TIME VARIATION IN THE CORRELATION STRUCTURE OF EXCHANGE RATES: HIGH-FREQUENCY ANALYSES

JAYARAM MUTHUSWAMY
SUDIPTO SARKAR*
AARON LOW
ERIC TERRY

The correlation structure of asset returns is a crucial parameter in risk management as well as in theoretical finance. In practice, however, the true correlation structure between the returns of assets can easily become obscured by time variation in the observed correlation structure and in the liquidity of the assets. We employed a time-stamped high-frequency data set of exchange rates, namely, the US$–deutsche mark and the US$–yen exchange rates, to calibrate the observed time variation in the correlation structure between their

The authors would like to thank an anonymous referee and the editor, Robert Webb, for their suggestions and recommendations. The authors are responsible for any remaining errors.
*Correspondence author, Finance and Business Economics Department, Marshall School of Business, University of Southern California, Los Angeles, California 90089-1427; e-mail: sudisarkar@cs.com

Received February, 2000; Accepted July, 2000

- *Jayaram Muthuswamy is an Associate Professor in the Department of Finance and the Faculty of Economics at the University of Sydney in Sydney, New South Wales, Australia.*
- *Sudipto Sarkar is in the Finance and Business Economics Department in the Marshall School of Business at the University of Southern California in Los Angelos, California.*
- *Aaron Low is the Head of Fixed Income and the Asia and Senior Fund Manager at Dresdner Asset Management in Singapore.*
- *Eric Terry is at Charis Consulting in Markham, Ontario, Canada.*

The Journal of Futures Markets, Vol. 21, No. 2, 127–144 (2001)

returns. We also documented time variation in the liquidity structure of these rates. We then attempted to link the observed correlations with the liquidity via an application of an illiquid trading model first developed by Scholes and Williams (1976). We show that the observed correlation structure is strongly biased by the liquidity and that it is possible to effect at least a partial rectification of the otherwise downward-biased observed correlation. The rectified sample correlation is, therefore, more appropriate for input into models used for forecasting, option pricing, and other risk management applications.

INTRODUCTION

The global foreign exchange market is operational on the 24-hr clock. After the opening quotes from major interbank players in New Zealand, Australia, Japan, and countries in the Far East, the markets of Europe kick in, followed by those in North America. The major capitals in this 24-hr cycle are, therefore, Sydney, Tokyo, Hong Kong, Singapore, Zurich, Frankfurt, London, and New York, in chronological order. An elegant description of this was given by Dacarogna et al. (1993) and Muller et al. (1990).

Bankers and corporate treasurers who perform hedging functions for their clients are constantly keeping track of fluctuations in foreign exchange rates. For interbank foreign exchange dealers, the task of maintaining and monitoring open speculative positions in currency is straightforward enough, if dealers are located in the various financial capitals that straddle the 24-hr clock. However, for intraday arbitrage-based dynamic hedging strategies, the viability of the hedging operations is also rather dependent on the ability to keep track of certain hedge ratios. For currency hedges, a necessary statistic to track is often the simple correlation coefficient between the rates of return with respect to the hedged assets. Usually, hedge ratios are estimated by simple regression analysis.

A more recent application arises in the pricing of exotic currency exchange options. In this case, the correlation coefficient between the returns of two exchange rates is actually an exogenous input pricing parameter, in much the same way that the standard deviation of a stock's rate of return is for the Black–Scholes model. Similarly, for the pricing of currency quanto options (Rubinstein, 1990), as well as options on the minimum or maximum of two assets (Stulz, 1982), the correlation between the rates of return of the two assets comprising the state variables is an input option pricing parameter. Consequently, any inaccuracies in measuring the true contemporaneous correlation will cause numerical biases to arise in the estimated value of the option. For instance, in valuing options based on the maximum or minimum of two assets, the value of the option depends very much on the extent of the correlation of the assets. If, as a result of liquidity-induced observation problems (e.g., the kind discussed later in this article), a treasurer were to estimate incorrectly the true correlation, inaccurate option values would result with potential consequences for hedging applications.

Similarly, when other exotic options such as term structure spread options are priced, the correlation between the rates of return on the long and short bonds is an input parameter. One example of such an option pricing model was given by Pearson (1995), who priced call options on the difference between two asset values.

Finally, in tests of multivariate asset pricing theory, the covariance matrix of asset returns is a crucially important structure to estimate. Any time variation in the correlation structure must, therefore, be properly treated econometrically if the test results are to be reliable.

The plan of this article is as follows: In the section entitled Time Variation in Exchange-Rate Correlations, we present sample evidence of time variation in the correlation structure of various pairs of exchange rates together with evidence on time variation in trading liquidity. We also discuss how trading liquidity can be a possible source of variation in observed correlation coefficients. In the section entitled A Model of Illiquid Trading, we describe a transactions-based theoretical infrequent trading model derived from Scholes and Williams (1976, 1977) and include a discussion on its relevance toward explaining time variation in the observed correlations. In the section entitled Data and Further Tests, we first review the nature of the special quotations-based high-frequency data used for this study and then fit the Scholes–Williams model to this quote data. The final section is a summary of the overall findings together with concluding comments.

Time Variation in Exchange-Rate Correlations

The basic notion of observed temporal variations in the contemporaneous correlation structure of exchange rates is illustrated in Table I, which shows correlation coefficients between the returns of the US$–yen and US$–deutsche mark (DM) exchange rates, according to the day of the week.[1] We can see that the correlation is not constant over the week. In fact, the correlation coefficient increases monotonically from Monday

[1] A detailed description of the data base used for this study is given in the Data and Further Tests section.

