Have Aluminium And Copper Futures Markets Become More Volatile Over Time?

Clinton Watkins and Michael McAleer

School of Economics and Commerce, University of Western Australia, E-Mail: Clinton.Watkins@rbnz.govt.nz

Keywords: *Volatility forecasting; GARCH; Rolling models; Futures contract; Industrial metals*

EXTENDED ABSTRACT

Within the industrial metals industry, there has been a great deal of interest surrounding trends in volatility over time. This paper uses a rolling AR(1)-GARCH(1,1) model to estimate and forecast the volatility processes for daily returns on the futures prices of two important non-ferrous metals, aluminium and copper.

The LME is the major international market for the main industrially-used non-ferrous metals, namely aluminium, aluminium alloy, copper, lead, nickel, tin, and zinc. Aluminium has the highest volume of spot and futures trade on the exchange, followed closely by copper. The two metals are also amongst the most important metals in an industrial sense.

Changes in the prices of aluminium and copper are often closely aligned with changes in global industrial production, but also reflect market specific events, compliment and substitute relationships between the physical metals in production, and financial market type influences. Brunetti and Gilbert [1995] characterise two sources of volatility in non-ferrous metals markets, those related to financial market considerations and those related to market fundamentals. Financial considerations include information effects, speculative pressures and hedging activity, usually giving rise to short-run volatility effects. Market fundamentals refer to the underlying availability, supply and demand of physical metal. Fundamentals are a source of long term volatility in metals markets, primarily due to the lags involved in demand and supply side changes.

In this paper, the rolling model is used to examine how the processes driving aluminium and copper returns volatility have evolved over a long sample. Settlement price data on 3-month futures contracts traded on the London Metals Exchange (LME) is used to calculate the daily returns series. The sample consists of daily data for 3-month futures settlement prices in US dollars for aluminium over the period 1 October 1982 to 15 July 2005, and for copper over the period 5 January 1976 to 15 July 2005. The model is estimated using a rolling window of 1000 observations, which iterates 4671 times for aluminium, and 6448 times for copper.

The models are used to examine when and in what manner the α and β coefficients, as estimated parameters of the volatility process, change over time. Estimates are presented graphically. The α estimates for both metals indicate that short-run persistence varies in magnitude through each sample. Similarly, the β estimates also vary markedly over time. The long run persistence of volatility, $\alpha+\beta$, is also non-constant over the sample. Moment conditions, specifically the second and fourth, are examined in order to evaluate the statistical properties of the empirical models. One-step-ahead forecasts are also generated and compared with a measure of the 'true' volatility, as defined by Pagan and Schwert [1990]. Several forecast evaluation criteria are also applied to the series of forecasts.

The results of the paper suggest that, while the volatility of returns does not appear to display an upward trend, relative to the 1980's there are periods over the following years where the process driving time-varying conditional volatility appears to have become more variable, and to some degree harder to model at some times using a simple GARCH specification. The variation over time seen in the volatility process as modelled by GARCH suggests that, while volatility in returns has not necessarily increased, volatility in metals markets is itself volatile when analysed over a long horizon. Of course, instability in the GARCH model may indicate that a more complex volatility model is required to better reflect the volatility process in the data.

1. INTRODUCTION

Within the metals industry, there has been a great deal of interest surrounding trends in volatility over time. These have been examined in a number of studies which look at variability in prices and market structure. Metals market participants suggest that volatility in non-ferrous metals prices has increased over time. Furthermore, this increase in volatility is thought to be associated with a change in market organization from pricing on a producer list basis, to exchange-based pricing. Although exchange-based pricing allows access to hedging products, such as exchange traded futures and options, some argue that the associated hedging costs and increased volatility has led to higher cost structures for metal produces than those connected with producer list pricing.

Slade [1991] provided empirical evidence to support the proposition that metals price volatility is higher during the 1980s relative to the 1970s, and that this increase in volatility was associated with the transition from producer list to exchange determined prices. However, Figuerola-Ferretti and Gilbert [2001] show that by extending the Slade's [1991] sample to include more recent data, there is no significant difference in the variability of exchange-based prices and producer list prices. Brunetti and Gilbert [1995, 1996] argue that while the volatility of non-ferrous metals prices is itself highly volatile, volatility does not trend stochastically over a sample covering 1972 to 1995. Periods of high volatility in metals markets are due to tighter metals balances, rather than speculative activities. Speculative and informational pressures are not precluded from influencing volatility. However, the effects are short lived, and fundamentals regarding the availability of metal generate persistence in volatility.

