97-30R

UNIVERSITY OF CALIFORNIA, SAN DIEGO

DEPARTMENT OF ECONOMICS

CORRELATIONS AND VOLATILITIES OF ASYNCHRONOUS DATA

BY

PATRICK BURNS

ROBERT ENGLE

AND

JOSEPH MEZRICH

DISCUSSION PAPER 97-30R
APRIL 1998

Correlations and Volatilities of Asynchronous Data

March, 1998
Patrick Burns, Robert Engle and Joseph Mezrich[1]
Burns Statistics, University of California, San Diego, Morgan Stanley &Co., Inc.

ABSTRACT

Asset prices are typically measured when markets close however the closing times may differ across markets. As a result the returns appear to have predictability and correlations are understated. This will distort the value of portfolios, value at risk measures, and hedge strategies. A solution is proposed. Prices can be "synchronized" by computing estimates of the values of assets even when markets are closed, given information from markets which are open. From these prices, synchronized returns are defined and can be used to perform standard calculations including measuring time varying volatilities and correlations with GARCH. The method is applied to G7 index data.

[1] The authors wish to thank the editor for very helpful comments on an earlier draft. The current affiliations of the authors are listed.

I. Introduction

Daily data are always measured from one point in time to the same point 24 hours later. However the time of measurement is often different in markets which do not have the same trading hours. In some cases, such as the US and Japan, there are no common open hours while in others, there is partial overlap. For example, the FTSE closes at 5:00 London, but this is only 12 noon in New York. Thus any news that occurs in New York during the afternoon will not show up in British prices until the next morning and will be measured as part of the next day returns. Even if the prices are only quoted at slightly different times, there are still biases; these have been emphasized in studies of individual markets where closing prices may be stale. See for example, Scholes and Williams (1977) and Lo and MacKinlay (1990). When the times differ by many hours the effects can be dramatic. In today's global markets, these problems take on a new importance.

Asynchronous data complicates or biases many of the tasks of financial management. Perhaps most important, the value of the portfolio is never known at a point in time and consequently measures such as value or value at risk may be misleading. P&L for a company or for a trading desk can be seriously biased or even manipulated by the use of stale prices. Second, hedging must be done when markets are open, and there is a problem determining what value is to be hedged and what is the cost of the hedge. Third, correlations are generally understated and this leads to further inaccuracies in value at risk, hedge ratios, and asset allocations. Thus any analyst using daily data where prices are not measured at the same time for all assets, is potentially making systematic errors and should consider some of the solutions proposed in this paper.

We report here on a statistical model designed to solve these problems. The general approach is to recognize that even when markets are closed, the asset values may change and that new values can be estimated for use before the market reopens. Such estimated data will be called "synchronized" and can be used to estimate value, value at risk, hedge parameters and correlations. As a byproduct, estimates of the term structure of volatilities and correlations can be computed using modern GARCH procedures. See for example Engle and Mezrich (1996) for a survey.

II. Predicting Prices for Closed Markets

At the end of trading in New York, the value of a portfolio which includes London stocks, should not be measured with the closing price in London but should be measured with an estimate of the value of the London portfolio. Even though this estimate will not be exactly right, it should have no systematic biases and thus is superior for most financial applications. Consider a day when the US market falls 1% after London closes. To value the London stocks at the closing price is highly unrealistic. It will result in the US share of the portfolio declining today while the London share declines tomorrow. Value at risk calculations, asset allocations, and hedging decisions should all be considered based on the estimated value of the London portfolio. The estimated prices should have the property that they are unbiased given the information available to the market and consequently simply reassign a portion of the London returns to today rather than tomorrow. That is, average returns are unchanged; only the timing is different.

A simple picture may be useful in keeping track of these times. Consider the S&P and the FTSE which close at 4:00 in New York and 5:00 in London respectively with a 5 hour time difference. Visualize a single time line with events in New York shown below and events in London above the line.

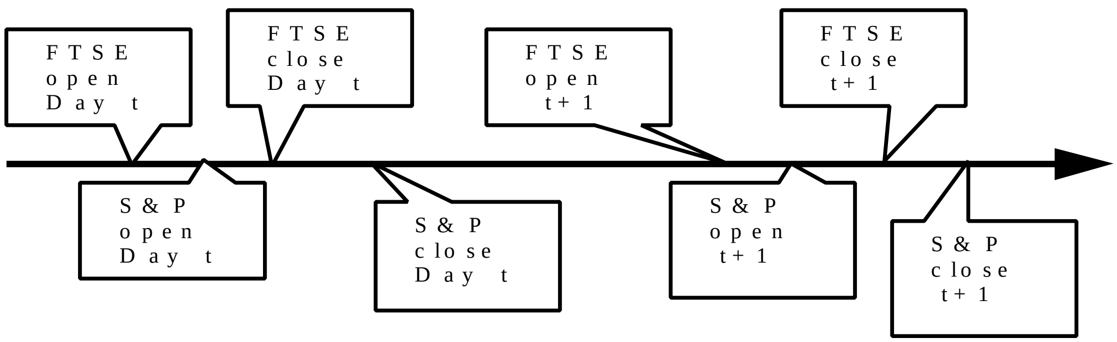

The day t accounting return is measured from close to close in both markets and consequently day t+1 returns are not perfectly synchronized. A day t return on the S&P has some overlap with day t+1 returns in London. When New York closes, there is information that can be used to predict what the London price would be if the market were open. Since it is no longer possible to trade at the London closing price, this prediction has no simple trading implications.