TABLE I

Contemporaneous Correlations between the US$–Deutsche and the US$–Yen Exchange Rates by the Day of the Week

Day of the Week	*Correlation Coefficient*	*Number of Observations*
Monday	.456	51
Tuesday	.503	50
Wednesday	.542	52
Thursday	.503	52
Friday	.318	50

Note: Sample period is October 1, 1992 through September 30, 1993.

through Wednesday, followed by a monotonic decline through Friday. Friday's correlation is particularly weak compared with the rest of the week.

In principle, the returns of exchange rates should be strongly contemporaneously correlated with one another, primarily because of the common informational shocks simultaneously affecting currencies. However, observed correlations may become attenuated by the high-frequency returns sampling process. Indeed, Epps (1979) showed that as the sampling interval progressively shrinks, the otherwise significant contemporaneous correlations between stock price returns are systematically reduced to white noise. Epps did not quite give the precise cause of this attenuation of the observed contemporaneous correlation. Muthuswamy (1990) attributed this systematic destruction of the sample correlation for high-frequency data to nonsynchronicities in the recording process.[2] This is illustrated in Figure 1, which shows the contemporaneous correlation coefficient between the US$–yen and US$–DM exchange-rate returns for various holding periods ranging from 5 to 2,500 min. This is interesting for several reasons. First, the Epps effect comes through clearly at very high sampling frequencies; when the returns are observed at intervals of less than 30 min, the contemporaneous correlations rapidly decay to zero. Second, once a certain horizon length is reached (about 2 hr), the correlations tend to remain fairly stable, suggesting that there exists a true equilibrium correlation that does not vary with time. As the horizon becomes very large, instability sets in again, but only because the effective sample size diminishes enough that the estimator's standard error becomes very large.

Figure 2 shows the correlation between the US$–DM and US$–yen exchange rates on a 5-min return basis observed over a 24-hr clock (i.e.,

[2]See Cohen et al. (1979), Lo and MacKinlay (1990), and Muthuswamy (1990) for other expositions of the effects of nonsynchronicities on measured returns.

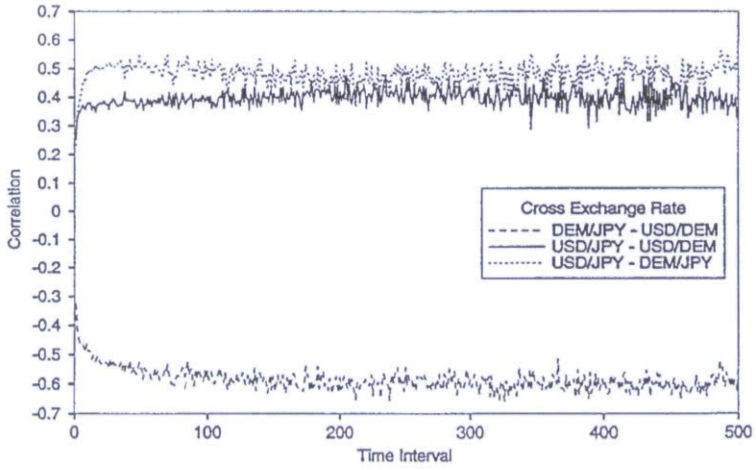

FIGURE 1

Observed contemporaneous correlations between returns for the US$–DM, US$–yen, and DM–yen exchange rates for varying holding periods. The correlations are based on the standard moments estimator and have been computed with high-frequency data with weekends and public holidays omitted. The holding period of the returns range from 5 min to 2,500 min.

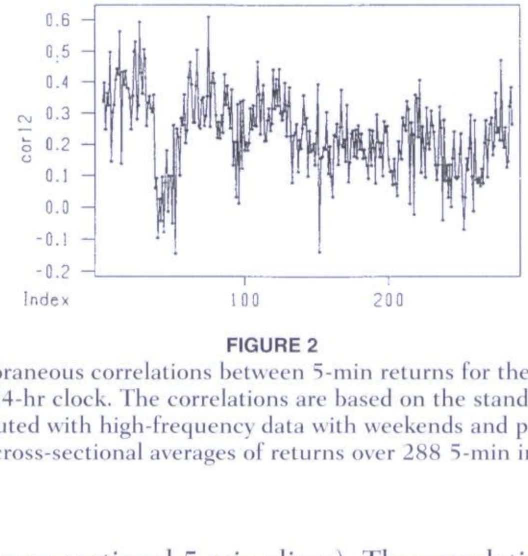

FIGURE 2

Observed contemporaneous correlations between 5-min returns for the US$–DM, US$–yen exchange rates on a 24-hr clock. The correlations are based on the standard moments estimator and have been computed with high-frequency data with weekends and public holidays omitted. The correlations are cross-sectional averages of returns over 288 5-min intervals covering 24 hr.

with 288 cross-sectional 5-min slices). The correlations fluctuate rather significantly, which is not surprising when one considers that the Epps effect mentioned earlier is operational. Despite this and with an effective sample size of 261 cross-sectional slices (corresponding to 261 actual trading days), a strong enough correlation structure is observed; indeed, the average 5-min correlation over the 288 slices of the day is .2388, compared to a correlation of .464 based on daily returns. The much higher correlation for daily returns is intuitive simply because there is

