x}_{T+1}, \boldsymbol{\theta}_0)$. Therefore $Y^*_{T+1} \xrightarrow{d} Y_{T+1}$. For $k = 2$,

$$Y^*_{T+2} = \mu_{T+2}(\boldsymbol{X}^*_{T+2}, \hat{\boldsymbol{\theta}}^*) + \hat{\varepsilon}^*_{T+2}\Omega^{\frac{1}{2}}_{T+2}(\boldsymbol{X}^*_{T+2}, \hat{\boldsymbol{\theta}}^*)$$

Applying the results for $k = 1$,

$$\mu_{T+2}(\boldsymbol{X}^*_{T+2}, \hat{\boldsymbol{\theta}}^*) \xrightarrow{d} \mu_{T+2}(\boldsymbol{X}_{T+2}, \boldsymbol{\theta}_0)$$

$$\text{and} \quad \Omega^{\frac{1}{2}}_{T+2}(\boldsymbol{X}^*_{T+2}, \hat{\boldsymbol{\theta}}^*) \xrightarrow{d} \Omega^{\frac{1}{2}}_{T+2}(\boldsymbol{X}_{T+2}, \boldsymbol{\theta}_0)$$

which implies $Y^*_{T+2} \xrightarrow{d} Y_{T+2}$. Induction completes the proof. This proof was for the non-parametric case and the same result is obviously true for the parametric case when ε_t is drawn from the true distribution with mean 0 and variance 1.[4]

It also follows that $Q^*(u) \longrightarrow Q(u)$ (pointwise for $0 < u < 1$), where Q^* is the theoretical bootstrap quantile function and Q is the true quantile function for Y_{T+k}. Since a finite value of B is used in practice, the endpoints of the interval are estimates of $Q^*[\alpha/2]$ and $Q^*[1 - \alpha/2]$. $Q^*_B(u) \xrightarrow{p} Q^*(u)$ as $B \longrightarrow \infty$ from the Glivenko-Cantelli theorem.

[3] Weaker conditions are sufficient for the GARCH(1,1) and IGARCH(1,1) models, see Lee and Hansen (1994) and Lumsdaine (1996).

[4] Relative to the general non-parametric case, weaker conditions are sufficient for the consistency of $\hat{\boldsymbol{\theta}}$ in the parametric case.

This establishes the asymptotic justification for the ARCH parametric and non-parametric $100(1-\alpha)\%$ bootstrap prediction interval:

$$\lim_{n\to\infty}[\lim_{B\to\infty} Pr\{L^*_{k,B}(\boldsymbol{x}_{T+1}) \leq Y_{T+k} \leq U^*_{k,B}(\boldsymbol{x}_{T+1})\}] = 1-\alpha$$

5 Monte Carlo Evidence

Two different Monte Carlo experiments are conducted to compare bootstrap prediction intervals with the BB asymptotic intervals.[1] The first is where the DGP is Model 1, and the Gaussian bootstrap prediction interval (p-boot) is compared to the Gaussian BB (1992) asymptotic prediction interval. Results for this appear in Tables 1 and 2.

$$\begin{array}{lll} \text{Model 1.} & y_t = 0.9y_{t-1} + u_t & t = 1,2,\ldots,T, \qquad y_0 = 1 \\ & \omega_t^2 = 0.1 + 0.3u_{t-1}^2 + 0.6\omega_{t-1}^2 & , u_0^2 = 0.25 \ , \ \omega_0^2 = 0.5 \\ & u_t = \omega_t\varepsilon_t & \varepsilon_t \text{ i.i.d. } N(0,1) \end{array}$$

The second Monte Carlo experiment is intended to mimic a situation where the researcher mistakenly applies the Gaussian BB prediction interval but the distribution is not normal. Thus, we compare the non-parametric bootstrap prediction interval (np-boot) with the Gaussian BB asymptotic prediction interval. The DGP is Model 2, and the results appear in Tables 3 and 4.

$$\begin{array}{lll} \text{Model 2.} & y_t = 0.9y_{t-1} + u_t & t = 1,2,\ldots,T, \qquad y_0 = 1 \\ & \omega_t^2 = 0.1 + 0.3u_{t-1}^2 + 0.6\omega_{t-1}^2 & , u_0^2 = 0.25 \ , \ \omega_0^2 = 0.5 \\ & u_t = \omega_t\varepsilon_t & \varepsilon_t \text{ i.i.d. standardized } t_5 \end{array}$$

The Gaussian BB asymptotic $100(1-\alpha)\%$ prediction interval for the following model:

$$\begin{array}{ll} y_t = \phi y_{t-1} + u_t & t = 1,2,\ldots,T \\ \omega_t^2 = \alpha_0 + \alpha_1 u_{t-1}^2 + \alpha_2\omega_{t-1}^2 & \\ u_t = \omega_t\varepsilon_t & \varepsilon_t \text{ i.i.d. } N(0,1) \end{array}$$

