International Review of Business Research Papers
Vol. 3 No.5 November 2007 Pp.88-122

Speed of Adjustment within Currency Markets

Terry Boulter*

One measure of market efficiency is the speed at which prices adjust to fundamental value with the arrival of information. This paper examines this issue by estimating speed of adjustment coefficients using three methodologies for eight currencies for the entire year of 1996 using half hourly non-overlapping return intervals. We find that the bulk of adjustment to fundamental value for all currencies occurs within the hour but then quickly deteriorates. Within the hour adjustment is sufficiently quick to be considered efficient but the lack of full adjustment to fundamental value is not what would be predicted within an efficient market. There is no evidence for any of the currencies studied of a tendency to over react. There is also little difference in the speeds of adjustment between actively and less actively traded currencies. There is however a definite difference in the speed at which currencies adjustment depending on whether they are free floating or managed exchange rates. Free floating rates adjust much quicker. Government intervention slows adjustment to fundamental value.

JEL Classification: F31; C32

Keywords: Market Efficiency, Exchange Rates, Speed of Adjustment, Non-synchronous Trading, Non-overlapping Returns

1. Introduction

The theory of efficient financial markets predicts that prices adjust “quickly” with the arrival of new information. Theory does not however define “quick” nor provide a methodology for measuring the speed at which this occurs. Being able to measure the speed at which exchange rates adjust means that currencies can be examined by comparing the speed at which they move to fundamental value. Comparisons can also be made between actively traded currencies and those that trade less actively or between currencies that are free-floating against those that are managed or between currencies of different liquidity. This paper makes these comparisons.

*Terry Boulter,Senior Lecturer: Finance,Deakin Business School,Faculty of Business and Law,Deakin University,336 Glenferrie Road,Malvern, 3156,
E-mail: Terry.Boulter@Deakin.edu.au

There are different ways to examine the speed at which currencies adjust. A common approach has been to use event study methodology so as to observe price dynamics around the time of a scheduled macroeconomic news announcement. A drawback of event studies is that speed can only be examined at the time of a predefined known event. Event studies are unable to confirm if the speed of adjustment that occurs at the time of an event is consistent with the speed at which prices adjust to all other information. It is possible that the market responds differently to expected announcements than it does to the random arrival of new information.

The purpose this paper is to estimate speed of adjustment within a continuous time series setting using high frequency data so as to gauge to what degree prices adjust, and to compare speeds of adjustment between currencies. This is done using a set of eight currencies observed at half hourly intervals over the entire year of 1996. The contribution of this paper is to measure speed of adjustment within currency markets using three simple methodologies which have not previously been used to examine adjustment within exchange rate markets. These methodologies are based on the description price innovation as explained by Amihud and Mendelson (1987). Prices are thought to evolve noisily and partially to fundamental value.

The structure of the paper is as follows. Section 2 discusses the theoretical underpinning of the empirical procedures used to generate speed of adjustment estimates. Section 3 discusses a set of practical issues that need resolution prior to estimating speed of adjustment coefficients. Section 4 identifies the data, its frequency, and its statistical properties. Section 4 also examines the possible impact that non-synchronous trading may have on speed of adjustment estimates. Section 5 generates speed of adjustment coefficients for eight currencies. These estimates are then used to make comparisons between currencies. It is found that the bulk of adjustment for all eight currencies occurs within the hour. Results also indicate that currencies do not adjust completely to fundamental value.

2. Measuring Speed of Adjustment

This section identifies a number of methodologies that have been designed to measure the speed of price adjustment process. These measures have traditionally been used to estimate the speed of adjustment within equities markets. They are adapted in this study to estimate speed of adjustment within currency markets. The idea that the speed of price adjustment is something that can be measured is found within the work of Amihud and Mendelson (1987). They argue that prices adjust partially and noisily towards fundamental value such that:

$$R_t = P_t - P_{t-1} = g(V_t - P_{t-1}) + u_t \quad (1)$$

$$V_t - V_{t-1} = \mu + \varepsilon_t \quad (2)$$

where (V_t) and (P_t) are in logarithms. ($P_t - P_{t-1}$) is the observed price change, ($V_t - P_{t-1}$) represents the information-induced change in price (i.e. noise) which adjusts by (g). The adjustment coefficient is assumed to be stationary[1] and satisfies the condition ($0 < g < 2$). Within this framework, it is assumed that (u_t) is a white noise sequence of i.i.d. random variables with a mean of zero and finite variance.[2] ($V_t - V_{t-1}$) is the change in logarithmic intrinsic values, (μ) the mean of the intrinsic value random walk process and (ε_t) the innovations in logarithmic intrinsic values which must be uncorrelated if the market is efficient.[3]

The Amihud and Mendelson (1987) partial adjustment model of price innovation does not, however, explain how to generate the estimate for (g). It has been left to other authors to develop methodologies to estimate the speed of adjustment coefficient. An initial procedure was developed by Damodaran (1993). His estimator is based on the decomposition of observed return variance into three distinct components:

$$Var(R_t) = v^2 + 2\sigma^2 + \left[(\frac{g}{2-g} - 1)v^2 + (\frac{2}{2-g} - 2)\sigma^2 \right] \quad (3)$$

Within this framework, (v^2) is the variance of the intrinsic value, ($2\sigma^2$) is the noise variance, and $\left[(\frac{g}{2-g} - 1)v^2 + (\frac{2}{2-g} - 2)\sigma^2 \right]$ is the variance associated with the price adjustment effect.