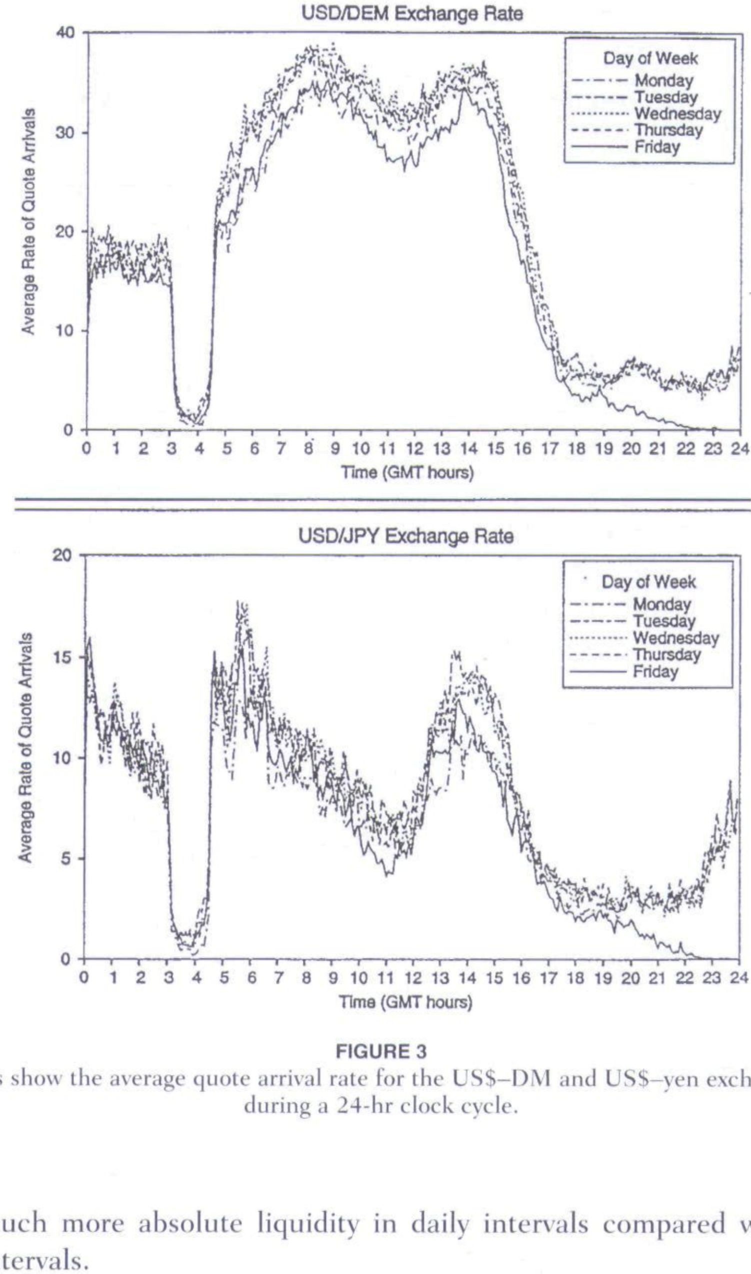

FIGURE 3
The plots show the average quote arrival rate for the US$–DM and US$–yen exchange rates during a 24-hr clock cycle.

much more absolute liquidity in daily intervals compared with 5-min intervals.

What is more interesting about Figure 2 is that certain systematic variations are inherent. For instance, observed 5-min correlations exhibit a pronounced dip (toward zero) at about the same time when the Asia–Pacific region has its lunch break. Other systematic patterns exist but have to be studied in conjunction with Figure 3.

Figure 3 shows the average number of arrivals of bid–ask quotes over 5-min intervals for the US$–DM and US$–yen exchange rates for each day of the week. Because of negligible activity, weekends have been omitted. The US$–DM exchange rate has a greater number of quote arrivals with a total of 1,472,241 bid–ask quotes for the entire sample period or an average of 4,034 per day, whereas the US$–yen rate has a daily arrival quote rate of 1,564. It is also interesting to qualitatively examine the intraday quote arrival patterns, as there seems to be a similar pattern for both rates. Trading starts to pick up just prior to 0000 hr Greenwich Mean Time (GMT; 0900 hr Tokyo time).[3] There is very little trading activity during the Asia–Pacific lunchtime break, around 0300 to 0430 hr GMT. Quote arrivals pick up again after lunch and continue until Europe (Zurich, Frankfurt, London) opens at around 0700 hr GMT. New York starts at approximately 1200 hr GMT (0800 Eastern Standard Time [EST]), and trading increases together with the other (European) banking centers. Quote arrivals begin tapering off just before noon New York time, which is also when London begins to wind down at 1600 hr GMT.

The upshot of Figures 2 and 3 is to demonstrate the existence of powerful 24-hr seasonals in both the observed correlation and the liquidity structures of exchange rates. Indeed, Figures 2 and 3 have commonalities that suggest a plausible link between observed correlation and liquidity.

Having presented this evidence, we primarily aim to show in this article that the liquidity process directly affects the observed contemporaneous correlation structure. Indeed, the observed correlation can be best thought of as a conditional correlation with the conditioning variables being the directly observable liquidity variables. We invoke the transaction-based model of Scholes and Williams (1976) as our choice for the conditional moments and then proceed to examine the role of the conditioning variables in affecting the (unobservable) unconditional moments such as the contemporaneous correlations.[4]

A MODEL OF ILLIQUID TRADING

Since the official identification of the infrequent trading problem by Fama (1965) and Fisher (1966), there have been several theoretical academic articles attempting to represent the problem. Perhaps the best known of these was written by Scholes and Williams (1977), who analyzed the ef-

[3]Official trading hours in Japan are 0900 to 1200 hr and 1430 to 1530 hr Tokyo time, although actual trading activity is widely acknowledged to start before and end after the official period. This will soon change to accommodate increased trading hours.

[4]The choice of model is actually somewhat arbitrary, but this article is not about the search for the definitive model.

fect of nonsynchronous trading on sample betas. They presented a formal model of infrequent trading based on random trading times. Scholes and Williams also prescribed instrumental variable-type estimators for beta estimates based on nonsynchronous data.[5] In contrast, Lo and MacKinlay (1990) employed the concept of nontrading probabilities to model infrequent trading, whereas Muthuswamy (1990) used an explicit concept of trading frequency. Also, Dimson (1979) and Miller, Muthuswamy, and Whaley (1994) employed time series methods to represent the effects of thin trading. All these approaches basically arrived at similar conclusions, albeit with rather different mathematical machinery.