is $(L_{k,A}, U_{k,A})$ for k-steps ahead, where

$$L_{k,A} = \hat{\phi}^k y_T + \Phi^{-1}(\frac{\alpha}{2})\sqrt{\frac{\hat{\omega}^2(1-\hat{\phi}^{2k})}{1-\hat{\phi}^2} + \frac{(\hat{\omega}_{T+1}^2 - \hat{\omega}^2)((\hat{\alpha}_1+\hat{\alpha}_2)^k - \hat{\phi}^{2k})}{\hat{\alpha}_1+\hat{\alpha}_2-\hat{\phi}^2}}$$

$$U_{k,A} = \hat{\phi}^k y_T + \Phi^{-1}(1-\frac{\alpha}{2})\sqrt{\frac{\hat{\omega}^2(1-\hat{\phi}^{2k})}{1-\hat{\phi}^2} + \frac{(\hat{\omega}_{T+1}^2 - \hat{\omega}^2)((\hat{\alpha}_1+\hat{\alpha}_2)^k - \hat{\phi}^{2k})}{\hat{\alpha}_1+\hat{\alpha}_2-\hat{\phi}^2}}$$

[1]The experiments were performed using Ox version 2.00 and the simulation code is available upon request.

$$\hat{\omega}^2 = \frac{\hat{\alpha}_0}{1-\hat{\alpha}_1-\hat{\alpha}_2} \quad , \quad \hat{\phi}^2 \neq \hat{\alpha}_1+\hat{\alpha}_2$$

and $\Phi^{-1}(p)$ refers to the p^{th} quantile in the standard normal distribution.

The following procedure is followed in our Monte Carlo experiments.

1. Simulate an ARCH process for a specified length, T and generate R = 2000 true future values of Y_{T+k} from this process with the true parameter values.
2. Use the prediction interval procedure to obtain a $100(1-\alpha)\%$ prediction interval denoted by (L^i, U^i).
3. Estimate the conditional coverage by $1-\hat{\alpha}_i = \frac{\#\{L^i \leq Y^r_{T+k} \leq U^i\}}{R}$, where Y^r_{T+k} $(r = 1, \ldots, R)$ are the true future values generated in Step 1.
4. Calculate the length of the interval $LEN_i = U^i - L^i$

This is done for both the bootstrap procedure and the BB procedure, using the same ARCH process and future values that are generated in step 1. For the bootstrap procedure $B = 2000$.

Steps 1-4 are then repeated $M = 100, (i = 1, \ldots, M)$ times to get a collection of summary statistics on $1-\hat{\alpha}$ and LEN for both prediction interval procedures.

$$\text{Mean} \ = 1 - \frac{1}{100}\sum_{i=1}^{100}\hat{\alpha}_i$$

$$\text{S.D.} \ = \{\frac{1}{99}\sum_{i=1}^{100}(\hat{\alpha}_i - \bar{\hat{\alpha}})^2\}^{\frac{1}{2}}$$

$$\text{R.M.S.E.} \ = \{\frac{1}{100}\sum_{i=1}^{100}(\hat{\alpha}_i - \alpha)^2\}^{\frac{1}{2}}$$

$$\text{Mean Length} \ = \frac{1}{100}\sum_{i=1}^{100}LEN_i$$

$$\text{S.D. of Length} \ = \{\frac{1}{99}\sum_{i=1}^{100}(LEN_i - \overline{LEN})^2\}^{\frac{1}{2}}$$

Table 1. Prediction Intervals for Model 1

Nominal Coverage	Method	T	k	Mean Coverage	S.D. of Coverage	R.M.S.E	Mean Length	S.D. of Length
0.8	p-boot	400	1	0.800	0.034	0.034	2.18	0.82
			2	0.800	0.032	0.032	2.82	0.99
			3	0.799	0.032	0.032	3.28	1.04
			4	0.796	0.036	0.036	3.62	1.08
			5	0.797	0.038	0.037	3.89	1.04
			6	0.797	0.039	0.039	4.10	1.08
	BB		1	0.800	0.032	0.032	2.18	0.82
			2	0.819	0.030	0.035	2.97	1.05
			3	0.822	0.029	0.037	3.51	1.16
			4	0.824	0.032	0.040	3.91	1.22
			5	0.827	0.035	0.044	4.23	1.25
			6	0.829	0.036	0.046	4.48	1.26
	p-boot	800	1	0.797	0.024	0.024	2.16	0.82
			2	0.798	0.022	0.022	2.80	0.99
			3	0.798	0.023	0.023	3.26	1.06
			4	0.797	0.026	0.026	3.60	1.11
			5	0.797	0.027	0.027	3.87	1.10
			6	0.796	0.028	0.028	4.08	1.08
	BB		1	0.798	0.021	0.021	2.16	0.82
			2	0.818	0.021	0.027	2.94	1.05
			3	0.821	0.021	0.030	3.48	1.16
			4	0.824	0.024	0.034	3.88	1.22
			5	0.826	0.025	0.036	4.20	1.24
			6	0.828	0.026	0.039	4.45	1.24
0.95	p-boot	400	1	0.950	0.019	0.019	3.35	1.25
			2	0.948	0.016	0.016	4.65	1.63
			3	0.948	0.016	0.016	5.49	1.80
			4	0.947	0.018	0.018	6.10	1.86
			5	0.947	0.019	0.019	6.62	1.92
			6	0.947	0.018	0.019	7.02	1.96
	BB		1	0.949	0.017	0.017	3.33	1.25
			2	0.944	0.016	0.017	4.54	1.60
			3	0.944	0.016	0.017	5.36	1.78
			4	0.944	0.017	0.018	5.98	1.87
			5	0.943	0.018	0.019	6.46	1.91
			6	0.943	0.019	0.020	6.85	1.93