The estimate of (g) is derived as a function of observed return variances that are constructed over different time intervals. For example, (R_{jt}) would be the return in time period (t) where each return interval is of length (j). The variance of these returns is:

$$Var(R_{jt}) = \left[\frac{g_j}{2-g_j} jv^2 + \frac{2}{(2-g_j)}\sigma^2 \right] \quad (4)$$

Assuming that a number of observed return unit-intervals are appropriate to calculate the variances, such that (*j* = 1, 2, ...*k*), and assuming that (*k*) is of sufficient length to allow (g) to equal one, then the variance in (*j*) interval returns and the variance in (*k*) interval returns can be calculated as:

$$Var(R_{jt}) - \frac{Var(R_{kt})}{k} = v^2\left[\frac{g_j}{(2-g_j)} - 1 \right] + 2\sigma^2\left[\frac{1}{(2-g_j)} - \frac{1}{k} \right] \quad (5)$$

---

[1] The implications of which are examined by Theobold and Yallup (1998).

[2] The size of u_t is believed to be a function of information-related factors (such as noise trading and trading based solely on liquidity) and market-related factors (such as the bid–ask spread, dealer inventory positions etc.).

[3] As discussed by Theobald and Yallup (2004), the intrinsic value is assumed to follow a random walk process, i.e. intrinsic values fully (efficiently) adjust to information shocks.

where the variance of the intrinsic value (v^2) and the noise variance (σ^2) are functions of the covariance and variance in (*k*) intervals such that:

$$v^2 = \frac{Var(R_{kt}) + 2Cov(R_{kt,}R_{kt-1})}{k} \quad (6)$$

$$\sigma^2 = -Cov(R_{kt}, R_{kt-1}) \quad (7)$$

Substituting (v^2) and (σ^2), in equation (6.4), and solving for (g), Damodaran calculates the speed of adjustment to be:

$$g_j = \frac{2\left[\frac{Var(R_{jt})}{j} + \frac{Var(R_{kt})}{k}(j-1) + \frac{Cov(R_{kt,}R_{kt-1})}{j}\right]}{\frac{Var(R_{jt})}{j} + \frac{Var(R_{kt})}{k}(2j-1) + \frac{2Cov(R_{kt,}R_{kt-1})}{k}} \quad (8)$$

Damodaran proposes that within efficient markets, (g) should be asymptotic to increasing intervals in time. In other words, the more efficient the market, the quicker (g) should converge to (1). Figure 6.1 is a graphical representation of Damodaran's view of the asymptotic nature of the price-adjustment process believed to occur within efficient financial markets.

