Correlation in Volatility Among Related Commodity Markets

Clinton Watkins and **Michael McAleer**

Department of Economics, University of Western Australia (cwatkins@iinet.net.au)

Abstract: Related commodity markets have two characteristics: (i) they may follow similar volatility processes; and (ii) such markets are frequently represented by a market aggregate or index. Indices are used to represent the performance and time series properties of a group of markets. An important issue regarding the time series properties of an index is how it reflects the time series properties of its components, particularly with regard to volatility. In this paper, correlation matrices are derived from rolling AR(1)-GARCH(1,1) model estimates to examine the second and fourth moment properties of ARMA processes with GARCH errors, and are also compared with the properties of the individual returns series. The correlations between the volatility of returns on several 3-month non-ferrous metals futures contracts traded on the London Metal Exchange are examined for aluminium, copper, nickel, lead, tin and zinc. Relationships between the volatility of individual metals returns and returns on the London Metal Exchange Base Metal Index are also examined.

Keywords: GARCH, futures, volatility, moments, cross-sectional aggregation, indexes, rolling windows.

1.   INTRODUCTION

Market aggregates (or indices) based on equities, commodities, or other classes of assets are used frequently in finance. Empirical modelling in this paper focuses on the recently launched London Metal Exchange Base Metals Index (LMEX) of the London Metal Exchange (LME), an index of spot and futures prices for the six primary metals traded on the LME. The LMEX is expected to perform an informational role for participants in the LME spot, futures and options markets. Moreover, the index is investible through futures and traded options contracts, based on the level of the index, that are available on the LME. In contrast to the LME metals futures and options, those for the index are familiar to financial market participants in that there is no element of physical delivery, and the contracts are cash settled. An index futures contract provides a convenient vehicle for investors to gain exposure to industrially-used non-ferrous metals markets without having to participate in one of the existing physical, futures or options markets at the LME. Such investors would be interested in the risk of the index relative to the risk in the primary non-ferrous metals markets, that is, the risk relationships between the markets on the LME and the index.

The volatility of the LMEX is compared with that of a subset of its components by analysing the correlations between rolling generalised autoregressive conditional heteroskedasticity (GARCH) processes for each individual series and the index. Estimated parameters, t-ratios, and moment conditions are generated using univariate rolling GARCH models. Correlation matrices are generated for the α and β estimates, their t-ratios, and the second and fourth moment conditions. The estimates, t-ratios and moment conditions of the

rolling GARCH model are treated as 'data' in the sense that inferences are drawn regarding the relationships between the index and its components, and between the components themselves, by examining the correlations between the series of estimates, t-ratios, and moment conditions.

2.   TIME-VARYING VOLATILITY MODEL

Bollerslev's (1986) GARCH model is used in the empirical analysis. The GARCH(1,1) specification is the most widely used model in the financial volatility literature, and is considered to represent adequately the observed symmetric intertemporal dependencies in daily returns of many financial time series. The conditional mean of futures price returns is given by the stationary AR(1) model:

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

and the conditional variance of ε_t is given by:

$$\varepsilon_t = \eta_t \sqrt{h_t}, \qquad (2)$$

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

where r_t denotes the returns on the futures price from period t-1 to t; ε_t is the unconditional shock; η_t is a sequence of 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 $\alpha + \beta$.

Several statistical properties have been established for the GARCH(p,q) process in order to define the moments of the unconditional shock. Ling and McAleer (2002a) derived a necessary and sufficient condition for the existence of moments of a family of GARCH processes, which includes the GARCH(1,1) model. Furthermore, Ling and McAleer (2002b) established the moment conditions for GARCH(p,q), and related the moment conditions to the statistical properties of the models, namely consistency and asymptotic normality. The implications of the non-existence of moment conditions, such as possible inconsistency of the parameter estimates and invalid inferences, have typically been ignored in the empirical literature on modelling volatility using GARCH-type processes.