Table 2. Prediction Intervals for Model 1

Nominal Coverage	Method	T	k	Mean Coverage	S.D. of Coverage	R.M.S.E	Mean Length	S.D. of Length
0.95	p-boot	800	1	0.948	0.013	0.013	3.29	1.26
			2	0.949	0.011	0.011	4.61	1.63
			3	0.948	0.012	0.012	5.44	1.80
			4	0.947	0.013	0.013	6.06	1.91
			5	0.946	0.015	0.015	6.54	1.95
			6	0.947	0.015	0.015	6.95	1.98
	BB		1	0.949	0.012	0.012	3.30	1.26
			2	0.945	0.010	0.012	4.50	1.61
			3	0.945	0.011	0.013	5.32	1.78
			4	0.944	0.013	0.014	5.93	1.87
			5	0.944	0.013	0.014	6.42	1.90
			6	0.944	0.013	0.015	6.81	1.90
0.99	p-boot	400	1	0.988	0.007	0.008	4.36	1.67
			2	0.989	0.006	0.006	6.63	2.49
			3	0.988	0.007	0.007	8.03	2.90
			4	0.988	0.007	0.007	9.06	3.03
			5	0.988	0.008	0.008	9.85	3.09
			6	0.988	0.008	0.008	10.59	3.25
	BB		1	0.989	0.006	0.006	4.38	1.65
			2	0.981	0.007	0.011	5.97	2.10
			3	0.981	0.008	0.012	7.05	2.34
			4	0.979	0.009	0.014	7.86	2.46
			5	0.978	0.009	0.015	8.50	2.51
			6	0.978	0.010	0.015	9.01	2.53
	p-boot	800	1	0.988	0.006	0.006	4.29	1.60
			2	0.988	0.005	0.005	6.52	2.26
			3	0.988	0.005	0.006	7.89	2.60
			4	0.988	0.005	0.006	8.92	2.91
			5	0.988	0.006	0.006	9.82	3.13
			6	0.988	0.005	0.006	10.34	3.05
	BB		1	0.989	0.005	0.005	4.33	1.65
			2	0.982	0.005	0.009	5.91	2.11
			3	0.981	0.006	0.011	6.99	2.34
			4	0.980	0.006	0.012	7.80	2.45
			5	0.979	0.006	0.012	8.43	2.50
			6	0.979	0.007	0.013	8.95	2.50

Table 3. Prediction Intervals for Model 2

Nominal Coverage	Method	T	k	Mean Coverage	S.D. of Coverage	R.M.S.E	Mean Length	S.D. of Length
0.80	np-boot	400	1	0.794	0.057	0.057	1.86	0.71
			2	0.793	0.053	0.053	2.48	0.82
			3	0.794	0.049	0.049	2.92	0.89
			4	0.794	0.047	0.047	3.25	0.91
			5	0.792	0.045	0.045	3.47	0.90
			6	0.793	0.046	0.046	3.68	0.90
	BB		1	0.834	0.049	0.059	2.08	0.77
			2	0.836	0.048	0.060	2.83	0.99
			3	0.837	0.046	0.059	3.35	1.11
			4	0.838	0.049	0.061	3.73	1.19
			5	0.838	0.048	0.061	4.04	1.24
			6	0.839	0.050	0.063	4.28	1.28
	np-boot	800	1	0.798	0.038	0.038	1.99	0.95
			2	0.795	0.037	0.037	2.64	1.17
			3	0.793	0.038	0.039	3.08	1.27
			4	0.793	0.039	0.039	3.42	1.37
			5	0.793	0.038	0.039	3.66	1.38
			6	0.793	0.038	0.039	3.85	1.35
	BB		1	0.833	0.033	0.047	2.21	1.11
			2	0.838	0.035	0.051	3.01	1.44
			3	0.841	0.035	0.053	3.56	1.63
			4	0.843	0.034	0.055	3.97	1.75
			5	0.845	0.035	0.057	4.29	1.82
			6	0.848	0.036	0.060	4.55	1.86
0.95	np-boot	400	1	0.943	0.032	0.033	3.22	1.21
			2	0.944	0.027	0.028	4.41	1.56
			3	0.944	0.026	0.026	5.23	1.73
			4	0.944	0.025	0.026	5.84	1.82
			5	0.944	0.025	0.025	6.31	1.87
			6	0.943	0.025	0.026	6.65	1.87
	BB		1	0.942	0.028	0.029	3.17	1.17
			2	0.941	0.026	0.027	4.33	1.51
			3	0.942	0.026	0.027	5.12	1.70
			4	0.940	0.026	0.027	5.71	1.81
			5	0.940	0.026	0.028	6.17	1.89
			6	0.940	0.027	0.028	6.55	1.95