In order to formulate this estimation problem in the simplest fashion, assume that the New York data are closing prices denoted S_t, t=1,…, T. Thus S_1 means the price on the NYSE at 4:00 on the first day. The exchange opens 17.5 hours later at 9:30 the next morning. If we allow fractions of a day, then $S_{1.73}$ is the opening price for day 2.
Of course when it is 4:00 in New York, it is 9:00 PM in London, and the 5:00PM close in London on day 1 occurred at 12:00 in New York which is 4 hours before the New York close. Thus the London closing price can be denoted $S_{0.83}^{LON}$. In this way the estimated price in market j at time t can be described in terms of the observed prices $S_{t_j}^{j}$ as

(1) $$\log(\hat{S}_t^j) = E(\log(S_t^j)\,|\,I_t), \text{ where } I_t = \left\{S_{t_j}^j \,\middle|\, t_j \le t, j = 1,\ldots,J\right\}$$

Such prices will be called *Asynchrony Adjusted* or simply *Synchronized* prices. The logs are used to be consistent with continuously compounded returns. Clearly if S is observed at t, then its expectation is just this value. However, if the market closed before t, then the past prices in this market and all other markets that have subsequently closed, are potentially useful in predicting S at t. The expression I_t refers to the complete information set of all recorded prices at time t. Typically, this will be the set of closing prices known at the time the New York market closes, but it could include other times as well.

The synchronized prices in (1) are also unbiased estimates of the next recorded price if future changes from the synchronized prices are unpredictable.

(2) $$\log(\hat{S}_t^j) = E(\log(S_{t_j+1}^j)\big|I_t),\ \text{for } t_j \in I_t\ .$$

The variance of this prediction can be denoted

(3) $$v_t^j = V(\log(S_{t_j+1}^j)\big|I_t),\ \text{for } t_j \in I_t$$

This variance will clearly depend upon the time from t to t_j+1, upon the current volatility of the asset, and upon the value of the intervening information up to time t in predicting S.

To illustrate the approach, consider a simple example of a perfectly hedged portfolio - long Japanese ADR's traded in New York and short the underlying stocks traded in Tokyo. For the illustration we ignore currency risk. When Tokyo stocks rise, this portfolio will fall, but when the New York market opens, hours later, the ADR's will open higher, off-setting the losses. Similarly, if there is positive news in New York, the portfolio will rise temporarily until the Japanese market opens and the short positions show losses. Thus there will be strong negative autocorrelation in the measured or *accounting* value of the portfolio and it will not have a zero variance. The variability and predictability of these returns are however purely illusory and result from the use of old prices in valuing the portfolio. There are clearly no gains to be made from recognizing the predictability of prices since by the time the forecasts are made, the prices used are no longer available.

If instead, the portfolio is evaluated with synchronized prices, then the portfolio will be much more stable. When New York closes, the ADRs have a closing market value but the current value of the Japanese stocks must be inferred. Since on average, the ADRs and underlying stocks are priced the same, the estimated value of the Tokyo stocks would be just the ADR values and the total portfolio value would be zero. Similarly, the ADR prices would be estimated when Tokyo closes. With synchronized prices, this portfolio would show no variability, but is it truly riskless?

The risk in such a portfolio is truly close to zero if it can be closed at a point in time. However, since the markets are not open at the same point in time, this is not possible. Thus there is always some risk in opening or closing such a portfolio since an unhedged position will necessarily be held for some hours. The variance for this portfolio is simply a fraction of the daily variance of the returns and the value at risk is proportional to this standard deviation. These standard deviations are calculated directly from the model as formalized in (3).

If this portfolio is held for several days, the difference between average synchronized and average accounting returns is reduced. Similarly, the value at risk measures would be close and correlations between different portfolios over multi-day periods will be close.

The task is to formulate a model that can give the expected prices at any point in time given the most recent information and also the covariance matrix of all future prices. In the next section the asynchronous GARCH model will be introduced which computes both an expected return and a covariance matrix for any horizon.

III. Asynchronous GARCH

Accounting returns show cross correlations that must be modeled along with time varying volatilities and correlations. Because the correlations theoretically only have predictability for one day in the future, a first order vector moving average is the natural model. Let the vector of returns occurring in different markets measured at various times on day t be denoted R_t. Since the individual markets close within 24 hours, $\max_j \{t_j\} - \min_j \{t_j\} \leq 1$.

$$(4) \qquad R_t = \begin{pmatrix} \log(S^1_{t_1} / S^1_{t_1-1}) \\ . \\ . \\ \log(S^J_{t_J} / S^J_{t_J-1}) \end{pmatrix} \equiv \log(S_t) - \log(S_{t-1})$$

where $t = \min_j \{t_j\}$. Since New York is the last market to close, t will generally simply refer to the day in New York. However the notation will allow the analysis to be rewritten from the perspective of any market.