The theoretical model of nonsynchronous trading that we adopt here is based on a relatively little known unpublished companion article to Scholes and Williams (1977), namely, Scholes and Williams (1976). In this precursor to their 1977 article, Scholes and Williams (1976) developed detailed algebraic expressions for observed moments of nonsynchronous returns as functions of true moments and nonsynchronous trading parameters.

The next section contains a brief description and relevant results from the Scholes–Williams (Scholes & Williams, 1976) model. A less detailed summary was also presented in Atchinson, Butler, and Simonds (1987). Formal proofs are omitted on account of their complexity, and the reader is asked instead to consult the Scholes and Williams (1976, 1977) articles directly.[6]

The Scholes–Williams Nonsynchronous Trading Model

In their seminal 1977 article, Scholes and Williams examined how illiquidity-induced asynchronicities in the trading process can cause observed (sample) moments of true returns to differ from true moments. Consider Figure 4, which shows a diagrammatical representation of two assets that trade infrequently. Returns are observed over two trading sessions of arbitrary length. Last observed trades are denoted by the dark blobs, whereas all trades prior to the last have been suppressed during a trading session. The basic intuition behind Figure 4 is that as an asset's trading intensity increases, its last trade time (relative to some arbitrarily defined trading session) converges to the actual close of trading during that session. As a result of illiquid trading, observed returns are no longer

[5]See Cohen et al. (1979), Lo and MacKinlay (1990), Muthuswamy (1990), and Low, Muthuswamy, and Whaley (1995) for different expositions of the effects of nonsynchronicities on measured returns.
[6]Detailed proofs are available on request from the authors.

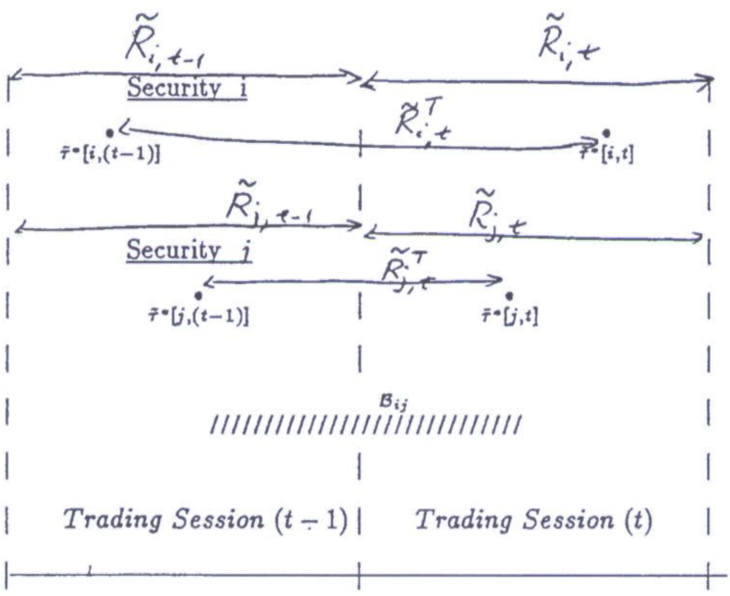

FIGURE 4
Representation of last trade times observed during consecutive trading sessions for two arbitrary securities.

definable over fixed session-length horizons but rather over stochastic intervals that depend critically on the nature of the quote (or transaction) arrival process. In some sense, the observed returns process can be said to be subordinated to the trading process. Observed moments can be shown to be related to true moments via some standard functions. In addition, observed covariances (and, consequently, observed correlations) are effectively measured only over common overlap intervals such as those depicted by the slashed lines (/////) in Figure 4.

The notation adopted is straightforward. Observed returns and moments are labeled with the transaction superscript T, whereas unobserved returns and moments are labeled without the superscript T. For instance, $\tilde{R}_{it}$ denotes asset i's true return during the session $(t - 1, t)$, whereas $\tilde{R}_{it}^{T}$ denotes the observed return during that session as a result of the illiquidity. Both $\tilde{R}_{it}$ and $\tilde{R}_{it}^{T}$ are labeled in Figure 4. Similarly, σ_{ij} denotes the true unconditional contemporaneous covariance between assets i and j, whereas σ_{ij}^{T} denotes the observed conditional covariance (based on a relevant conditional expectation with respect to $\tilde{R}_{it}^{T}$ and $\tilde{R}_{jt}^{T}$ instead of $\tilde{R}_{it}$ and $\tilde{R}_{jt}$).

Using λ_i and λ_j to denote quote (alternatively, transaction) arrival rates akin to that of Poisson distributions, Scholes and Williams (1976) derived a set of relations between the true and transactions-based moments:

$$\operatorname{Cov}(\tilde{R}_{it}^{T}, \tilde{R}_{jt}^{T}) = \sigma_{ij}^{T} = \gamma(\lambda_i, \lambda_j)\cdot\sigma_{ij} \tag{1}$$

and

$$\operatorname{Var}(\tilde{R}_{it}^{T}) = (\sigma_i^2)^T = \theta(\lambda_i, \nu_i)\cdot\sigma_i^2 \tag{2}$$

where $\alpha(\lambda_i,\lambda_j)$, $\gamma(\lambda_i,\lambda_j)$, $\theta(\lambda_i,\nu_i)$, and $\delta(\lambda_i,\nu_i)$ are functions of the nontrading parameters λ_i and λ_j, and the coefficients of variation ν_i and ν_j are defined by