Table 4. Prediction Intervals for Model 2

Nominal Coverage	Method	T	k	Mean Coverage	S.D. of Coverage	R.M.S.E	Mean Length	S.D. of Length
0.95	np-boot	800	1	0.945	0.018	0.019	3.41	1.58
			2	0.945	0.019	0.020	4.69	2.13
			3	0.946	0.018	0.019	5.56	2.50
			4	0.945	0.019	0.020	6.20	2.73
			5	0.945	0.019	0.020	6.64	2.77
			6	0.944	0.018	0.019	6.98	2.76
	BB		1	0.943	0.018	0.019	3.38	1.70
			2	0.943	0.019	0.020	4.61	2.21
			3	0.944	0.018	0.019	5.45	2.50
			4	0.944	0.018	0.019	6.07	2.67
			5	0.944	0.018	0.019	6.56	2.78
			6	0.945	0.018	0.019	6.96	2.84
0.99	np-boot	400	1	0.985	0.016	0.016	5.07	2.04
			2	0.986	0.012	0.013	7.15	2.80
			3	0.986	0.011	0.011	8.44	3.12
			4	0.986	0.010	0.011	9.45	3.38
			5	0.985	0.011	0.011	10.32	3.57
			6	0.984	0.012	0.013	10.96	3.91
	BB		1	0.976	0.016	0.021	4.17	1.54
			2	0.975	0.015	0.021	5.69	1.98
			3	0.974	0.014	0.021	6.72	2.23
			4	0.974	0.014	0.022	7.50	2.38
			5	0.973	0.015	0.022	8.11	2.49
			6	0.972	0.015	0.023	8.61	2.56
	np-boot	800	1	0.986	0.009	0.010	5.36	2.79
			2	0.987	0.008	0.008	7.58	3.95
			3	0.986	0.008	0.009	9.01	4.66
			4	0.987	0.007	0.008	10.11	4.73
			5	0.987	0.007	0.008	11.03	5.17
			6	0.987	0.008	0.008	11.75	5.46
	BB		1	0.976	0.010	0.017	4.44	2.23
			2	0.976	0.010	0.017	6.06	2.90
			3	0.975	0.010	0.018	7.16	3.28
			4	0.975	0.010	0.018	7.98	3.51
			5	0.975	0.009	0.017	8.63	3.65
			6	0.976	0.010	0.017	9.14	3.73

5.1 Results

For the parametric case, the 80% and 99% BB prediction intervals are only accurate for a one-step-ahead prediction interval. This was also found by Baillie and Bollerslev (1992), in which the prediction interval overestimated the 80% prediction interval and underestimated the 99% prediction interval. However, the parametric bootstrap prediction interval is quite accurate for all cases. The mean coverage is also very good for both sample sizes of 400 and 800. It becomes more accurate as the sample size becomes larger.

In the non-parametric setting the bootstrap does very well, relative to an interval based on an incorrect parametric assumption, particularly for the 80% and 99% nominal coverages. At the 95% nominal coverage level the BB prediction interval is competitive with the non-parametric bootstrap prediction interval.

The lengths of the bootstrap prediction intervals are generally good, whereas, the BB's length can often be too wide for the 80% nominal coverage and too narrow for the 99% nominal coverage.

6 Empirical Illustration

We illustrate the improved accuracy of the bootstrap prediction interval relative to the analytical Gaussian prediction interval by modelling the Yen/$US as an AR(1)-GARCH(1,1) process. The GARCH(1,1) model for the Yen/$US has been used in previous studies, see West and Cho (1995). We generate the one-step-ahead distribution of future values using both procedures and find the non-parametric bootstrap to be significantly more accurate than the analytical Gaussian procedure.

6.1 Data

Our exchange rate is the Yen/$US. Which is the number of Yen to buy one US dollar. The data are Wednesday, New York noon bid rates, as published in the *The Federal Reserve Bulletin.* When Wednesday was a holiday we used Thursday data. We take the logarithmic difference of the exchange rate and then multiply by 100. That is, we model

$$e_t = 100 * ln\left(\frac{Yen/\$US \text{ in week t}}{Yen/\$US \text{ in week t-1}}\right)$$

and thus has the interpretation of percentage change in the level of the Yen/$US.

Figure 1 graphs the level of the Yen/$US and Figure 2 graphs its percentage change. Our data set covers the period from 7 March 1973 to 10 November 1999.

6.2 Empirical Coverage for One-Step-Ahead Prediction

We now compare the empirical performance of the non-parametric bootstrap against the analytical Gaussian. We take rolling sample sizes of 500 observations to estimate AR(1)-GARCH(1,1) models beginning on 28 March 1979. We do this 500 times, moving through time. The last observation used was on 20 May 1998.

Our first estimation using data from 28 March 1979 through to 26 October 1988 produced:

$$\begin{aligned} e_t &= \underset{(0.0649)}{-0.0805} + \underset{(0.0508)}{0.0939 e_{t-1}} + u_t \\ \omega_t^2 &= \underset{(0.4372)}{0.8128} + \underset{(0.0633)}{0.1223 u_{t-1}^2} + \underset{(0.2195)}{0.5103 \omega_{t-1}^2} \\ u_t &= \omega_t \varepsilon_t \end{aligned}$$

Robust standard errors, as described in Bollerslev and Wooldridge (1992), are given below the parameter estimates.