The Asynchronous GARCH model is formulated as a vector first order moving average with a GARCH covariance matrix as in

$$(5) \qquad R_t = \varepsilon_t + M\varepsilon_{t-1}, \quad V_{t-1}(\varepsilon_t) = H_t$$

where M is the moving average matrix and H is the covariance matrix of ε, the unpredictable part of returns from the perspective of time t-1. For a first order vector moving average with J assets, there are J^2 elements in M which must be estimated to match the J^2 first order autocorrelation coefficients. In world of efficient markets, the diagonal and below diagonal elements of the autocorrelation matrix should all be zero since these are times with no overlap. This only implies that the M matrix should have this form if the innovations are uncorrelated. However, this will not generally be the case, so the implication is only that the last row of M must be zero. Nevertheless, empirically the M matrix is typically found to be upper triangular with some non-zero elements on the diagonal corresponding to markets which have serially correlated returns.

Defining the *synchronized returns* as the change in the log of the synchronized prices given by

$$(6) \qquad \hat{R}_t = \begin{pmatrix} \log(\hat{S}^1_{t_1} / \hat{S}^1_{t_1-1}) \\ . \\ . \\ \log(\hat{S}^J_{t_J} / \hat{S}^J_{t_J-1}) \end{pmatrix} = \log \hat{S}_t - \log \hat{S}_{t-1}$$

The estimated change in value from the end of one day until the next is the sum of the innovation ε and its impact on the future opening price Mε. Substituting (2)(4) and (5) into (6) gives the synchronized return as

$$\begin{aligned}\hat{R}_t &= E_t \log(S_{t+1}) - E_{t-1}\log(S_t) \\ &= E_t R_{t+1} - E_{t-1} R_t + \log S_t / S_{t-1} \\ &= M\varepsilon_t - M\varepsilon_{t-1} + R_t\end{aligned}$$

(7) $$\hat{R}_t = \varepsilon_t + M\varepsilon_t, \quad V_{t-1}(\hat{R}_t) = (I+M)H_t(I+M)'.$$

$V_{t-1}(\hat{R}_t)$ is a positive definite covariance matrix for all time periods, since H_t is positive definite, and the synchronized returns are serially uncorrelated since the epsilons are serially uncorrelated. Furthermore, the mean of the synchronized and the unsynchronized returns will be almost identical over any sample period since these differ only insofar as the means, $\bar{\varepsilon}_t \neq \bar{\varepsilon}_{t-1}$. Furthermore, any such differences are multiplied by M which is mostly zero with a few small non-zero elements. The variances and covariances however will be different and typically larger for the synchronized returns since some of the variability of the accounting returns is spread across days.

The unconditional moments of the accounting and synchronized returns can be directly computed from (7) and (5). Letting $E(\varepsilon_t\varepsilon'_t) = \Omega = E(\varepsilon_{t-1}\varepsilon'_{t-1})$,

(8) $$E(R_t R_t') = \Omega + M\Omega M', \quad E(R_t R_{t-1}') = M\Omega, \quad E(\hat{R}_t \hat{R}_t') = (I+M)\Omega(I+M)'$$

An estimate of the covariance matrix of synchronized returns can be obtained from the contemporaneous and lagged covariances of accounting returns. An unbiased estimate of the covariance matrix of synchronized returns can be formed as:

(9) $$\hat{E}(\hat{R}_t\hat{R}_t') = \frac{1}{T}\sum_{t=1}^{T}(R_t R_t' + R_t R_{t-1}' + R_{t-1}R_t')$$

However, there is no guarantee that this covariance matrix is positive definite. This result is simply a generalization of the Scholes and Williams(1977) result.

The term structure of correlations and volatilities of the synchronized returns can be directly calculated using the parameters of the GARCH model. In particular, letting $H_{t+k,t}$ be the covariance matrix of returns on day t+k given information on day t, then the term structure of correlations and volatilities is easily computed from the covariance matrices:

(10) $$V_t\left(\sum_{j=1}^{k}\hat{R}_{t+j}\right) = (I+M)\sum_{j=1}^{k}H_{t+j,t}(I+M)'$$

This differs from the calculation of the term structure of correlations of measured or accounting returns. For example assuming (5) to be correctly specified:

(11) $$V_t(R_{t+1}+R_{t+2}) = V_t(\varepsilon_{t+2}+(I+M)\varepsilon_{t+1}) = H_{t+2}+(I+M)H_{t+1}(I+M)'$$

or in general

(12) $$V_t\left(\sum_{j=1}^{k-1} R_{t+j}\right) = H_{t+k,t}+(I+M)\sum_{j=1}^{k-1} H_{t+j,t}(I+M)'$$

This covariance matrix therefore does not include the imputed value accruing after markets close on the last day of the forecast period. However since only one term is different between (11) and (12), the average volatilities and correlations will be very similar over multiple day holding periods.