Figure 6.1: The Assumed Asymptotic Nature of Price Adjustment

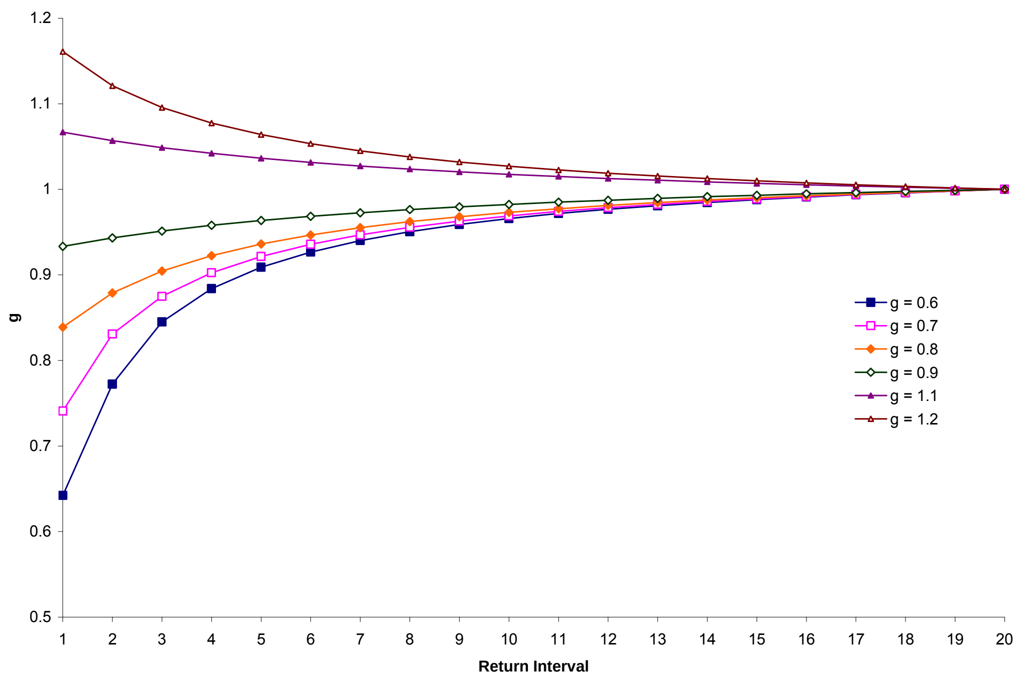

The implication is that the time it takes for (g) to converge on one (1) and that this is a good estimate of the "speed" at which prices adjust. Intuitively, the more efficient the market, the quicker (g) converges on one (1). In the foreign exchange market where traders are able to respond to information very quickly, twenty-four hours a day, it is likely adjustment takes place within minutes of information arrival. Within this framework, if $(0 < g < 1)$ within the first half hour then this would be interpreted as only a portion of the full value of the intrinsic value variance (v^2) being transmitted to observed returns. When ($g = 1$), the full extent of intrinsic variance would have been passed to observed returns, and if ($g > 1$) then this would be a case where intrinsic value variance is accentuated by over-reaction.

Damodaran uses his estimator of (g) to measure the speed of adjustment of American stocks using daily data for the period 1977 to 1986. His intention is to determine if there is evidence of lagged price adjustment to news arrival on an intra-day basis, and to establish whether there are differences between markets in the speed at which prices adjust to the arrival of information. He finds there is a lagged price adjustment within American securities markets and that there are differences in the speed of adjustment between stocks listed on the NYSE/AMEX and stocks listed on the NASDAQ. He concludes that the over-the-counter market reacts slower than markets that trade electronically.

Subsequent to Damodaran (1993), Brisley and Theobald (1996) confirm that the (g) estimator is a practical speed of adjustment estimation procedure which can be used to test market efficiency and from which comparisons between markets can be made. However, Brisley and Theobald (1996) uncovered an error within Damodaran's original formulation (refer to equation (8)) which they subsequently correct. They reformulate the estimator as:

$$g_j = \frac{2\left[\frac{Var(R_{j,t})}{j} + \frac{2Cov(R_{k,t}),R(_{k,t-1})}{j}\right]}{\frac{Var(R_{j,t})}{j} + \frac{Var(R_{k,t})}{k} + \frac{2Cov(R_{k,t}R_{k,t-1})}{k}} \qquad (9)$$

The problem with the original Damodaran (1993) formulation is that it has a tendency to overstate the price reaction towards intrinsic value when adjustment is incomplete, particularly at shorter differencing intervals. This is because as the differencing interval approaches its limiting value, the error term within the numerator and denominator increasingly dominate the other terms in the equation, which forces the Damodaran price adjustment coefficient to one at an increasing rate. The Brisley and Theobald's (1996) estimator corrects for this bias.

A remaining issue with the Damodaran estimator is the procedure for setting the limiting interval (*k*). Ariff and Chan (2002) develop an objective procedure for setting the length of the limiting interval (*k*). Previous to this work (*k)* had been set arbitrarily.

Safvenblad (1997) has investigated the econometric properties of the Damodaran (1993) estimator and concludes that it is easy to estimate and is practical, but that its performance is biased with individual estimates exhibiting unacceptably low precision. Safvenblad (1997) also points out that (g) does not have a readily derived sampling distribution from which tests of parameter significance and constancy can be examined. This means that the researcher can never really be sure whether the estimated speed of adjustment coefficient is statistically different from zero.

Collectively, the weaknesses of the Damodaran estimator add up to substantial shortcomings. This is irrespective of the number of researchers, including Eilifsen *et al.* (2001), who continue to use the Damodaran (1993) procedure and report interesting results. It may therefore be worthwhile estimating speed of adjustment coefficients using the Damodaran (g) within this study and to compare estimates against other methodologies as a way of gaining further insight into the Damodaran (1993) estimator. It also makes it possible to gauge how divergent speeds are, depending on the methodology used, and to establish if rankings of currencies based on different methodologies fundamentally change with the estimator used. In order to achieve these objectives, an alternate speed of adjustment procedure needs to be identified.

Theobold and Yallup (2004) develop two alternate estimators which do not rely on a limiting interval. Both also have a sampling distribution from which tests of statistical significance and parameter constancy can be examined. The two estimators are built on the assumption of a lagged price response. A lagged response is believed to be caused by traders initially under or over-reacting to information arrival. A lagged response is consistent with the Amihud and Mendelson (1987) view of partial adjustment. It is also a stylised feature of price behaviour that is discussed at length within the exchange rate determination literature. Many models of exchange rate determination account for under or over-reaction as a central tenet of the price innovation process. The Dornbusch (1976) model, for example, includes overshooting due to “sticky” prices. Barberis *et al.* (1998) develops a model where investors under-react because of a “conservative bias”. The intuition of these models is that partial price adjustment is a function of under or over-reaction to the arrival of new information.

Partial adjustment could be an effect of non-synchronous (infrequent) trading. This is what is suggested by Lo and MacKinlay (1990) who believe that price adjustment is a function of the autocorrelation and cross-autocorrelation terms within returns which is induced by trading which is cyclical in nature due to of recurring themes in the level of trading activity. Indeed, if exchange rates under or over-react to information arrival, or if they are influenced by cyclical patterns, then the persistence in prices through time will cause positive or negative autocorrelations in the return series. Within the context of efficient market theory, serial dependency is interpreted as evidence of incomplete adjustment. When prices under-react, the series will be positively correlated and when a currency over-reacts the series will be negatively correlated. The time it takes intrinsic value to re-establish is the time it takes for the positive or negative autocorrelation within the observed return series to dissipate. The time it takes to dissipate can be interpreted as a measure of the speed at which prices adjust to the arrival of information.