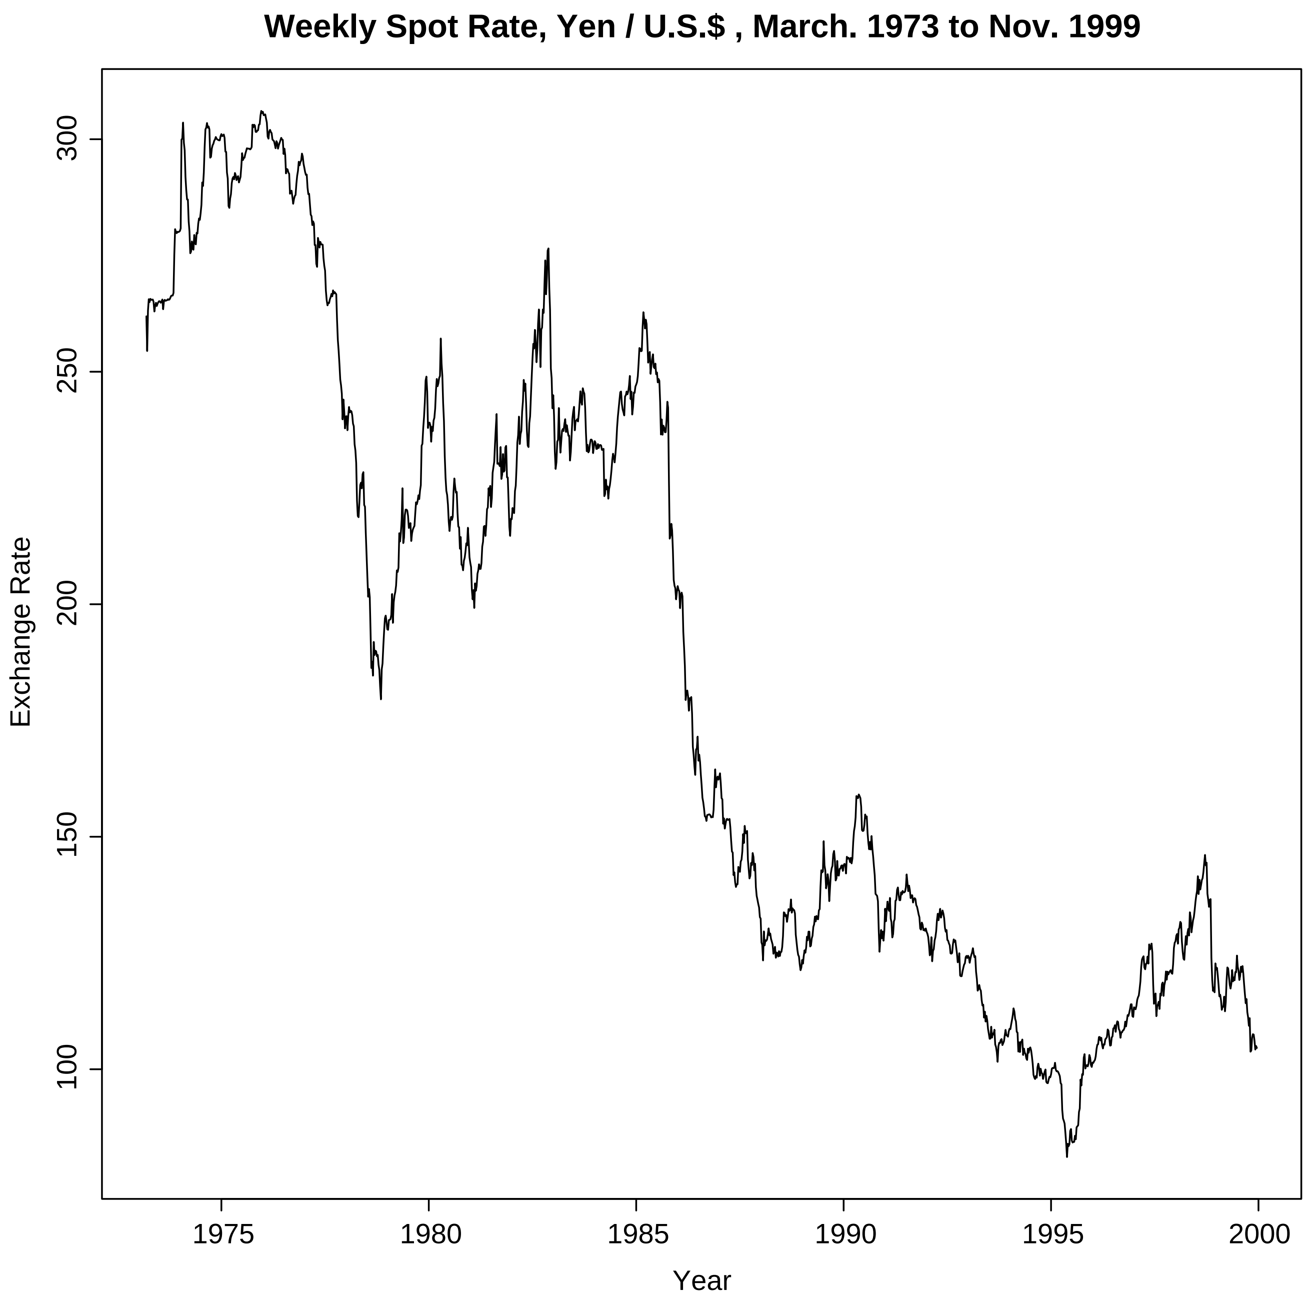
Weekly Spot Rate, Yen / U.S.$, March. 1973 to Nov. 1999
Exchange Rate
100
150
200
250
300
1975
1980
1985
1990
1995
2000
Year