IV. EMPIRICAL FINDINGS

The G-7 equity markets exhibit very clearly the problems with asynchronous data. As these markets span the globe they include markets with almost perfectly synchronous data and markets which are completely out of phase. Estimates of correlations will typically be too small when the markets are highly asynchronous. In Table I, for example the highest correlations are between the CAC, DAX and FTSE, and between the S&P500 and Toronto. Much lower are the correlations of the Nikkei with all others

TABLE I
CORRELATIONS OF G-7 EQUITY RETURNS
Jan 2,1990 to Oct 3,1996

	CAC	DAX	FTS	MIL	NIK	S&P	TTO
CAC	1.000	0.597	0.617	0.319	0.246	0.286	0.252
DAX	0.597	1.000	0.448	0.411	0.275	0.248	0.224
FTS	0.617	0.448	1.000	0.273	0.261	0.330	0.316
MIL	0.319	0.411	0.273	1.000	0.193	0.139	0.145
NIK	0.246	0.275	0.261	0.193	1.000	0.146	0.191
S&P	0.286	0.248	0.330	0.139	0.146	1.000	0.557
TTO	0.252	0.224	0.316	0.145	0.191	0.557	1.000

Of course there is no reason to believe that all correlations should be high. One indication of the correct average or unconditional correlation can be found from time aggregated data. Such data are relatively unaffected by the timing of markets since the degree of asynchronicity is much less. In the following table correlations between weekly returns of the same indices are presented, where a week is defined as five trading days. In most cases the numbers are larger. In particular, the correlation between S&P and NIK is now .298 rather than .146, and the correlation between S&P and the three big European indices, the FTSE, CAC and DAX rises from .2or .3 up to .4 or .5.   At the same time the within European correlations are little changed except for Milan.

TABLE II
WEEKLY CORRELATIONS G-7
Jan 2,1990 to Oct 3,1996

	WCAC	WDAX	WFTS	WMIL	WNIK	WS&P	WTTO
WCAC	1.000	0.663	0.576	0.465	0.308	0.447	0.359
WDAX	0.663	1.000	0.528	0.490	0.281	0.408	0.327
WFTS	0.576	0.528	1.000	0.373	0.311	0.495	0.423
WMIL	0.465	0.490	0.373	1.000	0.233	0.268	0.216
WNIK	0.308	0.281	0.311	0.233	1.000	0.298	0.296
WS&P	0.447	0.408	0.495	0.268	0.298	1.000	0.634
WTTO	0.359	0.327	0.423	0.216	0.296	0.634	1.000

If these differences are due to asynchronicity, then they should also appear in lag effects. It should appear that later closing markets forecast earlier markets. In table III, the daily first order correlations are tabulated. The largest elements are in the last two rows, revealing that US and Canadian returns forecast the earlier markets. The FTSE predicts the DAX and the CAC, DAX and FTSE all forecast the NIKKEI. Milan and Toronto both exhibit substantial autocorrelation which could be a consequence of stale quotes in the closing index. This is plausible in markets without high volume.

TABLE III
LAGGED CROSS CORRELATIONS OF G-7
Jan 2,1990 to Oct 3,1996

	CAC	DAX	FTS	MIL	NIK	S&P	TTO
LCAC	0.029	0.160	-0.001	0.223	0.161	0.039	0.090
LDAX	-0.018	0.010	-0.006	0.140	0.126	-0.014	0.063
LFTS	0.012	0.153	0.040	0.184	0.140	0.001	0.088
LMIL	0.024	0.021	0.013	0.201	0.063	0.010	0.033
LNIK	-0.035	-0.056	-0.042	0.024	0.019	-0.023	0.031
LS&P	0.208	0.306	0.267	0.246	0.223	0.053	0.206
LTTO	0.116	0.151	0.153	0.148	0.137	0.050	0.260

These results are simply for unconditional correlations. With the Asynchronous GARCH model discussed above, the correlations between accounting returns and the correlations between synchronized returns can both be calculated at each moment of time and forecasts of these can be computed. The long run forecasts of both synchronized and unsynchronized returns should be close to the weekly correlations tabulated above.

V. EMPIRICAL RESULTS FOR SYNCHRONIZED RETURNS

The model described in equation (5) requires specification of the GARCH covariance matrix as well as the moving average model for the mean. Using G7 equity data from July 1987 through October 1996, a model was formulated, estimated by maximum likelihood and subjected to diagnostic testing. The crash was included in the data set because it was an event with very important correlations, however in much of the analysis below, the focus will be on the 90's. The GARCH model is a component model (See Engle and Lee(1993) and Engle and Mezrich(1996)) of the BEKK type (See Engle and Kroner(1995)) which includes some leverage terms and the moving average means. Altogether there are 96 estimated parameters. The model and the diagnostic tests will be presented in the next section. In this section the moving average matrix and the synchronization of returns will be presented.