The necessary and sufficient condition for the second moment of the GARCH(1,1) model to exist, which is equivalent to the GARCH(1,1) process being strictly stationary and ergodic, is given by:

$$\alpha + \beta < 1. \qquad (4)$$

If the standardised (or conditional) shocks, η_t, are a series of normally, independently and identically distributed random variables, the fourth moment of the unconditional shock will exist if and only if the following condition is satisfied:

$$3\alpha^2 + 2\alpha\beta + \beta^2 < 1. \qquad (5)$$

Prior to modelling the volatility of the series using AR(1)-GARCH(1,1), a specification search was conducted for both the mean and the variance. The GARCH(1,1) model was chosen to represent the variance process of the series. A higher-order GARCH process or a fractionally integrated GARCH (FIGARCH) model may also represent the volatility process for metals returns (Teyssiere, Brunetti and Gilbert, 1997). The GARCH(1,1) model performs adequately and the trade-off for parsimony, tractability and moment conditions are reasons for preferring the GARCH model. In order to conduct a correlation analysis between estimates of the model for various markets, using an adequate and parsimonious model permits more meaningful interpretations of the rolling correlations. Estimates of the coefficients of the conditional variance were not sensitive to changes in the specification of the AR(1) conditional mean equation.

A procedure was programmed in EViews 3.1 to estimate the AR(1)-GARCH(1,1) model using a rolling sample window of 1000 observations, which is approximately 4 years of trading days, over the entire data set. Recent research on optimal window sizes in rolling GARCH models suggests using 3 or 4 years of daily data to estimate the model (Yew, McAleer and Ling, 2001), as windows of this size are shown to produce stable estimates and moments. The rolling window procedure begins with the first 1000 observations being used to estimate the model.

Then the estimation interval is moved one-day into the future by deleting the first trading day and adding an extra day at the end of the sample window, and the parameters of the model are re-estimated. Each model is estimated by the maximum likelihood method, with the Marquardt algorithm used to maximise the likelihood function numerically. In the absence of normality of η_t, Quasi-Maximum Likelihood Estimators are obtained. This procedure is repeated 902 times.

In order to examine the structure of relationships between the estimated returns volatility processes for the LME metals and LMEX, correlations between the rolling estimates of the different metals models are analysed. These correlations are generated for the matrices described below. Correlations of the estimates are used to gain insights into the relationships between the estimated GARCH volatility processes, the short and long run persistence of volatility in related non-ferrous metals markets, and the closeness of LMEX to the volatility properties of its underlying assets.

Correlation matrices of the following form are created for each set of rolling α and β estimates, their robust t-ratios, and the second and fourth moment conditions:

$$P = \begin{bmatrix} 1 & \cdots & \rho_{1,j} & \cdots & \rho_{1,LMEX} \\ \vdots & 1 & \vdots & & \vdots \\ \rho_{i,1} & \cdots & 1 & \cdots & \rho_{i,LMEX} \\ \vdots & & \vdots & 1 & \vdots \\ \rho_{LMEX,1} & \cdots & \rho_{LMEX,j} & \cdots & 1 \end{bmatrix} \qquad (6)$$

where P denotes the correlation matrix, and ρ is the correlation between the estimates, robust t-ratios, moments of metals i and j, and metal I and LMEX.

3. THE DATA

The LMEX, which was launched in April 2000, is a linear combination of 1-, 2- and 3-month futures prices for six non-ferrous metals traded on the exchange, namely aluminium, copper, nickel, lead, tin and zinc. However, the LMEX data used in this paper were constructed retrospectively by the LME, which provides a sufficiently large sample for the analysis. Construction of the index is given as:

$$I = c_0 \sum_{m=1}^{6} \left[w_m \times \sum_{j=1}^{3} p_{j,m} \right] \qquad (7)$$

where I is LMEX, c_0 a constant chosen to normalise the index to 1000 on 4 January 1999, m represents the 6 metals, w_m is the weight for each metal, and $p_{j,m}$ is the futures price for metal m with maturity j.