2. ROLLING VOLATILITY MODEL

This paper uses a rolling volatility model to examine how the processes driving aluminium and copper volatility have evolved over a long sample. Bollerslev's [1986] GARCH model is used, specifically AR(1)-GARCH(1,1). This is one of the most widely used time-varying financial volatility models in practice. In this model, the conditional mean of futures price returns is given by the AR(1) model:

$$r_t = \mu + \varphi r_{t-1} + \varepsilon_t, \quad |\varphi| < 1, \tag{1}$$

and the conditional variance of ε_t is:

$$\varepsilon_t = \eta_t \sqrt{h_t}\ , \tag{2}$$

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}\, , \tag{3}$$

where r_t denotes returns on the futures price from period t-1 to t; ε_t is the unconditional shock; η_t is a sequence of normally, independently and identically distributed random variables, with zero mean and unit variance; and h_t is the conditional variance of returns. For the GARCH process to exist, $\omega > 0$, $\alpha \geq 0$ and $\beta \geq 0$ are sufficient conditions for the conditional variance to be positive. The ARCH coefficient, α, measures short-run persistence in volatility, and the GARCH effect, β, measures the contribution to long-run persistence, namely α+β.

Several statistical properties have been established for the GARCH(p,q) process in order to define the moments of the unconditional shock (see for example, Ling and McAleer [2002]). The necessary and sufficient condition for the second moment to exist for the GARCH(1,1) model, guaranteeing that the process is strictly stationary and ergodic, is given by:

$$\alpha + \beta < 1. \tag{4}$$

If the conditional shocks, η_t, are iid random variables, the fourth moment of the unconditional shock will exist if the following condition is satisfied:

$$3\alpha^2 + 2\alpha\beta + \beta^2 < 1. \tag{5}$$

The model is estimated using a rolling window of 1000 observations, which iterates 4671 times for aluminium, and 6448 times for copper. Each model is estimated by maximum likelihood. Non-normality in the residuals is likely in the presence of extreme returns observations. Robust t-ratios based on the standard errors of Bollerslev and Wooldridge [1992], which are designed to be robust to non-normality, are used. Estimates from the rolling samples are treated as "data" in the descriptive discussion below.

An empirical measure of volatility is required against which to evaluate the rolling one-step-ahead forecasts generated by the model. The measure defined in Pagan and Schwert [1990] is used at the "true" or actual empirical volatility:

$$v_t = |r_t - \bar{r}| \tag{7}$$

where v_t refers to the 'true' volatility at time t, and $\bar{r}$ is defined as the mean return over the estimation window for the sample used. The 1-day ahead forecast error, u_{t+1}, is defined as:

$$u_{t+1} = \hat{h}_{t+1} - v_{t+1}. \quad (8)$$

A positive forecast error implies that the GARCH model has over-forecast the empirical volatility, while a negative forecast error means the empirical volatility has been under-forecast.

3. NON-FERROUS METALS DATA

The sample consists of daily data for 3-month futures settlement prices in US dollars for aluminium over the period 1 October 1982 to 15 July 2005, and for copper over the period 5 January 1976 to 15 July 2005. Each returns series is calculated as:

$$r_{t-1,t} = \left(f_t - f_{t-1}\right) / f_{t-1} , \quad (6)$$

where $r_{t-1,t}$ is the return over period t-1 to t, and f_t is the futures price at time t. This gives a sample of 5671 returns on aluminium futures and 7448 returns on copper futures, noting that 79 zero returns observations were eliminated from the aluminium sample.

Plots of the price and returns series for aluminium and copper are presented in Figures 1 and 7. While copper and aluminium sometimes share periods of clustered volatility at similar times, each market also contains periods of volatility not occurring in the other. Industrial metals markets, while being affected by macroeconomic shocks, are also strongly influenced by financial and physical market-specific events. The extent to which these permeate between metals markets depends on a number of factors, including the complimentary and substitute relationships between the particular metals.