The moving average matrix M from equation (5) was specified to have mostly zeroes with a few non-zero estimated parameters. The estimated return equation is given by:

$$(13) \qquad R_t = e_t + Me_{t-1}, \quad M' = \begin{bmatrix} 0 & .08 & 0 & .09 & .07 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & .11 & 0 & .08 & .10 & 0 & 0 \\ 0 & 0 & 0 & .21 & 0 & 0 & 0 \\ -.01 & -.05 & -.03 & 0 & 0 & 0 & 0 \\ .29 & .49 & .28 & .27 & .28 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & .21 \end{bmatrix}, \quad R = \begin{bmatrix} r_{CAC} \\ r_{DAX} \\ r_{FTS} \\ r_{MIL} \\ r_{NIK} \\ r_{S\&P} \\ r_{TTO} \end{bmatrix}$$

where the variables are ordered (CAC, DAX, FTSE, MIL, NIK,S&P, TTO) and rows predict columns just as in the correlation tables. All coefficients have t-ratios greater than three except the three negative coefficients in Japan and the .07 coefficient in France. The row with the largest coefficients is the S&P row which means that there is substantial predictability of all other markets given US returns. This is natural because the US closes last. Although Canada also showed predictability, this appears to be primarily a result of autocorrelation and a somewhat similar story applies to Italy.

Based on this moving average matrix and the computed innovations, ε, the synchronized returns can be directly calculated.

TABLE IV
CORRELATIONS OF SYNCHRONIZED RETURNS
Jan 2,1990 to Oct 3,1996

	SCAC	SDAX	SFTS	SMIL	SNIK	SSPX	STTO
SCAC	1.000	0.703	0.650	0.437	0.343	0.441	0.308
SDAX	0.703	1.000	0.578	0.481	0.298	0.482	0.312
SFTS	0.650	0.578	1.000	0.381	0.321	0.530	0.381
SMIL	0.437	0.481	0.381	1.000	0.236	0.295	0.200
SNIK	0.343	0.298	0.321	0.236	1.000	0.298	0.250
SSPX	0.441	0.482	0.530	0.295	0.298	1.000	0.565
STTO	0.308	0.312	0.381	0.200	0.250	0.565	1.000

Examination of these correlation matrices reveals that in most cases, the synchronized return correlations are much closer to the weekly correlations than the daily. For example, the S&P correlation with the NIKKEI in daily data is only .146 while in both the weekly data and the synchronized data it is .298. The correlation with the FTSE is .330 in daily data, .495 in weekly data and .530 in the synchronized data. Even in markets with relatively small non-synchronicity, the effects are noticeable. The CAC – DAX correlations are .597 daily, .663 weekly and .703 synchronized, while MIL-FTS correlations rise from .273 to .373 and .381.

What other properties do synchronized data have? Table V presents the annualized means and volatilities of the accounting and synchronized returns.

TABLE V
MEANS AND VOLATILITIES OF
ACCOUNTING AND SYNCHRONIZED RETURNS
Jan 2,1990 to Oct 3,1996

	CAC	DAX	FTS	MIL	NIK	SPX	TTO
Mean Acc.	0.93	5.75	7.17	-1.00	-8.59	9.62	4.34
Mean Syn.	0.93	5.73	7.16	-1.00	-8.61	9.62	4.34
Vol. Acc.	17.49	16.80	12.42	18.47	23.14	11.35	8.54
Vol. Syn.	18.23	18.76	13.16	22.50	23.87	11.35	9.97

It is clear that the means are virtually unchanged by calculating returns on a synchronized or accounting basis, but the volatilities are slightly higher for the synchronized returns. This is not surprising since volatilities of asynchronous data are theoretically too low as first pointed out by Scholes and Williams (1977).

VI. DYNAMIC CORRELATIONS

The GARCH model can be used to forecast volatilities and correlations of either the accounting data using equation (12) or the synchronized data using equation (10). For longer horizons however the results should be very similar. Several examples illustrate this point. In Figure 1, the term structure of correlations between MILAN and S&P and between MILAN and FTSE are presented for the last date in the sample, October 3, 1996. In both cases, the very short term correlations are much lower than the longer horizons as was already illustrated by the daily vs. weekly unconditional correlations. The curves approach the unconditional or theoretical synchronized correlations as the horizon gets large. The FTSE on this date overshoots the long run but then decays toward it while the S&P moves rather quickly to the long run value.

If such a correlation term structure is computed starting on each day of the sample period then by selecting the 30 day correlation, a constant maturity graph can be constructed. In Figure 2, the 30 day correlations between DAX and MIL and S&P are presented over the entire sample period. As can be seen the DAX-MIL correlations are far more volatile ranging from .25 to .6 while the S&P-DAX correlations drop only to .4 and rise only to .6 except just after the ’87 crash.