Weights for each metal in the index are based on the proportions of global production and LME trading

volumes (liquidity) of each metal over the previous five years. Index weights are revised annually, with the new weights applying on the first trading day of July each year. Aluminium and copper are the most important metals in the index, weighted at 0.4 and 0.34, respectively. The weights of all the components remain relatively unchanged over the sample period. Daily data were obtained from the LME for 3-month futures prices for six non-ferrous metals, namely aluminium, copper, nickel, lead, tin and zinc, and for the LMEX. The 1-month and 2-month futures price series used in conjunction with the 3-month futures price to create the index were not available. Data cover the period 1 July 1992 to 10 January 2000, to provide 1902 trading days.

The logarithmic returns were calculated as follows:

$$r_{i,t} = \ln p_{i,t} - \ln p_{i,t-1}, \qquad (8)$$

where $r_{i,t}$ is the return for metal i from t-1 to t, and $p_{i,t}$ is the futures price, or index value series for metal i. Calculating the returns provides a sample of 1901 observations for the empirical analysis.

Typical attributes of commodity time series are apparent in the plots of the prices and returns. For the prices, these include periods of upward and downward trend, turning points, and structural breaks. The index follows a pattern loosely resembling the two major constituent parts, aluminium and copper, which are somewhat similar in terms of trends and turning points. For aluminium, copper and nickel, the index coincides with periods of upward and downward trends in the components. Lead and tin price series follow a similar pattern, while the zinc futures price follows a different trend. Within a long period of downward trend in copper prices, a steep decline in copper prices occurs between May and June 1999 as a result of the collapse of manipulation activities in the copper market by Sumitomo Corporation of Japan.

All the component series exhibit a significant amount of volatility and volatility clustering. Relative to its constituents, the LMEX returns series contain no extreme observations, so that no large returns shocks are present in the index. However, to some extent, component outliers influence the index. For example, two extremes in the copper market (observations 997 and 1000) appear to affect the index. Copper dominates the index during this time, due primarily to there being little activity in the other markets. While volatility clustering can be seen in the aggregate series, the degree is less than is evident in the components.

4.   CORRELATION OF THE ESTIMATES

4.1   Rolling α Estimates
Table 1 gives the correlation matrix for the seven sets of rolling α estimates. Correlations between different elements of the correlation matrix vary between a low positive correlation of 0.0594 and a high positive correlation of 0.9259, and a low negative correlation of –0.0603 to a moderate negative correlation of –0.4637. The α estimates for three metals returns series show a high and positive correlation with the α estimates for the index returns, namely copper, nickel and tin. Moreover, the α estimates for these metals are highly correlated with each other. Aluminium is the most important metal in the index by weight. However, aluminium is only moderately correlated with the index. It is also moderately correlated with every other component metal. The correlation is negative between aluminium and lead, while for the other metals it is positive. Both the lead and zinc α estimates show a small and negative correlation with those of the index, while being moderately negatively correlated with each other. Lead has a low positive correlation with metals that are highly correlated with the index, namely copper, nickel and tin, while zinc shows a low negative correlation with the same metals. Zinc is negatively correlated with all metals except for aluminium.

The correlation matrix for the α estimates shows that the short run volatility effects are highly correlated between some markets but not others. Four of the six metals have highly correlated short run volatility effects with the index, while two have very small negative correlations. The two largest constituents of the index, namely aluminium and copper, have substantially different correlations between their α estimates, and that of the index.

4.2   Rolling α Estimate Robust t-Ratios
As the residuals are suspected not to be conditionally normally distributed, quasi-maximum likelihood (QML) covariances and standard errors using the methods described by Bollerslev and Wooldridge (1992) are used. Even when the residuals are not conditionally normal, the α parameter estimates are consistent, provided the mean and variance functions are correctly specified. Under these circumstances, the estimate of the covariance matrix is consistent using QML covariances, so that the t-ratios and standard errors will be valid.