Theobold and Yallup (2004) develop two unique yet related procedures for estimating the speed of price adjustment based on the assumption of persistent under and over-reaction. Specifically, they note that auto-covariances for lags one and two can be derived as:

$$Cov\{R_t, R_{t-1}\} = \frac{g}{2-g}\left[(1-g)v^2 - \sigma^2\right] \qquad (10)$$

(see Amihud and Mendelson, 1987 equation (3.6)) and lag two (-2) as:

$$Cov\{R_t, R_{t-2}\} = g\,\frac{(1-g)}{2-g}\left[(1-g)v^2 - \sigma^2\right] \qquad (11)$$

(see Theobold and Yallup, 2004 equation (3b)). If changes in returns and the noise process are both stationary stochastic processes and the cross-covariance between them are zero at all lags then the speed of adjustment coefficient (g) can expressed as a function of the auto-covariance structure. The speed of adjustment coefficient can be written as:

$$1-g = \frac{Cov\{R(t), R(t-2)\}}{Cov\{R(t), R(t-1)\}} \qquad (12)$$

Within this framework, (g) is simply a function of the auto-covariance structure within lagged returns. An additional benefit of this procedure is that the auto-covariance estimator is as easy or easier to estimate as the Damodaran (g) coefficient. An estimation problem inherent within this framework, however, is the effect non-synchronous trading may play. Non-synchronous trading occurs because of thin trading which happens when there are little or no changes in prices over extended periods of time. This causes runs of zeros within the return series. This persistence could be a source of autocorrelation that is distinct from traders under or over-reaction to information arrival.

The likelihood of thin trading within the foreign exchange market is a definite possibility. Even though currency markets are open twenty-four hours, the bulk of trading occurs when dominant markets such as London and New York are both open. When New York closes, after London and before Sydney reopens, trading activity slows significantly. This decrease in trading may cause changes within the quoted data that would generate runs of zeros within the return series. These runs of zeros become a potential source of serial correlation (see Miller *et al*., 1994). This would be distinct from the autocorrelation caused by trader under or over-reaction, and would need to be removed before estimating a speed of adjustment coefficient.

If the empiricist did not want to test nor remove the effects of non-synchronous trading, Theobald and Yallup (2004) allow for its existence by making adjustments to equation (13). This can be achieved in the following way:

$$1-g = \frac{Cov\{R(m,t), R(m,t-2-q)\}}{Cov\{R(m,t), R(m,t-1-q)\}} \qquad (13)$$

Where $R(m,t)$ are observed returns that are subject to the effects of non-synchronous trading and (q) is the longest lag in returns that impacts upon $R(m,t)$.

Theobald and Yallup (2004) also demonstrate that there is another estimator that can be formulated. They build this estimator on intuition that is within the auto-covariance estimator. After first differencing and rearranging, the Amihud and Mendelson (1987) partial adjustment model (1) can be rewritten as:

$$R_t = (1 - g)R_{t-1} + g\Delta V_t + \Delta u_t \quad (14)$$

by substituting for (ΔV_t) from equation (2), equation (13) becomes:

$$R_t = g\mu + (1 - g)R_{t-1} + g\varepsilon_t + u_t - u_{t-1} \quad (15)$$

Within this modelling framework, the autocorrelations induced by under or over-reactions are manifest within an ARMA (1,1) model where price adjustment effects occur within the AR (1) coefficient. When adjustment is complete (g =1), the process is simply a MA (1) process[4]. In this state innovations are purely a function of noise. The AR process will be stationary as long as $|1-g| < 1$ which is an original precondition identified by Amihud and Mendelson (1987) within their partial price adjustment model (see equation 1).[5]

The intuition of the AR (1) estimator is that the quicker the AR (1) coefficient disappears (the quicker (g) converges on 1), the quicker that adjustment occurs and thus the more efficient is the market.Unlike any other study undertaken within the currency markets, this study uses all three estimators. To the best of this author's knowledge, estimates of this sort have not appeared within the literature previously.

Before proceeding to the estimation of speed of adjustment coefficients it is necessary to examine a set of practical issues connected to the estimation procedures. These include the setting of the limiting interval *(k)*, the length of the holding period return and the effect of intervalling. These important issues are discussed in Section 3. It is also necessary to examine the statistical properties of the sample data. This is undertaken to ensure that return characteristics for individual currencies are not fundamentally different from one another. If this was the case, it would be difficult to make comparisons between currencies which is a central theme of the current study. It is also necessary to examine the potential extent of autocorrelation for each currency that could be caused by non-synchronous trading. This issue of non-synchronous trading (infrequent trading) is discussed in Section 4.

---

[4] Which is assumed to be a stationary process within an efficient market.

[5] The pre-condition ensures that the price innovation is not a process which is explosive.