$$\alpha(\lambda_i,\lambda_j) = \frac{\left[\dfrac{\lambda_j(1 - e^{-(\lambda_i+\lambda_j)})}{\lambda_i(\lambda_i + \lambda_j)} - e^{-(\lambda_i+\lambda_j)} - \left[1 + \dfrac{1}{\lambda_i} - \dfrac{1}{\lambda_j}\right]e^{-\lambda_j}(1 - e^{-\lambda_j})\right]}{(1 - e^{-\lambda_i})(1 - e^{-\lambda_j})}$$

$$\gamma(\lambda_i,\lambda_j) = 1 - \alpha(\lambda_i,\lambda_j) - \alpha(\lambda_j,\lambda_i)$$

$$\theta(\lambda_i,\nu_i) = 1 - 2\delta(\lambda_i,\nu_i)$$

$$\delta(\lambda_i,\nu_i) = \frac{-\left[\dfrac{1}{\lambda_i^2} + \dfrac{1}{1 - e^{-\lambda_i}} - \dfrac{1}{(1 - e^{-\lambda_i})^2}\right]}{\nu_i^2}$$

and

$$\nu_i = \frac{E(\tilde{R}_{it})}{\sqrt{\operatorname{Var}(\tilde{R}_{it})}} = \frac{\mu_i}{\sigma_i}$$

From the previous equations, it is straightforward to show that for observed returns R_{it}^T and R_{jt}^T, the observed (transactions-based) correlation ρ_{ij}^T is related to the true correlation ρ_{ij} by

$$\begin{aligned}\rho_{ij}^T &= \frac{\sigma_{ij}^T}{\sigma_i^T\cdot\sigma_j^T} = \frac{\gamma(\lambda_i,\nu_i)\sigma_{ij}}{\sqrt{\theta(\lambda_i,\nu_i)\sigma_i^2\cdot\theta(\lambda_j,\nu_j)\sigma_j^2}}\\ &= \frac{1 - \alpha(\lambda_i,\lambda_j) - \alpha(\lambda_j,\lambda_i)}{\sqrt{[1 - 2\delta(\lambda_i,\nu_i)][1 - 2\delta(\lambda_j,\nu_j)]}}\,\rho_{ij} \end{aligned}\tag{3}$$

or, equivalently,

$$\rho_{ij}^T = B(\lambda_i,\lambda_j,\nu_i,\nu_j)\cdot\rho_{ij}$$

where

$$0 \leq B(\lambda_i,\lambda_j,v_i,v_j) = \frac{1 - \alpha(\lambda_i,\lambda_j) - \alpha(\lambda_j,\lambda_i)}{\sqrt{[1 - 2\delta(\lambda_i,v_i)][1 - 2\delta(\lambda_j,v_j)]}} \leq 1 \qquad (4)$$

is a bias function attributed to the observed correlation as a result of the nonsynchronous trading parameters λ_i and λ_j and their respective coefficients of variation v_i, v_j. By definition, the bias is restricted to lie between 0 and 1. A bias of 1 indicates perfect neutrality between the observed and true correlations, whereas a bias near 0 indicates that the observed correlation understates the true correlation by a multiplicative inverse of the bias. For added insight, Figure 5 shows a plot of the correlation bias function $B(\lambda_i, \lambda_j, v_i, v_j)$ for various configurations of parameter values. It is clear from this figure that as the trading parameters λ_i and λ_j increase[7] indefinitely, the bias goes to unity, regardless of the values of v_i and v_j. This is intuitive because in the limit of continuous trading, there is full synchronicity, so observed correlations based on transactions data provide a good approximation for true correlations. With v_i and v_j kept at plausible levels, Figure 5 also shows that the worst biases, values of $B(\lambda_i, \lambda_j, v_i, v_j)$ close to zero, occur when one of the trading parameters is large and the other is small. In this case, observed covariances are likely to significantly understate true covariances.[8] Indeed, for small values of $\lambda_i(\lambda_j)$, large values of $\lambda_j(\lambda_i)$ can all but annihilate the true correlation. The correlation bias is also significant whenever both λ_i and λ_j are small, as it should be. The effects of v_i and v_j are somewhat second order in nature because no direct notion of synchronicity is involved.[9]

Having presented the Scholes–Williams (Scholes & Williams, 1976) model, we now examine the empirical links between observed and true correlations and consider the question of how to estimate the true (unobservable) correlations with minimal measurement error.

Rectification of the Mismeasured Observed Correlation

Consider again the relation between true correlations ρ_{ij} and observed correlations ρ_{ij}^T (based on the quotation data) described earlier in Equa-

[7]Increasing (decreasing) a trading parameter such as λ_i is tantamount to increasing (decreasing) the differencing interval. This is simply because there are more transactions occurring if the sampling interval increases. The Scholes–Williams (Scholes & Williams, 1976) model, therefore, allows for a very convenient mapping between the trading interval and the intensity of trades in that interval, effectively a linear scaling of λ_i.

[8]Scholes and Williams (1976, 1977) were the first to document these effects.