In Table 2, the correlation matrix for the t-ratios of the seven sets of rolling α estimates is presented. The correlations between the rolling robust t-statistics of each series of α estimates reveal several interesting relationships. Movements in the t-ratios of the α estimates are positively correlated among the individual metals, and between each metal and the index. In general, these positive correlations are also high in magnitude. However, the exception to this general observation is for any correlations involving lead, which are always low to moderate. The high positive relationship between changes in the robust t-ratios for the α estimates in most models indicates that changes in t-ratios and the importance

of the short run effects in volatility, are closely related between the models in the rolling windows.

The correlation matrix indicates that there is generally a high correlation between the short run volatility effects, in that the significance and change in significance of the α estimates are highly correlated. However, this is in contrast to the correlation between the rolling α estimates themselves, which presented a less consistent set of relationships between the various individual metals and the index. In terms of the t-ratios, shocks have similar short run effects on each market, and this translates through to the index as the aggregation of similar component effects means the index has comparable properties. The exception is lead, where the change in significance of the short run effect of shocks appears to be different.

4.3   Rolling β Estimates

The correlation matrix for the rolling β estimates is provided in Table 3. Correlations between the rolling β estimates range from –0.3920 to 0.8977, and highlight the substantially different behaviour of long run persistence in various markets. Overall, the level of correlation between the β estimates appears to be lower and less homogeneous than the correlation between short run persistence. Only the copper β estimate has a high positive correlation with that of the index. Aluminium, tin and zinc β estimates show a moderate and positive correlation with the aggregate series. There is a low and negative correlation between the index and both nickel and lead. Clearly, there are disparate GARCH effects among the components when compared with the index.

Generally, those components with a moderate or high positive correlation with the index are themselves moderately correlated. This is the case for aluminium and copper, aluminium and zinc, copper and tin, and copper and zinc. Exceptions include aluminium and tin, and tin and zinc, both of which have a low positive correlation. The two major components of the index themselves are only moderately positively correlated. Metals with a moderate or high positive correlation with the index have a low positive, low negative or moderate negative correlation with those metals which have a low negative correlation with the index. Aluminium and tin, and zinc and tin, display low positive correlations. While moderately correlated with the index and copper, tin is only slightly correlated with the other major component, aluminium. Copper and nickel, copper and lead, and tin and lead, have low negative correlations between their rolling β estimates. Zinc has a moderate positive correlation between its estimates and those of both nickel and lead. The only component metals to have a high positive correlation with each other are nickel and lead at 0.8771, which are those metals that displayed both a low and negative correlation with the index. Lead and nickel also exhibit low correlations with

the major components of the index, positive in the case of aluminium and negative for copper.

The correlation matrix for the rolling β estimates highlights a number of interesting relationships in the contribution of β to the long run volatility persistence in metals returns and returns on a metals index. Obviously, the autoregressive (or declining memory) effects of volatility shocks on the different markets can be systematically dissimilar for the various metals, possibly relating to different underlying fundamentals in these markets, different relationships with demand generated by industrial production, stocks and supply-side factors, and different supplementarity and complementarity relationships among the metals. The GARCH effect in the index is most closely related to the copper market, and somewhat less related to the aluminium, zinc and tin markets. In this regard, the lead and nickel markets bear little relation to the index.

4.4   Rolling β Estimate Robust t-Ratios

Table 4 provides the correlation matrix for the β estimate robust t-ratios for each model. The correlation between the t-ratios of the index and of its components is negative for copper, tin and zinc, and positive for aluminium nickel and lead. Lead shows the largest correlation in absolute magnitude, followed by aluminium, both of which show a moderate correlation with the index at 0.3920 and 0.3268, respectively. In absolute magnitude, there is a low level of correlation between the index and each of copper, nickel, tin and zinc. None of the models for the components produces a t-ratio that has a high correlation with the index, either positive or negative. While the copper β estimates are highly positively correlated with those of the index, the correlation between their respective t-ratios is negative and close to zero. Similarly, the β estimate for tin is moderately positively correlated with that of the index, but there is a low negative correlation between the respective t-ratios.