4 PARAMETER ESTIMATES

Plots of the rolling α coefficient estimates and their t-ratios are provided in Figures 2 and 8 for aluminium and copper, respectively. Rolling β estimates and their t-ratios are shown in Figures 3 and 9. Dates on the x-axis of these figures indicate the last trading day contained in the estimation window.

The α estimates for both metals indicate that short-run persistence varies in magnitude through each sample. A number of the features apparent in the rolling α estimates coincide with extreme observations that occur as a result of shocks to the returns series. There appears to be an upward bias in the α estimates when large extreme observations are within the estimation window. A high α implies that the GARCH process will allow the level of estimated conditional variance to increase quickly in response to periods of high volatility in returns. The majority of the rolling β estimates for both metals series are greater than 0.8 and less than 1, as expected. However, β estimates also display varying characteristics over time.

Periods in which the α estimates appear biased upward often coincide with periods of apparent downward bias in β. A low β estimate allows for the rapid decay of volatility, as might be expected following an extreme observation. However, the presumption under the GARCH model is that observations from high volatility periods and low volatility periods follow the same parametric volatility process. If extreme observations generate a lower degree of persistence, and that cannot be captured in the GARCH(1,1) parameterisation, it would be expected that biases will occur in the estimates. In this case, the β estimate may be biased downward following an extreme observation.

5. ROBUST ROLLING T-RATIOS

For the results prior to December 1998, all aluminium α estimates are significant at the 5 percent level. Toward the end of the sample, the t-ratios generally sit just above the 5 percent significance level, however while the majority of estimates appear significant, the null is not rejected for some estimates. Copper t-ratios for the estimated α are above 2 for much of the sample. However there are some periods where the t-ratio varies around the 5 percent critical value. There is an extended period over which the null is not rejected, from March 2003 to May 2004. Interestingly, a level shift in the alpha estimate around September 1987 coincides with noticeably higher significance levels.

Typically, the β t-ratios for both series are large, particularly so when the coefficient estimates are between 0.8 and 1. The copper β estimates are significant over the entire sample. Between July 2000 and December 2001, when the aluminium β estimate becomes extremely variable, its t-ratio becomes low and most estimates appear not to be significant.

6. ROLLING MOMENT CONDITIONS

The percentage of the rolling estimation windows for which the second and fourth moment conditions are satisfied is shown in Table 1. Both the second and fourth moment conditions are satisfied more frequently in models for copper than for the aluminium returns data.

Moment	Aluminium	Copper
Second	85.38	89.41
Fourth	83.41	86.74

Table 1. Moment conditions satisfied (in %)

Plots of the moments are given in Figure 4 for aluminium and Figure 10 for copper. The aluminium moments exceed one for several periods over the sample, including for over two years of trading days from November 1987 to December 1990. Similarly for copper, there is an extended period where the moment conditions are violated, between October 1987 and October 1990.

Clearly, market conditions surrounding the fallout of the equity market melt-down at that time have influenced volatility in both aluminium and copper prices. It is interesting to note that, at the time when market participants would be expected to be most concerned with accurate estimates of returns volatility, inferences the simple rolling GARCH model cannot be relied upon. For copper, there are also instances when the forth moment condition is violated, between December 1991 to June 1993.

Despite these moment condition violations, the long-run persistence in volatility is generally closer to unity for copper than it is for aluminium. Long-run persistence for aluminium is low in the early 1990s and from September 2000 to December 2001, while for copper it is low in the mid-1980s.

7. VOLATILITY FORECASTS

Forecasts of volatility generated by the models, and the actual of 'true' volatility, are shown in Figures 5 and 11. Table 2 compares the forecasts using mean error (ME), mean absolute error (MAE), root mean squared error (RMSE), smoothed mean absolute percentage error (SMAPE), smoothed weighted median absolute percentage error (SMedWAPE), and smoothed weighted mean absolute percentage error (SWMAPE). Forecast errors are graphed in Figures 6 and 12.