[9]Lambda is not allowed to be equal to zero because doing so might lead to nondefined solutions under this formulation. This is like imposing a minimum of one-trade-per-session restriction.

$$B(\lambda_i, \lambda_j, v_i, v_j) = \frac{\gamma(\lambda_i, \lambda_j)}{\sqrt{[1 - 2\delta(\lambda_i, v_i)][1 - 2\delta(\lambda_j, v_j)]}}$$

where

$$\gamma(\lambda_i, \lambda_j) = 1 - \alpha(\lambda_i, \lambda_j) - \alpha(\lambda_j, \lambda_i)$$

$$\alpha(\lambda_i, \lambda_j) = \frac{\left[\frac{\lambda_j(1-e^{-(\lambda_i+\lambda_j)})}{\lambda_i(\lambda_i+\lambda_j)} - \left[1 + \frac{1}{\lambda_i} - \frac{1}{\lambda_j}\right] e^{-\lambda_i}(1 - e^{-\lambda_j}) - e^{-(\lambda_i+\lambda_j)}\right]}{\left[(1 - e^{-\lambda_i})(1 - e^{-\lambda_j})\right]}$$

$$v_i = \frac{\mu_i}{\sigma_i} \qquad and \qquad v_j = \frac{\mu_j}{\sigma_j}$$

$$\delta(\lambda_i, v_i) = \frac{-\left[\frac{1}{\lambda_i^2} + \frac{1}{1-e^{-\lambda_i}} - (1 - e^{-\lambda_i})^{-2}\right]}{v_i^2}$$

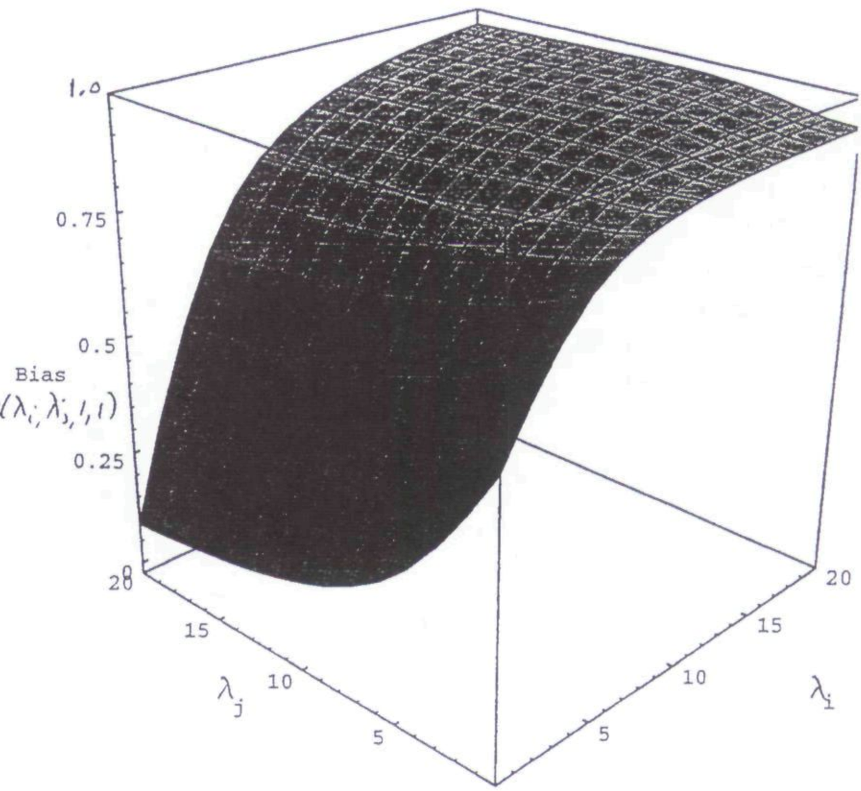

FIGURE 5

Bias function $B(\lambda_i, \lambda_j, \nu_i, \nu_j)$ for the theoretical correlation coefficient ρ_{ij}^{I} between the observed returns R_i^T and R_j^T for two stocks i and j when both stocks are measured with Poisson trading intensities λ_i and λ_j and true coefficients of variation ν_i and ν_j.

tion 3. For a sample of N observations based on the returns R_{it}^T, R_{jt}^T ($t = 1,2,3, \ldots, N$) of traded assets i and j, the standard moments estimator for the contemporaneous correlation is

$$\hat{\rho}_{ij}^T = \frac{\sum_{t=1}^{N} (R_{it}^T - \bar{R}_i^T)(R_{jt}^T - \bar{R}_j^T)}{\sqrt{\sum_{i=1}^{N} (R_{it}^T - \bar{R}_i^T)^2 \cdot \sum_{t=1}^{N} (R_{jt}^T - \bar{R}_j^T)^2}} \tag{5}$$

with the following probability limit (assuming that the Scholes–Williams model represents the illiquidity process correctly):

$$\operatorname*{plim}_{N \to \infty} \rho_{ij}^T = B(\lambda_i, \lambda_j, \nu_i, \nu_j)\rho_{ij}$$

The usual moments estimator $\hat{\rho}_{ij}^T$ is biased downward according to Equation 4, the bias stemming from the nonoverlap band (Figure 4) that induces a certain amount of nonsynchronicity-related leakage of the true covariational structure. An unbiased estimator may be obtained through rectification of the standard estimator by the bias. A rectified and unbiased estimator $\hat{\rho}_{ij}^R$ is given by

$$\begin{aligned} \hat{\rho}_{ij}^R &= \frac{\hat{\rho}_{ij}^T}{\hat{B}(\lambda_i, \lambda_j, \nu_i, \nu_j)} \\ &= \hat{\rho}_{ij}^T \frac{\sqrt{[1 - 2\delta(\hat{\lambda}_i, \hat{\nu}_i)][1 - 2\delta(\hat{\lambda}_j, \hat{\nu}_j)]}}{1 - \alpha(\hat{\lambda}_i, \hat{\lambda}_j) - \alpha(\hat{\lambda}_j, \hat{\lambda}_i)} \end{aligned} \tag{6}$$

Note that $\hat{\alpha}(\lambda_i, \lambda_j)$ and $\hat{\delta}(\lambda_i, \nu_i)$ are estimated by the moment estimator functions $\alpha(\lambda_i, \lambda_j)$ and $\delta(\lambda_i, \nu_i)$, with the parameters replaced by their estimates, where

$$\hat{\nu}_i = \frac{\hat{\mu}_i^T}{\hat{\sigma}_i^T} = \frac{\sum_{t=1}^{N} (R_{it}^T - \bar{R}_i^T)}{\sum_{t=1}^{T} (R_{it}^T - \bar{R}_i^T)^2}$$

and

$$\bar{R}_i^T = \frac{\sum_{t=1}^{N} R_{it}^T}{N}$$

In the previous equation, $\hat{\lambda}_i$ and $\hat{\lambda}_j$ are estimates based on the sample means of the quote arrival rates for the particular differencing intervals being examined.