Most component metal β estimate t-ratios exhibit low (absolute) correlations, and of these nine, seven are negative. Those with low negative correlations are aluminium and tin, aluminium and zinc, copper and nickel, nickel and zinc, lead and tin, lead and zinc, and tin and zinc. Copper and zinc, and nickel and tin, have low positive correlations with each other. Zinc β estimate t-ratios have a low correlation with the other six series, and these correlations are negative for all the series except copper. Similarly, low correlations exist between the tin t-ratios and those of the other metals, and all but the correlations with copper and nickel are negative. Lead and aluminium t-ratios are the most highly correlated at 0.7049, followed closely by lead and nickel. Furthermore, nickel t-ratios are moderately positively correlated with those of aluminium and lead. Interestingly, the copper t-ratios have a moderate and negative correlation with the other major component, aluminium, and also with lead.

Comparing the β estimate correlations with the β t-ratio correlations reveals that relationships between the estimates of specific metals are not necessarily present between their respective t-ratios. Aluminium and copper t-ratios are negatively correlated, while their β estimates are positively correlated. Although the correlation between the copper and LMEX β estimates is high, the correlation between their t-ratios is close to zero. Aluminium, tin and zinc β estimates are moderately positively correlated with the index, but their t-ratios show a substantially lower correlation, which is negative for tin and zinc. However, nickel and lead have the highest correlation between the β estimates, and the second highest correlation between their t-ratios. Clearly, the pattern of correlations between the β estimates and their t-ratios is dissimilar in many cases.

4.5   Second Moment Condition
Second moment conditions were satisfied for all metals and the index. Table 5 contains the correlation matrix between the rolling second moments for seven variables. The correlation matrix for the second moments indicates the relationship between long run persistence for the seven series.

Aluminium and copper both have high correlations in volatility persistence with the index. While both are highly correlated with the index, they only show a moderate positive correlation with each other. Comparisons can be made with the α and β estimate correlations shown in Tables 1 and 3. Correlations in the second moments of the index and aluminium are much higher than for either the α or β estimates, while relationships between the index and copper, and between copper and aluminium, are similar.

Correlations between the component second moments and those of the index are positive. This contrasts with the α estimate correlations, in which lead and zinc had negative relationships with the index. The β estimate correlations involving LMEX were also not uniformly positive, with nickel and lead having negative relationships with the index. Tin and zinc show moderate second moment correlations with the index, while nickel and lead second moments have a low correlation with the index moment. An explanation for the low correlation between both nickel and tin with the index may lie in the structural change observed in the second moments of both component metals. Not surprisingly, the second moments of nickel and lead are high and positively correlated. At 0.9194, nickel and lead have the highest correlation among any of the components, followed by copper and zinc with 0.8825, and aluminium and copper with 0.6482.

For tin and zinc, the long run persistence correlation dominates the second moment correlations with the index. For the α estimate correlations, tin has a high positive relationship with the index, while zinc has a low negative correlation with the index. For nickel

and lead, neither the α nor β estimate correlations with LMEX bears any resemblance to the second moment correlations. While each has a low positive correlation in the second moments with the index, nickel has a high positive and low negative correlation in the α and β estimates, respectively, and lead has a low negative correlation for both.

Second moment correlations involving only pairs of the minor components are low (either negative or positive), except for the correlation between nickel and lead. All the correlations between aluminium and the other components are moderate and positive, except with zinc, which is low and positive. The correlations between the tin second moments and the other components are also all positive, and mostly low in magnitude. Copper, nickel, lead and zinc show both positive and negative correlations with the other components, the magnitude of which can vary between 0.9194 and 0.0177.