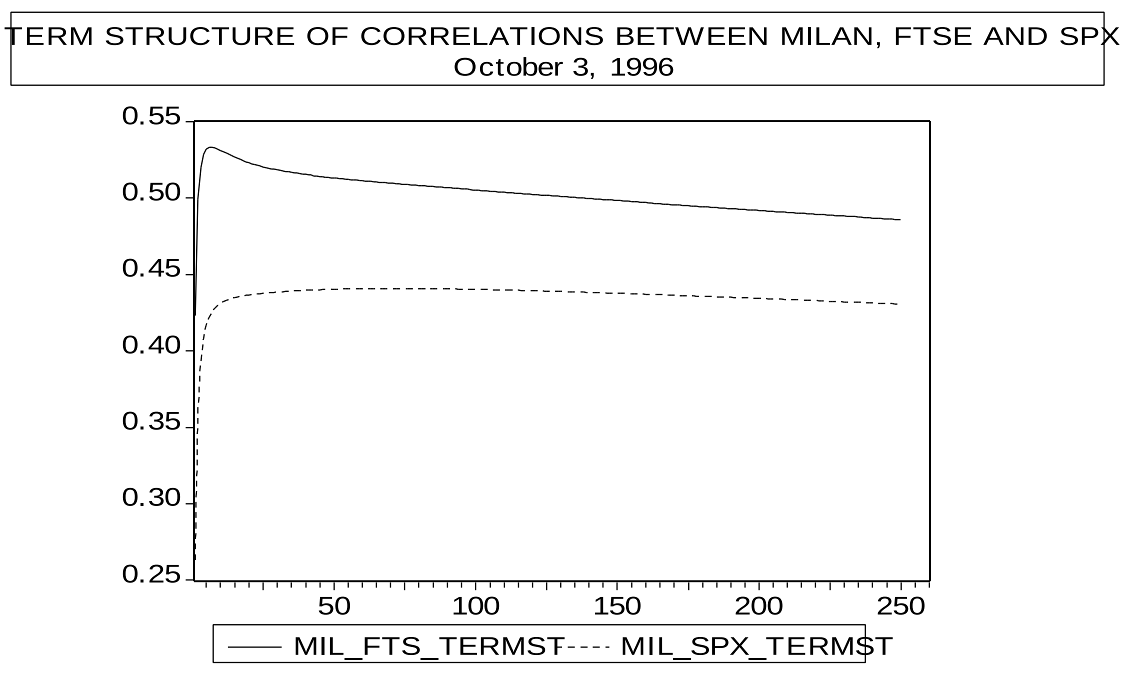

Figure 1

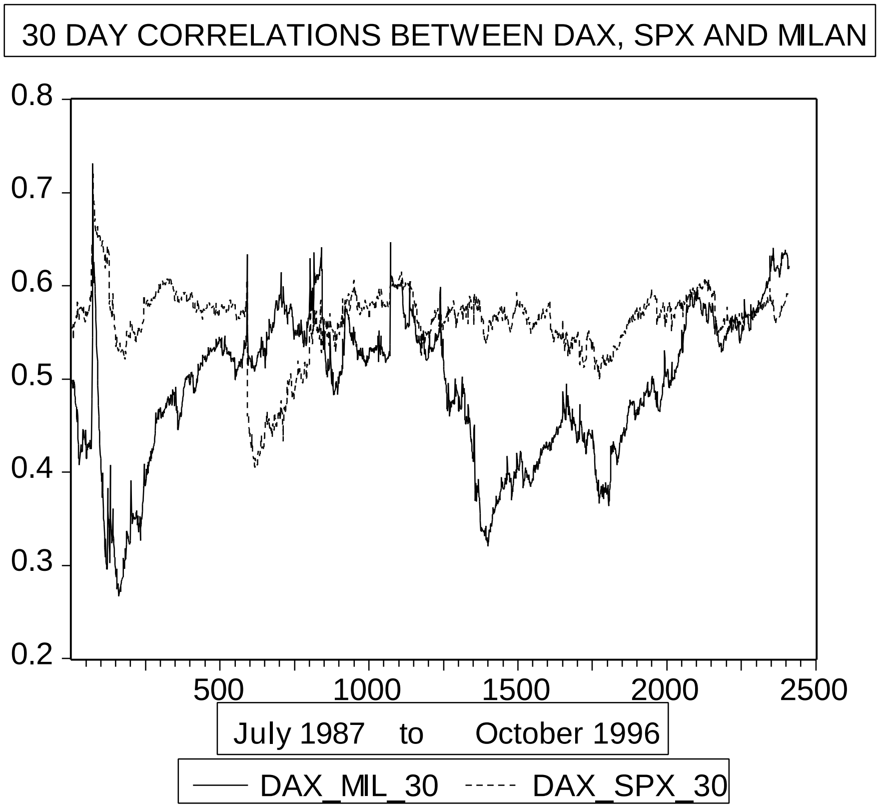

Figure 2.