3. Practical Considerations

Before proceeding to estimation, a number of practical issues need to be resolved. Each is related to the length of time used within each of the methodological procedures. These are decisions that the researcher must make when (i) setting the limiting interval (*k*); (ii) determining the length of the holding period return and (iii) whether the holding period return is calculated as overlapping or non-overlapping. Each of these issues is discussed in turn.

(i) The limiting interval *(k)*:

The Damodaran (1993) estimator requires that a limiting interval *(k)* be identified before estimation. Within the estimation procedure, it is assumed that full adjustment occurs within interval *(k).* Damodaran (1993) does not, however, identify the longest or shortest interval or a procedure for setting the optimal length of the interval. It is imperative that the interval be within the correct range.
Damodaran (1993) experiments using different intervals for *(k)* and settles on twenty days (20) when examining the equities market. He tests the sensitivity of his results to limiting intervals greater than twenty days and finds no improvement in the results beyond this time frame. He does not, however, identify the procedure adopted to select the twenty day *(k)* interval.
Ariff and Chan (2002) attempt to identify a procedure for setting *(k)* by eliminating intervals that are too long. They make the assumption that the inherent link between values of *(g)* and *(k)* means that a statistically significant relationship between them will exist up to the point where the limiting interval is no longer related to *(g)*. To find this interval, they incrementally test the relationship between *(g)* and *(k)* to the point where the relationship becomes statistically insignificant. They conclude that full adjustment occurs up to that point and then use that interval plus one (+1) as their best choice for *(k)*.Their incremental approach for estimating *(g)* becomes insensitive to the choice of the limiting interval *(k)* at nineteen days. Based on this finding, they set interval *(k)* to twenty days, which coincidently is the same limiting interval identified without clarification by Damodaran (1991; 1993).

The Ariff and Chan (2002) procedure is able to successfully identify the widest feasible interval but it is unable to identify the optimal range. Since the limiting interval determines the magnitude of the speed of adjustment coefficient and thus how much noise is within it, determining the correct interval remains a critically important issue.

When determining the limiting interval *(k)* within currency markets, the fact that these markets are open continuously means that this characteristic needs to be accounted for when setting the limiting interval. The twenty-four hour nature of currency markets means that there is greater likelihood that “speedy” adjustment will take place quicker than it does within equities markets which are open for a limited number of hours per day. For this reason, the limiting interval *(k)* for the current study has been set in terms of hours instead of days.

The current study sets the maximum limiting interval *(k)* at two and a half hours. This range has been identified based on two considerations. The first is a pragmatic reason based on the fact that the frequency of observations within the data set is half hourly. This precludes the opportunity of using observations at a higher frequency. The second is based on findings of other studies which have used event study methodology. These studies find that adjustment to predefined news events continues to take place outside ten minutes (see Ederington, L. and Lee, J. 1993). It is also the case that other estimators that do not require a limiting interval find that the optimal range is longer than fifteen minutes but within two and a half hours.

To best determine what the longest interval should be, other speeds of adjustment estimators are used. These are taken from the Theobold and Yallup (2004) procuedure. These estimators have the advantage of not having to use a limiting interval. They are, therefore, able to provide a good estimate of the optimal range that should be used for the Damodaran (1993) estimator. It also provides the additional advantage of allowing for comparisons generated from the different estimators in terms of the magnitudes of adjustment over the *(k)* interval.

The Theobald and Yallup (2004) estimators find that adjustment for all eight currencies begins to deteriorate after the one and a half hour time frame. This is an indication that full adjustment occurs within that time frame. So as to ensure that complete adjustment for the Damodaran estimator happens within the interval, *(k)* is set slightly longer to two and a half hours. A longer interval is needed to ensure full adjustment for the Damodaran (1993) estimator can take place within that interval.

(ii) The length of the holding period return:

The second issue requiring resolution is the length of the holding period return. When the frequency of observations within the data set is half hourly, the choice of holding period return can be hourly or greater. The decision as to which is optimal needs to be determined by the empiricist. When deciding on what to use, the researcher must weigh up the costs and benefits of using longer or shorter intervals. The longer the interval, the less impact noise will have on the final estimate. The cost, on the other hand, is that a longer return interval decreases sample size and thus weakens the power of the estimate. It is also the case that a longer return interval means that it is more difficult to capture short-term price dynamics. When making a decision based on this trade-off, the empiricist must take account of the market structure being observed and the research questions being addressed.

For example researchers studying capital markets often use daily data. They argue that a daily unit interval is appropriate because of the presence of a distorting amount of intra-day noise. This occurs for a number of reasons. In part, it is due to the fact that the market is open only for a limited number of hours each day and not for every day of the week. When the market is closed (which is more often than not), information continues to arrive which on reopening must be adjusted into price. This often causes trading to be hectic at the beginning of each day as the market responds to new information that arrived when the market was not open. After this initial hectic period the market often settles only to become hectic again at the end of the day as traders adjust their positions before the market closes. These opening and closing adjustments reflect a non constant flow of information that causes intra-day activity to be erratic and noisy. To minimise these disturbances, empiricists often use closing daily prices which reduces the distortions caused by intra-day erratic behaviour.