The criteria suggest a broadly similar forecast performance in each market. Median errors are always smaller than the comparable mean errors, and errors are slightly greater for copper volatility than for aluminium. When RMSE calculated separately for positive and negative forecast errors, RMSE(-) is substantially larger than RMSE(+) for both markets. Volatility forecast errors from the GARCH(1,1) model are negatively biased. The weighted forecast performance measures are considerably lower than their non-weighted counterparts. This suggests that a large proportion of forecast errors occur when the actual volatility of the forecast period is lower than the average volatility in the entire sample.

Table 2: Forecast evaluation criteria

Criteria	Aluminium	Copper
ME	0.00315	0.00365
MAE	0.00680	0.00785
MedAE	0.00590	0.00683
RMSE	0.00898	0.01028
RMSE(-)	0.01131	0.01262
RMSE(+)	0.00804	0.00931
RMedSE	0.00590	0.00683
SMAPE	74.13	74.61
SMedAPE	63.36	62.85
SMWAPE	51.89	51.11
SMedWAPE	32.97	33.30
R2	0.15	0.12
R2E	0.70	0.74
Forecasts Under	25.26	26.05
Forecasts Over	74.74	73.95

R^2 is obtained by regressing ex-post volatility on forecast volatility, and being relatively low for both metals it suggests a poor overall predictive performance. Similarly, R^2_e obtained by regressing the forecast errors on the ex-post volatility show that the actual volatilities have a high degree of explanatory power over forecast errors. The final two criteria show that the GRARCH model tends to over-predict volatility for both metals.

Looking at figures 5 and 11, forecasts from the GARCH models capture the major features of the actual volatility in aluminium and copper returns over the sample.

A clear period of heightened volatility is evident between October 1987 and early 1990, coinciding with the October 1987 equity market melt-down, and flow on volatility in global financial markets and the major macroeconomies. However, during this period, the forecast errors suggest that the model tends to over-predict the persistence of volatility in both markets.

While for aluminium, the October 1987 crash corresponds with the highest volatility forecasts, the largest copper forecasts occur in early 1996. The LME copper market was systematically manipulated by a trader in the Sumitomo Corporation of Japan during the early and mid 1990s [see Gilbert 1996]. Substantial volatility can be observed in returns around early to mid 1996, when conditions in the copper market made Sumitomo's position untenable. At that time, hedge funds saw an opportunity to attack the inflated copper price, which fell by USD700 per tonne over a four-week period.

In general, large negative forecast errors tend to be associated with the model under predicting when a substantial shock to returns hits the market, while positive forecast errors tend to be associated with the model over predicting the persistence of these shocks. In periods of relatively low volatility, the model tends to under-predict actual volatility. Extreme observations associated with a period of relatively higher volatility appear to be the source of adverse effects on the predictive ability of the GARCH model in subsequent periods.

8 EMPIRICAL SIGNIFICANCE

The α and β estimates for both aluminium and copper indicate that short- and long-run volatility persistence can vary over a relatively wide range as the models move through time. At times the estimates exhibit a high degree of stability. However shocks in returns can move the process dramatically. At some points, the α estimates are close to zero, suggesting no short-term persistence. In general, the β contribution to long-run persistence becomes relatively more important than short-run persistence, towards the end of each sample.

Tail events have a clear effect on the likelihood function, and thus on estimation of volatility models using maximum likelihood methods. Biased GARCH model estimates occur because the (quasi) maximum likelihood estimator seeks a model in which the estimated conditional variance may increase quickly, while also allowing for rapid volatility decay. The means by which a GARCH(1,1) model achieves this is with a high α parameter which allows a quick increase in conditional variance and a low β parameter which provides for rapid decay.

The rolling GARCH models produce a number of results that have interesting economic implications, particularly with regard to the perception within the non-ferrous metals industry that metals prices and returns became more volatile over time. Inspection of the aluminium and copper returns plots does not support this proposition. On the contrary, it would appear that returns are less volatile in both the aluminium and copper markets in the later portion of the sample, particularly in terms of the prevalence of tail or extreme events. Moreover, Brunetti and Gilbert [1996] contend that volatility itself has become more volatile. Several of the results presented in this paper support their conclusion. While the volatility of returns does not appear to display an upward trend, relative to the 1980’s there are periods over the following years where the process driving time-varying conditional volatility appears to have become more variable, and to some degree harder to model at some times using a simple GARCH specification. The variation over time seen in the volatility process as modelled by GARCH suggests that, while volatility in returns has not necessarily increased, volatility in metals markets is itself volatile when analysed over a long horizon. Of course, instability in the GARCH model may indicate that a more complex volatility model is required to better reflect the volatility process in the data.