Figure 1:

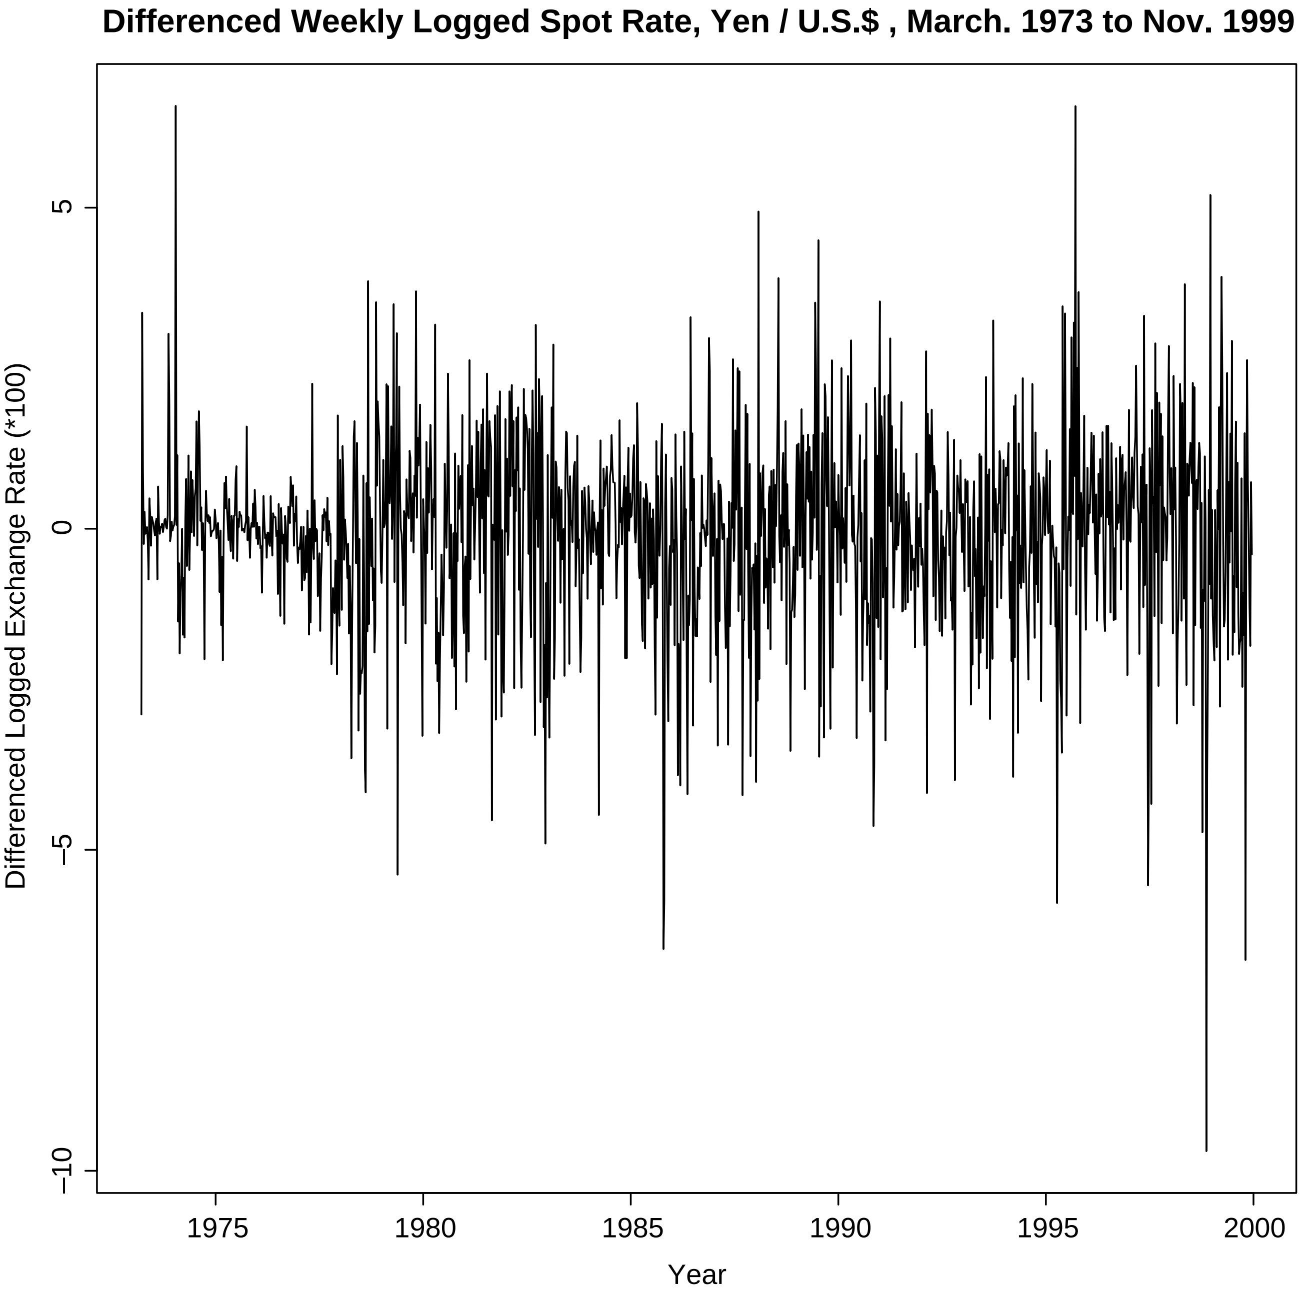

Differenced Weekly Logged Spot Rate, Yen / U.S.$, March. 1973 to Nov. 1999
Differenced Logged Exchange Rate (*100)
5
0
−5
−10
1975
1980
1985
1990
1995
2000
Year

Figure 2:

Our last estimation using data from 19 October 1988 through to 20 May 1998 produced:

$$\begin{array}{rcl} e_t & = & \underset{(0.0576)}{0.0208} + \underset{(0.0485)}{0.0203 e_{t-1}} + u_t \end{array}$$

$$\omega_t^2 = \underset{(0.1842)}{0.2679} + \underset{(0.0351)}{0.1031 u_{t-1}^2} + \underset{(0.0989)}{0.7777 \omega_{t-1}^2}$$

$$u_t = \omega_t \varepsilon_t$$

For each estimation for a sample size of 500 we calculate the one-step-ahead distribution of future values. This is done using the non-parametric bootstrap with $B = 2000$, and is also done using the analytical Gaussian solution.

We evaluate the empirical performance of these two procedures by examining where the actual exchange rate was relative to the prediction. The number of times out of 500 that the actual exchange rate was less than the predicted quantile gives us an empirical estimate of the empirical coverage of this predicted quantile. This is done for both the bootstrap and the Gaussian and is plotted in Figure 3, with the bootstrap being significantly more accurate.

Two-sided one-step-ahead prediction intervals are also constructed and their empirical performance measured as well. Table 5 displays the results with the superior performance of the bootstrap clearly evident.

Table 5. Empirical Coverage for Centered Two-Sided One-Step-Ahead Prediction Intervals

Desired Coverage	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Bootstrap Coverage	0.094	0.200	0.296	0.386	0.494	0.590	0.688	0.792	0.896	0.946
Gaussian Coverage	0.128	0.258	0.350	0.444	0.558	0.648	0.738	0.834	0.902	0.940