Using daily data to study currency markets however would, be too long a unit interval. These markets are open twenty-four hours a day, every day of the week. Irrespective information flow over the twenty four hour period is not constant. There are definite patterns throughout the day of high and low activity as markets around the globe open and close and throughout the week as the weekend approaches and ends. Regardless of the high and lows in activity, the fact that there is always another market open means traders have an opportunity to respond to new information at any time of any day or week. They never have to wait for the market to re-open. Given the speed at which traders have the ability to react and that they can do so twenty-four hours a day means that intra-day data needs to be used when examining adjustment within currency markets.
An intra-day interval based on half hourly returns is adopted within the current study. This level of frequency ensures that short-term price dynamics are captured. Half hourly frequency also minimises the impact of the loss of sample size as the holding period return interval is increased.

(iii) Overlapping versus non-overlapping returns:

The last issue of a practical nature that needs to be resolved is the impact that overlapping holding period returns have on the speed of adjustment estimate. Muller (1993) demonstrates that overlapping returns will cause serial dependence within the financial return time series. Serial dependence will bias speed of adjustment estimates. The impact that this will have on the Theobald and Yallup (2004) estimators is worth exploring in some detail.

Figure 2 is a plot of the log of a simulated return series of twenty thousand observations. The simulated series was generated within the Amihud and Mendelson (1987) partial adjustment framework using equations (1) and (2). The plot of the randomly generated series is a good depiction of a free floating exchange rate series.

Figure 2: Simulated Series of Log Price

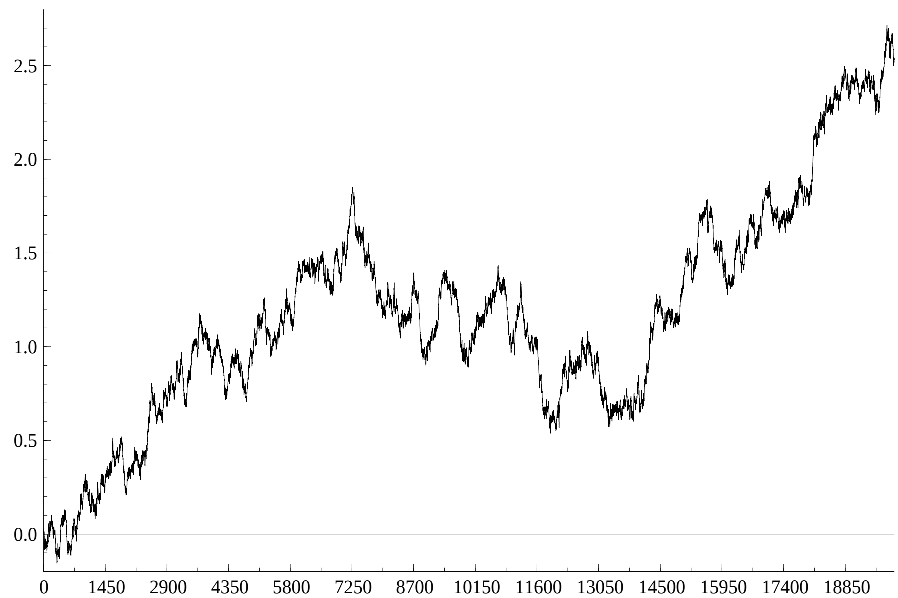

Using the simulated series it is possible to generate a set of five holding period overlapping returns. These return series can be used to estimate speed of adjustment coefficients using the Theobald and Yallup (2004) ARMA estimator. It is then possible to analyse how well the estimator performs.

Figure 3 is a plot of the autocorrelation and partial autocorrelation of the simulated series. Visually, it is evident that there is serial correlation within the overlapping returns.