ACKNOWLEDGEMENTS

The first author would like to acknowledge the Australian Research Council (ARC) for an Ad Hoc PhD Scholarship and the University of Western Australia for a C.A. Vargovic Memorial Fund Award. The second author wishes to acknowledge the financial support of the ARC.

REFERENCES

Bollerslev, T. (1986). Generalised autoregressive conditional heteroskedasticity, *Journal of Econometrics*, 31, 307-327.

Bollerslev, T. & Wooldridge J.M. (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances, *Econometric Reviews*, 11, 143-172.

Brunetti, C. & Gilbert C.L. (1995). Metals price volatility, 1972-95, *Resources Policy*, 21, 237-254.

Brunetti, C. & Gilbert C.L. (1996). Are metals' prices becoming more volatile?, *Proceedings of the Mineral Economics and Management Society*, 58-72.

Figuerola-Ferretti, I. & Gilbert, C.L. (2001). Price variability and marketing method in non-ferrous metals: Slade’s analysis revisited, *Resources Policy*, 27, 169-177.

Gilbert, C.L. (1996). Manipulation of metals futures: lessons learned from Sumitomo, *Discussion Paper No. 1537*, Centre for Economic Policy Research.

Ling, S. & McAleer, M. (2002). Necessary and Sufficient Moment Conditions for the GARCH(r,s) and Asymmetric Power GARCH(r,s) Models, *Econometric Theory*, 18, 722-729.

Pagan, A.R. & Schwert, G.W. (1990). Alternative models for conditional stock volatility, *Journal of Econometrics*, 45, 267-290.

Slade, M.E. (1991). Market structure, marketing method, and price instability, *Quarterly Journal of Economics*,106, 1309-1339.