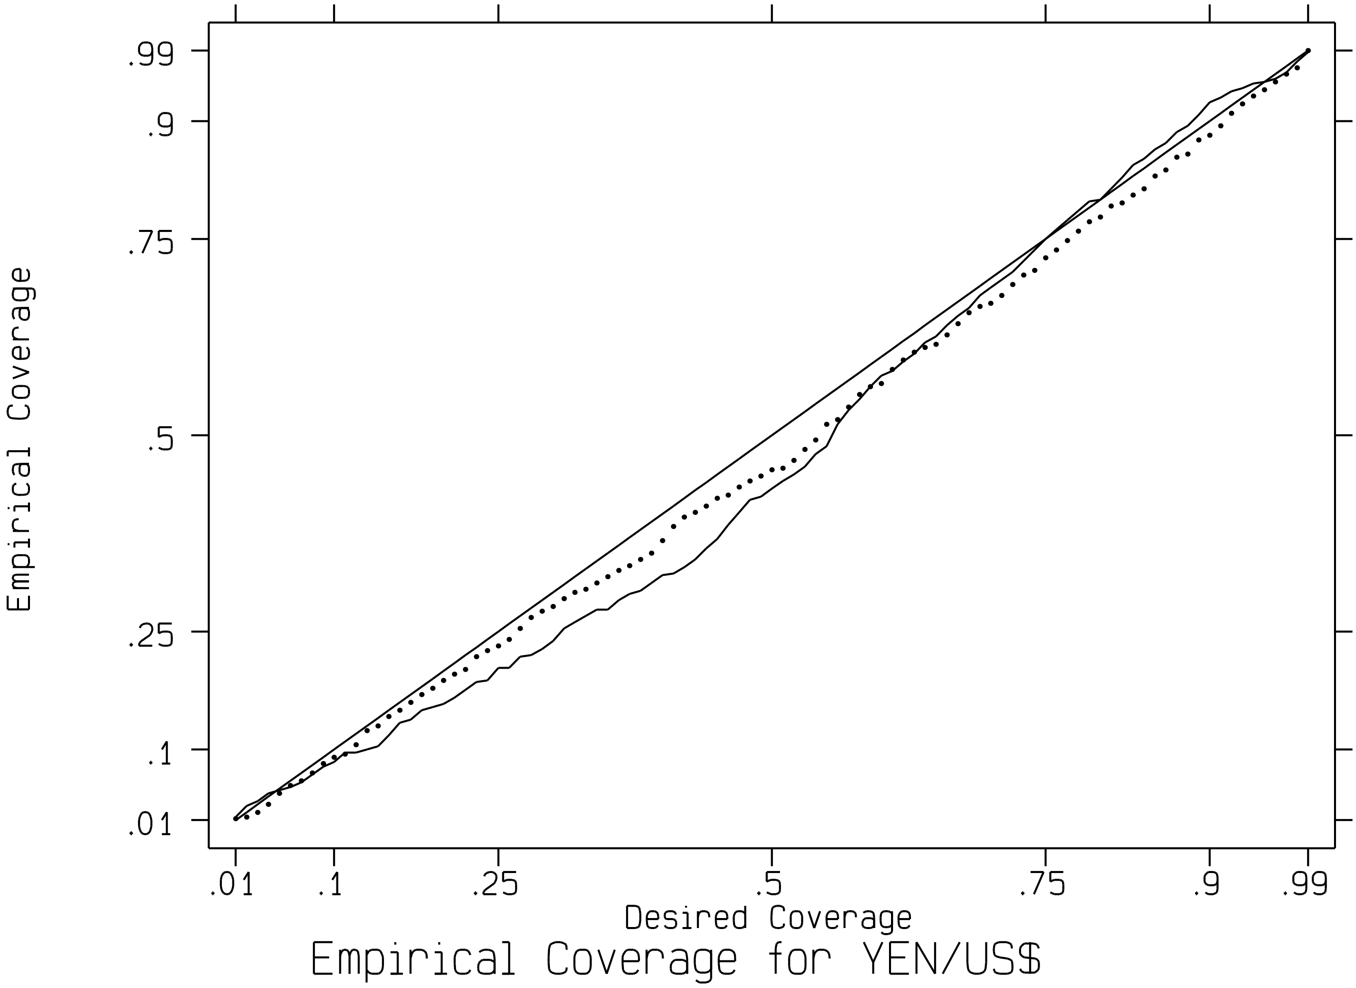

Figure 3: The bootstrap empirical coverage is the dotted curve and the Gaussian coverage is the smooth curve. The straight line is the 45 degree line.

7 Concluding Remarks

We have shown that the bootstrap prediction interval applied to ARCH processes is asymptotically justified, and performs well for the finite sample sizes often encounted in empirical finance. The procedure is applicable to the general class of ARCH models, takes account of parameter uncertainty, and one is not required to make a parametric assumption. Interesting future research would be on using these procedures in a multivariate setting, where again in most cases analytical closed form formulas for prediction intervals do not currently exist.

The accuracy of the AR(1)-GARCH(1,1) non-parametric bootstrap forecast distribution for the Yen/$US should lead to improvements in portfolio allocation decisions involving this asset.

REFERENCES

Baillie, R. T. and Bollerslev, T., (1992), "Prediction in Dynamic Models with Time-Dependent Conditional Variances," *Journal of Econometrics*, **52**, 91-113.

Beran, R., (1990), "Calibrating Prediction Regions," *Journal of the American Statistical Association*, **85**, 715-723.

Bickel, P. J. and Freedman, D. A., (1981), "Some Asymptotic Theory for the Bootstrap," *The Annals of Statistics*, **9**, 1196-1217.

Bollerslev, T. and Wooldridge, J. M., (1992), "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances," *Econometric Reviews*, **11(2)**, 143-172.

Diebold, F. X., (1988), Empirical Modeling of Exchange Rate Dynamics, (Springer, Verlag, New York, NY).