Consider a high-frequency data set that has been partitioned into N subperiods, where each subperiod is the smallest session that is convenient for observation. Denote the total number of quote arrivals for the tth subperiod ($t = 1,2,3, \ldots ,N$) by V_t. Then, the mean quote arrival rate λ_i for this smallest interval is well estimated by its sample mean:

$$\hat{\lambda}_i = \frac{\sum_{i=1}^{N} V_i}{N}$$

If the differencing interval is reduced, for example, to half the previous value, the new sample size is simply doubled to $2N$ observations and the associated quote arrival intensity is simply adjusted[10] to be $\hat{\lambda}^2$, where

$$\hat{\lambda}^2 = \frac{\sum_{i=1}^{T} V_i}{2T}$$

In general, if the differencing interval is a multiple (e.g., n) of the smallest interval, then the associated quote arrival rate is given by

$$\hat{\lambda}^n = n \cdot \hat{\lambda}_i$$

For the rest of this article, we omit the superscript n for the differencing interval, and it can be assumed that the relevant quote arrival intensities λ_i have already been adjusted for the appropriate differencing interval. The rectified estimator is then given by

$$\hat{\rho}^{R_{ij}} = \frac{\left[\dfrac{\sum_{t=1}^{N} (R_{it}^T - \bar{R}_i^T)(R_{jt}^T - \bar{R}_j^T)}{\sqrt{\left\{ \sum_{i=i}^{N} (R_{it}^T - \bar{R}_i^T)^2 \sum_{t=1}^{N} (R_{jt}^T - \bar{R}_j^T)^2 \right\}}} \right]}{B(\hat{\lambda}_i, \hat{\lambda}_j, \hat{\nu}_i, \hat{\nu}_j)} \tag{7}$$

Equation 7 is the rectified estimator that we use from now on. It is per-

[10]This assumes that the distribution of trades is uniform over some interval. To accommodate more complex trading patterns, such as that documented by Wood, McInish, and Ord (1985), we need to fine-tune the λ_is to the time variation present. We leave this aspect for future research and concentrate for the moment, for the sake of parsimony, on a simpler backdrop for the trading process.

haps important to state that the rectification of ρ_{ij}^{T} is only partial. Part of the reason for this is that Equation 7 suffers from the problem of not being necessarily positive definite; rectified correlations can exceed one in absolute value when sample sizes are not large enough. Korkie (1989) described this somewhat technical issue in great detail.

DATA AND FURTHER TESTS

The currency exchange-rate data for this study were obtained from the high-frequency database supplied by Olsen & Associates (O & A) of Zurich, Switzerland.[11] The exchange-rate data consists of real-time (GMT) stamped (to the nearest second) bid–ask quotes for three major exchange rates, namely, the US$–DM, Japanese yen–US$, and the cross Japanese yen–DM rate. However, we used only the first two exchange rates for the main part of this study. Although observations for the 1-year period from October 1, 1992 to September 30, 1993 included weekends, we deleted weekend quotes on account of the very thin quote arrivals. Returns were computed based on a basic 5-min sampling interval. The use of 5-min returns yielded 73,728 observations for each price after the deletion of weekends and public holidays, a very sizable data set indeed.

Although the use of quote data alleviates the liquidity problem to some extent, we also mention in passing that quote data are not necessarily the perfect solution and may even pose new problems. For instance, a problem can arise if there are many trades taking place within the bid–ask quotes, but these quotes do not change after the trades have taken place. In this situation, quote returns will suffer a positive serial correlation bias. However, this is more of a problem for stock data. The advantage in currency data lies in the greater intensity of quotes and trade activity compared with equity markets.

Before proceeding to implement the Scholes–Williams (Scholes & Williams, 1976) model of infrequent trading, we considered the general issue of trading liquidity and its temporal variations for the case of our exchange-rate data. Recall that Figure 3 is a 24-hr plot of 5-min × 5-min average quotation arrival rates (estimated values of the λ_is in the Scholes–Williams infrequent trading model) for the US$–DM and US$–yen exchange rates over the sample period. We also interpreted the nature of the quote arrival process in terms of the trading activity in the various financial centers of the world. There was, therefore, a legitimate reason

[11]O & A is a research institute that provides real-time forecasts, historical analysis, and trading advice for various foreign exchange rates.

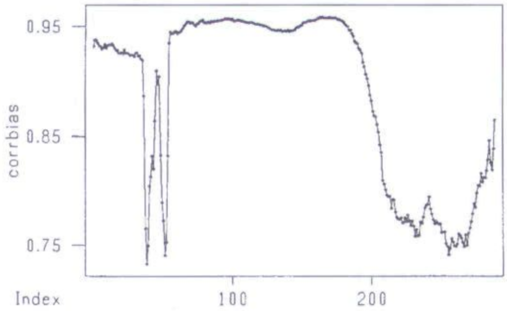

FIGURE 6

Empirical correlation bias function between returns for the US\$–DM and US\$–yen exchange rates. The plot shows the empirical bias function $B(\lambda_i, \lambda_j, v_i, v_j)$ for the Scholes–Williams (1976) infrequent trading model. Actual sample values for the parameters are used to compute the correlation bias, and the results are plotted on a 24-hr clock with 5-min returns for the US\$–DM and US\$–yen exchange rates.

to expect the quote arrival (liquidity) process to behave in accordance with its strong observed 24-hr seasonality.