4.6   Fourth Moment Condition
The fourth moment condition was satisfied for all the rolling models for aluminium, copper, nickel, lead and LMEX. Exceptions were tin and zinc, where the condition for the existence of the fourth moment was satisfied in 97% of the rolling models in each series. The correlation matrix for the fourth moments is given in Table 6. Substantial differences exist between the individual correlations for the fourth moments in Table 6 and the rolling second moments in Table 5. Aluminium has the highest fourth moment correlations with the index, followed closely by copper. However, the major components display only a moderate positive correlation with the fourth moments of LMEX, as compared with a high positive correlation in second moments. All minor components exhibit low positive correlations with the index. The relationship between the index and both nickel and lead is slightly stronger for the fourth moments, but the opposite is true for tin. In the case of zinc and the index, the correlation between the fourth moments is substantially lower than between the second moments.

The fourth moments for aluminium and copper have a low positive correlation with each other, even though they are the major components of the index, and are both moderately correlated with the index. Furthermore, the correlation between the aluminium and copper fourth moments is substantially lower than the correlation between their respective second moments. However, the correlation between aluminium and both nickel and lead is high in terms of the fourth moments, but moderate for the second moments. Copper, the other major component, shows a low and negative correlation in fourth moments with both nickel and lead. The correlation between nickel and lead themselves is the highest observed among the component metals, at 0.9426. This relationship is even stronger for the fourth moment than for the second moment. Negative fourth moment correlations exist between zinc and

all metals except copper. Tin is moderately correlated with aluminium, nickel and lead, but has a low positive correlation with copper, and a low negative correlation with zinc.

5. CONCLUSION

While not reflecting shocks to a metal or a sub-group of metals in the data set, the LMEX shows reasonably high correlations with short run volatility effects, but substantially lower correlations with the contributions to long run volatility effects. While the α estimates are small, the estimates and t-ratios are generally highly correlated. The β estimates and t-ratios are somewhat less correlated. In response to a shock, the ARCH volatility effects are more similar than the GARCH volatility effects across metals, and between each metal and the index. Long run persistence correlations among non-ferrous metals are frequently much greater than the corresponding correlations for the β estimates. Furthermore, the long run persistence correlation between the index and both of the major components is high, although the correlation for copper is lower than are the correlations for the α and β estimates. In contrast, the long run persistence correlation between the index and aluminium is substantially greater than for both the α and β estimates.

Long run persistence correlations may reveal relationships between fundamentals driving each metals market, such as consumption, production and stock levels, macroeconomic influences, or market related factors such as liquidity. Furthermore, aluminium and copper are the largest markets on the LME, and have high liquidity, while nickel, lead, tin and zinc do not have the same levels of liquidity.

REFERENCES

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

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

Ling, S. and M. McAleer, Stationarity and the existence of moments of a family of GARCH processes, *Journal of Econometrics*, 106, 109-117, 2002a.

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

Teyssiere, G.P., C. Brunetti and C.L. Gilbert, Non-ferrous metals price volatility: a fractionally integrated process?, in R. Smith (ed.), *Managing Metals Price Risk*, Risk Publications, London, pp. 141-152, 1997.

Yew, X.C.H., M. McAleer and S. Ling, Determining an optimal window size for modelling volatility, in A. Ullah, A. Wan and A. Chaturvedi (eds.), *Handbook of Applied Econometrics and Statistical Inference*, Marcel Dekker, New York, pp. 443-467, 2001.

Table 1: Rolling α estimate correlations

Series	Aluminium	Copper	Nickel	Lead	Tin	Zinc
Aluminium	1.0000	0.4897	0.5092	-0.4055	0.6032	0.4389
Copper	0.4897	1.0000	0.8599	0.1233	0.9206	-0.1709
Nickel	0.5092	0.8599	1.0000	0.1287	0.8217	-0.1517
Lead	-0.4055	0.1233	0.1287	1.0000	0.0594	-0.4637
Tin	0.6032	0.9206	0.8217	0.0594	1.0000	-0.0603
Zinc	0.4389	-0.1709	-0.1517	-0.4637	-0.0603	1.0000
LMEX	0.6662	0.9259	0.8660	-0.0597	0.9072	-0.0596