VII. THE ASYNCHRONOUS GARCH MODEL AND DIAGNOSTICS

The model estimated is of the Component Form – See Engle and Lee(1993) and Engle and Mezrich (1996) which can be described concisely as

(14)
$$H_t = \sum_{k=1}^{K}\left(A_k'[\varepsilon_{t-1}\varepsilon'_{t-1} - Q_t]A_k + B_k'[H_{t-1} - Q_t]B_k + C_k'[\eta_{t-1}\eta'_{t-1} - \frac{1}{2}Q_{t-1}]C_k \right)$$
$$Q_t = \sum_{k=1}^{K}\left(R_k'[Q_{t-1} - \hat{\Omega}]R_k + F_k'[\varepsilon_{t-1}\varepsilon'_{t-1} - H_t]F_k \right)$$

where H is the conditional covariance matrix of the innovations, ε, Q is the permanent part of this covariance process so that H-Q is the transitory component, $\hat{\Omega}$ is the sample covariance matrix of ε, and $\eta = \min(\varepsilon, 0)$ where the minimum is calculated for each element of the vector. In this model, K=3 and the dimension of ε is 7x1 so that all the matrices are 7x7. While it may appear that such a model has 735 parameters, most are taken to be zero. Only 80 parameters plus the moving average parameters are estimated. The non-zero parameters corresponding to this structure along with their standard errors and t-ratios are given in the Appendix to Burns, Engle and Mezrich(1997). The model is estimated by maximizing the log likelihood using an efficient genetic algorithm.

(15)
$$L = -\frac{1}{2}\sum_{t=1}^{T}\left(\log\left|H_t\right| + \varepsilon_t' H_t^{-1}\varepsilon_t \right)$$

Associated with this model is a collection of diagnostic tests. These are tests against various forms of dependence in the standardized residuals. If the model is correct, the innovations, ε, should be serially uncorrelated and have a conditional covariance matrix, H. Thus the standardized residuals defined by $\tilde{\varepsilon} = H^{-1/2}\varepsilon$, should have mean zero and identity covariance matrix conditional on any past information set. Several tests of this hypothesis are presented. Each of these tests examines the performance in this sample and therefore could naturally be supplemented by out-of-sample tests of exactly the same form.

The first test, in Table VI.a, simply looks at the contemporaneous covariance matrix to check whether it is approximately the identity matrix.

Table VI.a
Covariance of Standardized Residuals:

INDEX	NIK	DAX	MIL	CAC	FTS	S&P	TTO
NIK	1.030386	-0.040918	-0.042275	-0.00606	-0.0227	0.00277	-0.05615
DAX	-0.040918	0.969820	0.008062	-0.01707	-0.0216	-0.01452	0.002657
MIL	-0.042275	0.008062	1.011082	-0.00216	-0.0135	-0.02169	-0.01669
CAC	-0.006062	-0.017076	-0.002163	1.00902	0.00745	-0.03888	-0.02127
FTS	-0.022707	-0.021689	-0.013560	0.00745	1.0130	-0.01685	-0.06074
S&P	0.002771	-0.014523	-0.021697	-0.03888	-0.0168	1.02411	-0.02066
TTO	-0.056157	0.002657	-0.016698	-0.02127	-0.0607	-0.020665	1.032501

The Chi square statistic that tests whether this is an identity matrix is 41.85 on 28 degrees of freedom and has a p-value of 0.0448. Thus, this matrix of unconditional covariances is very close to satisfying the identity assumption for the standardized residuals.

Subsequent tests are all dynamic tests of various types of temporal dependence. In each case the Ljung-Box statistics test for various kinds of time dependence. All of these test statistics are computed using 15 lags so the asymptotic 5% critical value is 25. These are computed based on the *rank* autocorrelations and cross correlations, rather than the conventional autocorrelations; consequently the tests are far more robust to outliers in the data set and typically exhibit more power. In each case the table is organized so that rows predict columns.

Table VI.b presents tests of autocorrelation in standardized returns. If the moving average model truly captures the dynamics of these markets, then there should be no remaining autocorrelation or cross correlation There are 5 out of 49 entries exceeding this value which could suggest the need for a slightly richer moving average model.

Table VI.b
Ljung-Box Tests for Cross Correlation of Standardized Residuals

INDEX	NIK	DAX	MIL	CAC	FTS	S&P	TTO
NIK	13.17	7.05	8.862	22.80	23.73	15.12	15.905
DAX	16.31	16.57	18.824	11.84	17.62	13.97	16.970
MIL	11.91	10.81	14.411	14.64	23.95	18.75	16.808
CAC	16.48	13.75	18.228	11.11	41.56	18.91	17.454
FTS	11.05	19.27	12.236	22.14	13.87	12.65	6.903
S&P	26.28	44.07	14.064	14.40	14.39	21.44	26.424
TTO	12.55	16.15	18.219	14.30	16.60	15.10	31.818

Table VI.c gives similar results for squared standardized residuals. Significant diagonal elements therefore indicate failures with respect to own information sets and off-diagonal elements suggest cross country causality in variance. Now there are 6 significant entries out of 49. There is slight weakness in the univariate models for Canada, UK, Italy, and Japan as well as some predictability of the US and Canadian volatility based on German volatility. Overall however, these statistics are pretty good.