Next, we estimated quote arrival rates for the US\$–DM and US\$–yen exchange rates and their respective coefficients of variation (i.e., λ_i, λ_j, v_i, and v_j). We then computed the bias functions for the observed correlations based on Equation 4 and plotted them on a 5-min $\times$ 5-min basis over a 24-hr clock in Figure 6.

Figure 6 is extremely interesting on several counts. Although it represents a theoretical model of the bias function and is based on a smooth mathematical function as such, the input parameters are anything but smooth, being based on sample data, which Figures 2 and 3 show contain very significant time varying patterns. Despite this smoothness property, Figure 6 indicates that this empirical correlation bias function exhibits considerable time variation, and in particular, Figure 6 is highly similar to Figure 3; indeed, the correlation between the 5-min quote arrival rate λ_i for the US\$–DM rate and the empirical correlation bias function is extremely high at .901. This has the implication that the empirical bias function is able to track and account for some of the liquidity variation, despite the fact that it is a purely mathematical function.

Finally, the empirical function plotted in Figure 6 was applied to rectify the observed correlations of Figure 2, both being for the case of 5-min returns over a 24-hr clock. Over the 288 5-min slices of the 24-hr clock, the average of the empirical correlation bias function was found to be .888 with a minimum of .733 and a maximum of .959. Therefore, on average, the observed correlations in Figure 2 can understate the true correlations by as much as 36%, this arising only from liquidity considerations.

SUMMARY AND CONCLUSIONS

The global foreign exchange markets straddle many time zones and include several financial centers. Trading liquidity varies from minute to minute according to which financial center is predominant at that point in time.

For corporate treasurers and foreign exchange arbitrageurs, the successful application of certain intraday dynamic hedging strategies or the accurate pricing of certain exotic currency options requires a constant evaluation of the correlation between the rates of return of the associated currency values.

In this study, we used high-frequency quotation-based currency data to first examine the time varying nature of observed contemporaneous correlations between US$–DM and US$–yen exchange rates. We then proceeded to examine the quote arrival process for these currencies. The observed correlations were found to vary somewhat with the quote arrival process over a 24-hr clock. Arguing that observed correlations are highly dependent on the liquidity structure of the traded assets, we invoked a liquidity-based model derived from Scholes and Williams (1976). Use of this model allowed for an explicit specification of the observed conditional correlation as a joint multiplicative function of the true unconditional correlation and a liquidity bias function. The mathematical correlation bias function was studied along with its estimated empirical counterpart. It was shown that the correlation bias was directly linked to the liquidity process.

Although the particular liquidity model adopted in this study was somewhat arbitrary, the main point brought forth is that some form of correction factor should be maintained between true and observed correlations. It remains for future research to deal with the more technical and econometric issue of how best to gauge more fully equilibrium values of these correlations and to separate the true time variation from the more spuriously induced variation engendered by liquidity considerations. Much work, therefore, remains on the research agenda on this score.

BIBLIOGRAPHY

Atchison, M., Butler, K., & Simonds, R. (1987). Nonsynchronous security trading and market index autocorrelation. Journal of Finance, 42, 111–118.

Cohen, K., Maier, G., Schwartz, S., & Whitcomb, D. (1978). The returns generation process, returns variance, and the effect of thinness in securities markets. Journal of Finance, 33, 149–167.

Dacarogna, M., et al. (1993). A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. Journal of International Money and Finance, 12, 413–438.

Dimson, E. (1979). Risk measurement when shares are subject to infrequent trading. Journal of Financial Economics, 7, 197–226.

Epps, T. W. (1979). Comovements in stock prices in the very short run. Journal of the American Statistical Association, 74, 291–298.

Fama, E. (1965). Tomorrow on the New York Stock Exchange. Journal of Business, 38, 285–299.

Fisher, L. (1966). Some new stock market indices. Journal of Business, 39, 191–225.

Korkie, B. (1989). Corrections for trading frictions in multivariate returns. Journal of Finance, 44, 1421–1434.

Lo, A., & MacKinlay, C. (1990). An econometric analysis of infrequent trading. Journal of Econometrics, 45, 181–211.

Low, A., & Muthuswamy, J. (1996). Information flows in high frequency exchange rates. In C. Dunnis (Ed.), Forecasting in financial markets. New York: Wiley.

Low, A., Muthuswamy, J., & Whaley, R. (1994). The effects of infrequent trading on the time-series structure of portfolio excess returns. Manuscript in preparation.

Miller, M., Muthuswamy, J., & Whaley, R. (1994). Mean reversion of Standard & Poor's index basis changes: Arbitrage-induced or statistical illusion? Journal of Finance, 49, 479–513.

Muller, U., et al. (1990). Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. Journal of Banking and Finance, 14, 1189–1208.

Muthuswamy, J. (1990). Nonsynchronous trading and the index autocorrelation problem. Unpublished doctoral dissertation, University of Chicago, Graduate School of Business.

Pearson, N. (1995). An efficient approach for pricing spread options. Journal of Derivatives, 3, 76–91.

Rubinstein, M. (1990). Exchange options. Unpublished manuscript, University of California, Berkeley.

Scholes, M., & Williams, J. (1976). Estimating betas from daily data. Unpublished manuscript, University of Chicago, Graduate School of Business.

Scholes, M., & Williams, J. (1977). Estimating betas from non-synchronous data. Journal of Financial Economics, 5, 309–327.

Stulz, R. (1982). Options on the minimum or the maximum of two risky assets: Analysis and applications. Journal of Financial Economics, 10, 161–185.

Wood, R., McInish, T., & Ord, J. (1985). An investigation of transactions data for NYSE stocks. Journal of Finance, 40, 723–741.