Figure 3 Correlogram of Overlapping Returns

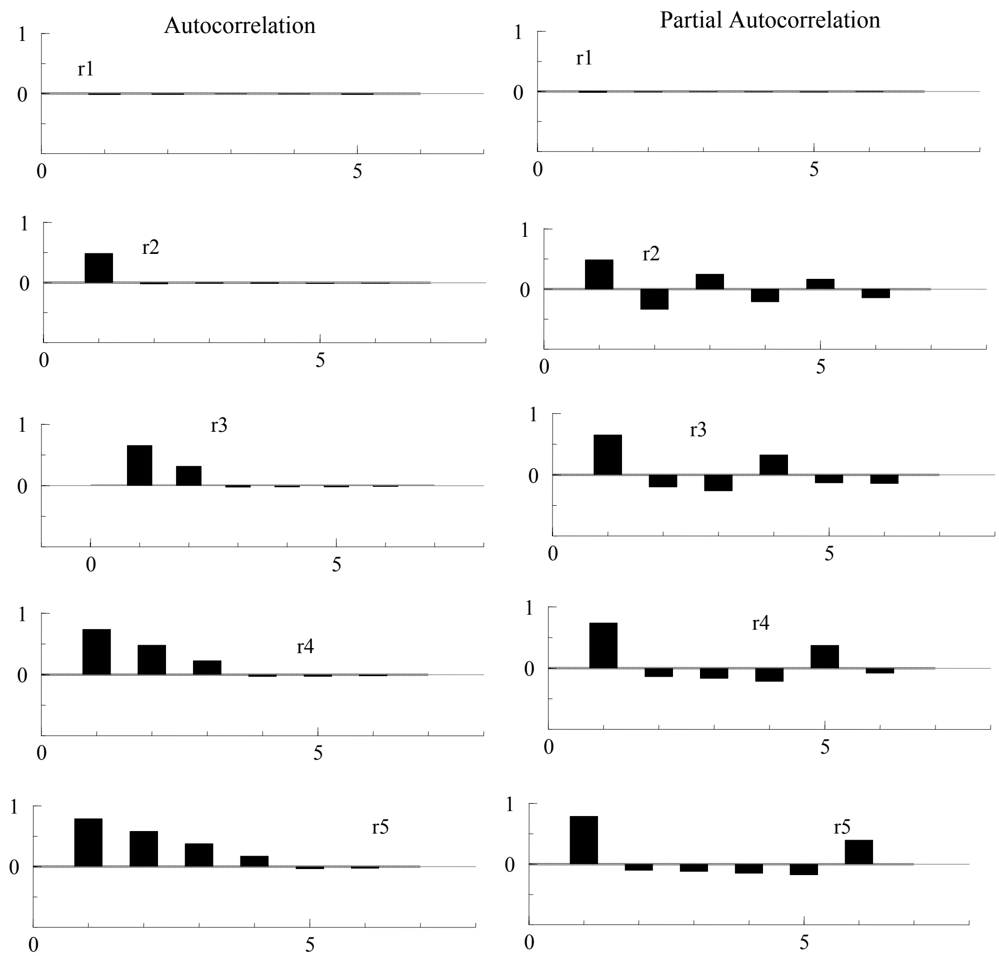

When an ARMA model is estimated (results shown in Table 1) which accounts for the MA structure, within Figure 6.3, the value of *(g)* does not deviate significantly from (.15) which is the initial setting for *(g)* used to generate the simulated data. As the interval over which the overlapping holding period return is calculated increases, the greater the order of the MA term that needs to be included within the ARMA estimator, but in each case a consistent estimator of the initial setting for the speed of adjustment parameter is returned. In other words, when the MA structure is modelled correctly, the AR coefficient simply reflects its initial setting; the use of overlapping data cannot offer any insight into the speed of adjustment. This is demonstrated in emphatic fashion in Table 1 which reports the estimates of *(g)* at different holding period return intervals.

Table 1: ARMA Simulated Results Using Overlapping Returns

	AR (1)	MA(1)	MA(2)	MA(3)	MA(4)	MA(5)	g
R1	.709** (.133723)	-.727** (.13022)					.29
R2	.821** (.082995)	.161* (.079875)	-.833** (.079943)				.18
R3	.860** (.064514)	.126** (.062031)	.127** (.061705)	-.868** (.061891)			.14
R4	.842** (.061496)	.144** (.059235)	.143** (.059274)	.144** (.058785)	-.851** (.059143)		.16
R5	.816** (.065083)	.168** (.062947)	.167** (.062953)	.167** (.062712)	.167** (.062364)	-.827** (.062609)	.18

**Denotes statistical significance of parameter estimate at 95% confidence interval. Standard errors are reported in brackets.

It is now worth examining the impact on *(g)* when non-overlapping returns are used to estimate the AR (1) parameter. Figure 4 plots the autocorrelation and partial autocorrelation for five holding period returns for the same data used to generate results reported in Figure 3 and Table 1. The only difference is that the holding period returns are calculated as non-overlapping returns. Figure 4 confirms that the autocorrelation and partial autocorrelation is removed at all return intervals.