Figure 1: Aluminium prices and returns

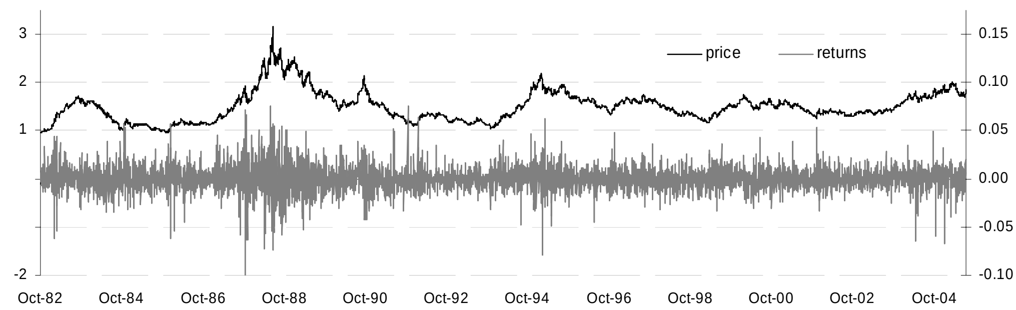

Figure 2: Aluminium alpha estimates and t-ratios

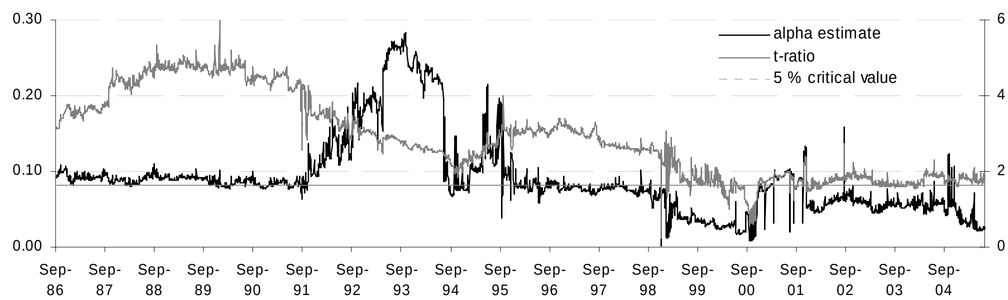

Figure 3: Aluminium beta estimates and t-ratios

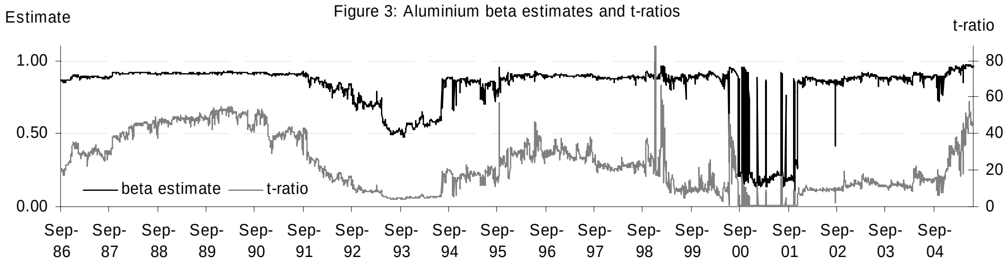

Figure 4: Aluminium second and fourth moments

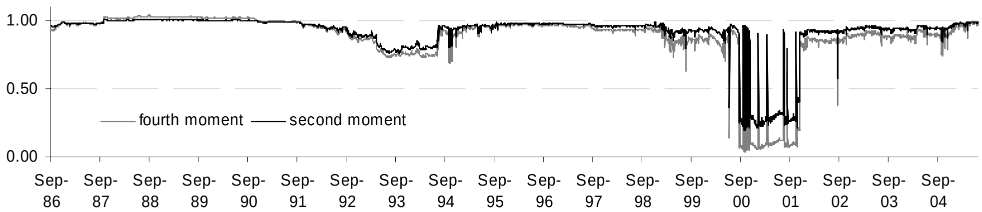

Figure 5: Aluminium forecast and actual volatility

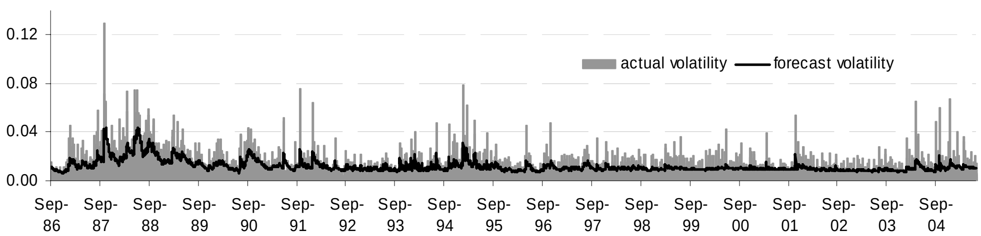

Figure 6: Aluminium forecast error