Engle, R. F., (1982), "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, **50**, 987-1007.

Engle, R. F., and Kraft, D. F., (1983), Multiperiod Forecast Error Variances of Inflation Estimated from ARCH Models, in: Applied Time Series Analysis of Economic Data (Bureau of the Census, Washington, DC).

Engle, R. F., Lilien, D.M., and Robins, R. P., (1987), "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model," *Econometrica*, **55**, 391-407.

Geweke, J., (1989), "Exact Predictive Densities for Linear Models with ARCH Disturbances," *Journal of Econometrics*, **40**, 63-86.

Granger, C. W. J., White, H., and Kamstra, M., (1989), "Interval Forecasting: An Analysis Based Upon ARCH-Quantile Estimators," *Journal of Econometrics*, **40**, 87-96.

Koenker, R., and Zhao, Q., (1996), "Conditional Quantile Estimation and Inference for ARCH Models," *Econometric Theory*, **12**, 793-813.

Lee, S., and Hansen, B. E., (1994), "Asymptotic Theory for the GARCH(1,1) Quasi-Maximum Likelihood Estimator," *Econometric Theory*, **10**, 29-52.

Lumsdaine, R. L., (1996), "Consistency and Asymptotic Normality of the Quasi-Maximum Likelihood Estimator in IGARCH(1,1) and Covariance Stationary GARCH(1,1) Models," *Econometrica*, **64**, 575-596.

Schachter, B., (1998), "An Irreverent Guide to Value at Risk," *Risks and Rewards*, **March**, 17-18.

Thombs, L. A. and Schucany, W. R., (1990), "Bootstrap Prediction Intervals for Autoregression," *Journal of the American Statistical Association*, **85**, 486-492.

West, K. D. and Cho, D., (1995), "The Predictive Ability of Several Models of Exchange Rate Volatility," *Journal of Econometrics*, **69**, 367-391.