Table VI.c
Ljung-Box Tests for Squared Standardized Residuals

INDEX	NIK	DAX	MIL	CAC	FTS	S&P	TTO
NIK	25.97	11.119	11.70	15.270	14.281	17.65	8.719
DAX	17.40	21.198	23.36	16.270	10.180	30.89	34.313
MIL	17.10	11.251	30.84	16.958	13.903	16.82	15.405
CAC	21.83	19.889	13.49	7.407	8.295	24.24	17.316
FTS	11.04	12.908	13.36	18.493	28.365	15.79	6.118
S&P	13.58	8.852	11.78	12.156	17.518	14.79	15.471
TTO	22.98	15.503	19.67	17.936	12.537	16.09	29.055

Table VI.d examines multivariate leverage effects by checking whether the level of return in one market is correlated with subsequent squared returns in another market. There are 4 out of 49 significant entries here. The intriguing statistics here are the suggestions that directional moves in France predict volatility in Germany and Italy.

Table VI.d
Ljung-Box Tests for Leverage Effects of Standardized Residuals

INDEX	NIK	DAX	MIL	CAC	FTS	S&P	TTO
NIK	22.156	7.831	15.65	19.41	14.640	20.32	14.532
DAX	18.054	13.979	10.95	16.98	9.818	14.87	15.708
MIL	15.798	15.300	26.70	16.45	16.480	18.54	25.260
CAC	21.532	29.215	28.44	22.32	7.731	15.64	14.528
FTS	9.412	8.557	11.50	16.31	14.843	6.92	9.822
S&P	5.414	16.372	10.81	20.41	5.958	20.41	10.453
TTO	20.638	12.491	16.67	13.10	14.178	11.30	15.987

Table VI.e examines the autocorrelation of cross products of standardized residuals. For example, the first element is an autocorrelation test of the series made up of returns in Japan times returns in Germany, each standardized by the model. The test is expected to reveal failures in the estimates of covariances. The statistics look fine.

Table VI.e
Ljung-Box Autocorrelation Tests of Cross Products of Standardized Residuals:

INDEX	DAX	MIL	CAC	FTS	S&P	TTO
NIK	21.7856	11.4359	34.3772	16.9797	9.4629	6.5749
DAX		8.7119	8.7200	15.8421	22.5017	6.7044
MIL			17.7490	12.0109	18.8546	18.2795
CAC				19.7531	25.1769	7.3293
FTS					18.3938	13.3379
S&P						14.9570

Overall there are 16 entries which exceed 25 in the "standardized residual" tables out of a total of 169. Several of these are only a fraction above 25 and thus the results broadly support the specification. The model was selected from simpler models by examining diagnostic tables such as these; this is a natural method for picking a parsimonious model, but it does leave the exact significance level of the tests uncertain due to pretest effects.

VIII. CONCLUSIONS

This paper has presented an approach to computing synchronized returns and their term structures of volatilities and correlations from asynchronous data. The central task is estimating the value of assets that are not trading so that the portfolio value can be estimated at a moment of time. From these estimated prices, synchronized returns can be defined which have similar properties to accounting returns when aggregated over time, but can also be used for modeling time varying correlations, volatilities and hedge ratios.

Synchronized returns and their time varying characteristics are computed from a multivariate GARCH model with a first order vector moving average called an Asynchronous GARCH model. From the same model, characteristics of the accounting returns can be calculated.

Any analyst using daily data where prices are not measured at the same time for all assets, is potentially making systematic errors by using observed or accounting prices since some of these will be stale. The use of weekly data will reduce these biases, though not to zero, but may hide some of the detail which is of interest. Global portfolios have distorted values when accounting data are used, and estimated value at risk measures will consequently be incorrect. P&L for a company or for a trading desk can be seriously biased or even manipulated by the use of stale prices. Because correlations are on average too low, hedge ratios will likely be too low unless the data are synchronized. Portfolio rebalancing should be done with a good measure of the value of each asset at a common time; otherwise the consequences could be different from those desired. It is commonly said that international diversification is not as valuable as it would appear, and part of this may be the underestimation of correlations.

References

Burns, P., R.F. Engle and J. Mezrich,(1997) “Correlations and Volatilities of Asynchronous Data” U.C.S.D. Discussion Paper 97-30

Engle, R. F. and G.J. Lee, (1993) "A Permanent and Transitory Component Model of Stock Return Volatility", Manuscript, UCSD

Engle, R.F. and K. Kroner, "Multivariate Simultaneous GARCH," *Econometric Theory*, 11 (1995): 122-150.

Engle, R.F. and J. Mezrich (1996) “GARCH for Groups”, *Risk*, August, Vol. 9, No 8: pp36-40

Lo, A. and A.C.MacKinlay, (1990), “An Econometric Analysis of Nonsynchronous Trading,” *Journal of Econometrics*, 40, pp 203-238

Scholes, M. and J. Williams, (1977), “Estimating Betas from Nonsynchronous Data,” *Journal of Financial Economics*, 5, 309-328