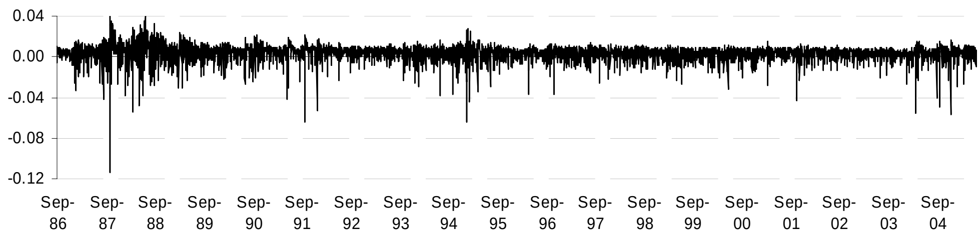

Figure 7: Copper prices and returns

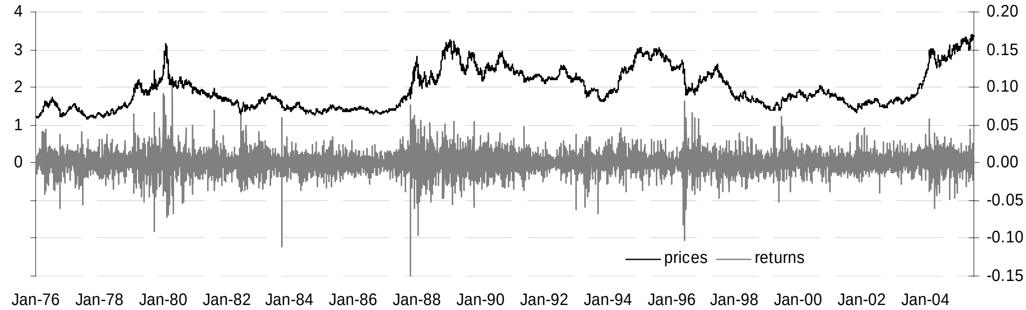

Figure 8: Copper alpha estimates and t-ratios

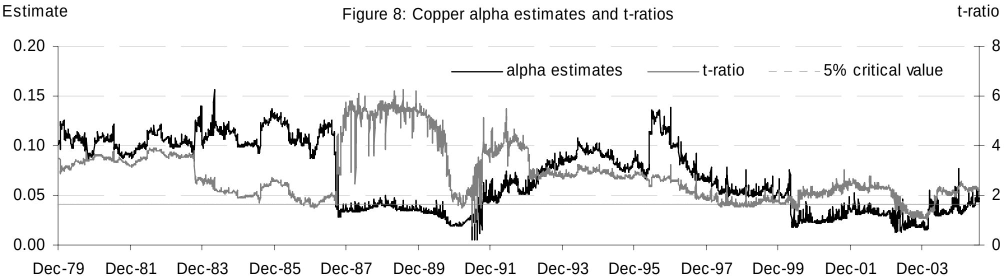

Figure 9: Copper beta estimates and t-ratios

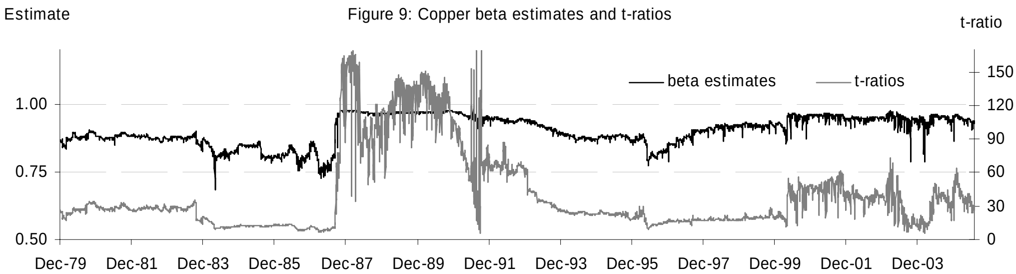

Figure 10: Second and fourth moment

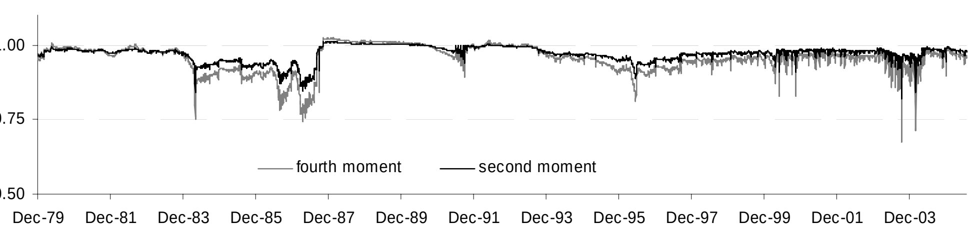

Figure 11: Copper forecast and actual volatility

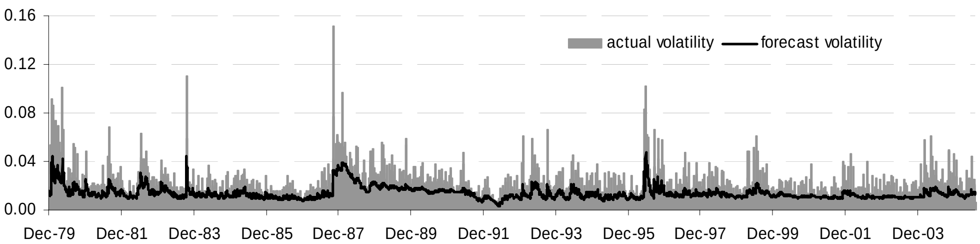

Figure 12: Copper forecast error