Table 2: Rolling α estimate t-ratio correlations

Series	Aluminium	Copper	Nickel	Lead	Tin	Zinc
Aluminium	1.0000	0.7085	0.8830	0.3082	0.9035	0.7636
Copper	0.7085	1.0000	0.8309	0.6502	0.7748	0.7709
Nickel	0.8830	0.8309	1.0000	0.4460	0.9309	0.8368
Lead	0.3082	0.6502	0.4460	1.0000	0.4087	0.4947
Tin	0.9035	0.7748	0.9309	0.4087	1.0000	0.8446
Zinc	0.7636	0.7709	0.8368	0.4947	0.8446	1.0000
LMEX	0.8953	0.8575	0.9694	0.4468	0.9472	0.8684

Table 3: Rolling β estimate correlations

Series	Aluminium	Copper	Nickel	Lead	Tin	Zinc
Aluminium	1.0000	0.4492	0.0778	0.1396	0.1784	0.4889
Copper	0.4492	1.0000	-0.1902	-0.2407	0.6247	0.5888
Nickel	0.0778	-0.1902	1.0000	0.8771	0.0421	-0.3920
Lead	0.1396	-0.2407	0.8771	1.0000	-0.0612	-0.3585
Tin	0.1784	0.6247	0.0421	-0.0612	1.0000	0.2574
Zinc	0.4889	0.5888	-0.3920	-0.3585	0.2574	1.0000
LMEX	0.5883	0.8977	-0.1611	-0.1810	0.4952	0.5629

Table 4: Rolling β estimate t-ratio correlations

Series	Aluminium	Copper	Nickel	Lead	Tin	Zinc
Aluminium	1.0000	-0.5648	0.5002	0.7049	-0.1050	-0.0212
Copper	-0.5648	1.0000	-0.2750	-0.5468	0.3056	0.0671
Nickel	0.5002	-0.2750	1.0000	0.6515	0.0813	-0.0649
Lead	0.7049	-0.5468	0.6515	1.0000	-0.0959	-0.0244
Tin	-0.1050	0.3056	0.0813	-0.0959	1.0000	-0.1980
Zinc	-0.0212	0.0671	-0.0649	-0.0244	-0.1980	1.0000
LMEX	0.3268	-0.0695	0.1887	0.3920	-0.0345	-0.1154

Table 5: Rolling second moment correlations

Series	Aluminium	Copper	Nickel	Lead	Tin	Zinc
Aluminium	1.0000	0.6482	0.5469	0.5895	0.4869	0.3634
Copper	0.6482	1.0000	-0.0785	-0.0177	0.3985	0.8291
Nickel	0.5469	-0.0785	1.0000	0.9194	0.3681	-0.2765
Lead	0.5895	-0.0177	0.9194	1.0000	0.3019	-0.2167
Tin	0.4869	0.3985	0.3681	0.3019	1.0000	0.2516
Zinc	0.3634	0.8291	-0.2765	-0.2167	0.2516	1.0000
LMEX	0.8241	0.8825	0.1822	0.2424	0.4358	0.6198

Table 6: Rolling fourth moment correlations

Series	Aluminium	Copper	Nickel	Lead	Tin	Zinc
Aluminium	1.0000	0.2929	0.7592	0.7617	0.5570	-0.1511
Copper	0.2929	1.0000	-0.1864	-0.1441	0.1443	0.6618
Nickel	0.7592	-0.1864	1.0000	0.9426	0.5393	-0.4295
Lead	0.7617	-0.1441	0.9426	1.0000	0.4564	-0.3806
Tin	0.5570	0.1443	0.5393	0.4564	1.0000	-0.0987
Zinc	-0.1511	0.6618	-0.4295	-0.3806	-0.0987	1.0000
LMEX	0.6934	0.6655	0.3299	0.3490	0.3805	0.1484