Figure 4 Correlogram of Non-overlapping returns

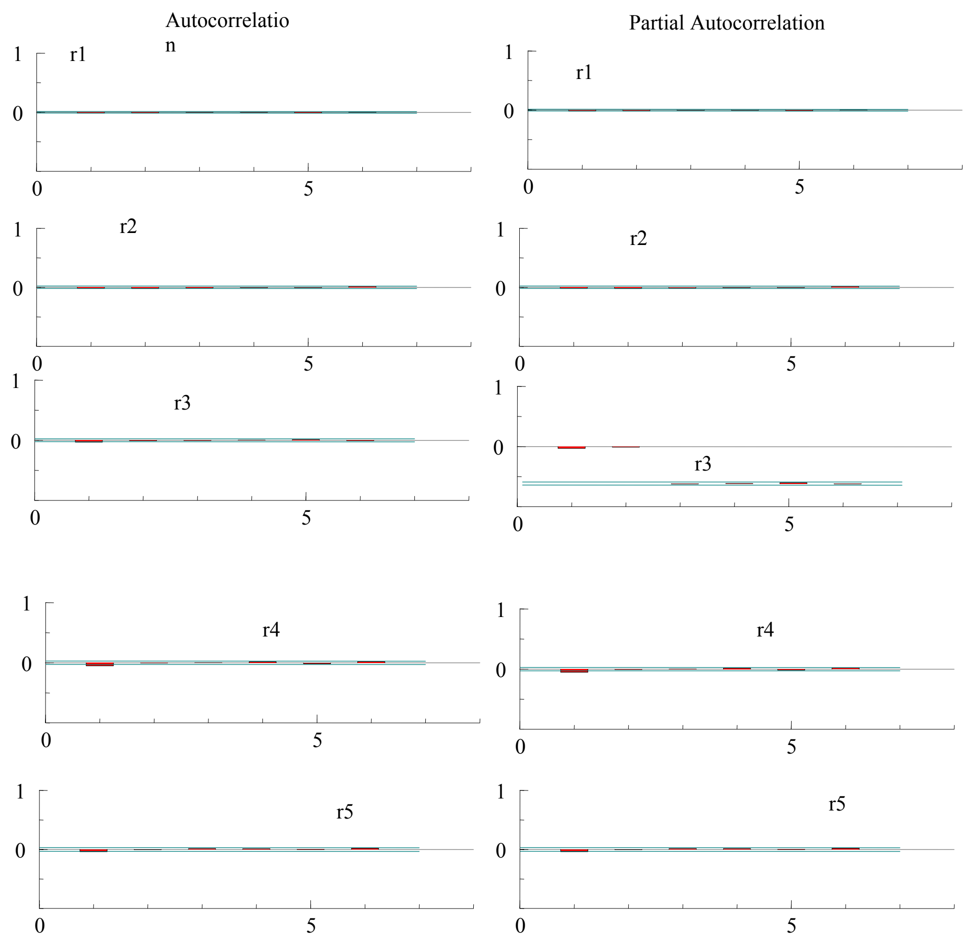

When non-overlapping simulated return series are used to generate estimates of speed of adjustment at different return intervals using the ARMA (1,1) estimator, adjustment proceeds incrementally towards full adjustment from the initial setting of *(g)*. The incremental adjustment is shown in Table 2 where *(g)* adjusts to .29 when returns are calculated as the first difference, and .44 (an increment of .15) when (r2) is calculated, and this incremental adjustment continues until full adjustment is reached at return interval five (r5). Convergence is as it should be given the simulated data was randomly generated using a starting value for (g) of (.15).

Table 2: ARMA Simulated Results using Non-overlapping Returns

	AR (1)	MA(1)	g
R1	.709** (.133644)	-.727** (.130142)	.29
R2	.565** (.204293)	-.592** (.199607)	.44
R3	.209 (.329696)	-.242 (.327105)	.71
R4	.017 (.268963)	-.068 (.268385)	.98
R5	.005 (.387049)	-.044 (.386701)	1

What the above simulation exercise confirms is that speed of adjustment coefficients must be estimated using non-overlapping returns even at the cost of sample size (see Richardson and Smith, 1991). Non-overlapping half hourly observations are used within this study.
Before reporting speed of adjustment estimates, it is necessary to examine the statistical properties of the sample data and to examine the possible effect that non-synchronous trading may have on speed of adjustment estimates. Statistical descriptions of the data and the influence of infrequent trading are matters discussed within Section 4.

4. Data

The nature of the current study means the choice of observation frequency is critically important. It is important because if it is too low (i.e. daily), then speed of adjustment estimates may miss the true speed at which prices adjust. If the frequency is too high (i.e. at intervals of minutes) estimates may be too noisy to be of value. When deciding on what is optimal, the objective is to use a frequency that minimises the loss in sample size but which is simultaneously able to capture short-term price dynamics with minimal noise.

There are a number of factors that need to be considered when deciding what is optimal. The first is the necessity to set a frequency that is capable of measuring the speed of both actively traded currencies as well as those less actively traded. Currencies adjust at different speeds, so a frequency needs to be used that is capable of measuring the fastest as well as the slowest. Determining what frequency is difficult because as it increases, the amount of noise within the estimates also increases. This is in part due to news providers constantly streaming headline news. Constant streaming does not mean that every headline contains “new” information. Often it is the case that information is repeated, updates provided and analysis revised. These updates, revisions and analysis generate noise within the market. To reduce the level of noise within the speed estimates, frequency should therefore not be observed at intervals of minutes. A large number of alternate frequencies are possible; they could be daily, hourly, or even half hourly. The decision on which to use needs to be decided upon in reference to two criteria. The first relates to the market being observed and the second depends on the nature of the individual study.

The nature of the current study is such that it is not known in advance when new information will arrive. What is known, however, is that information arrives frequently and randomly. It is also true that traders are able to respond to new information at any time of the day and any day of the year. In light of these two facts, if observation frequency is too long (e.g. daily) then the estimate will not be able to account for intra-day information arrival. Frequency therefore needs to be intra-day. Given that the foreign exchange market is open twenty-four hours and that market participants can respond almost instantaneously within any hour, it is likely that hourly observations will be too long. Since minute-by-minute observations are too short and hourly observations too long, half hourly observations appear to be the best possible compromise. The data used within the current study is based on half hour observations for every day of the entire year of 1996. There is no adjustment made within the data for quiet periods, weekends or holidays. This is because it is the intention of the study to examine the continuous nature of the market. It is also recognised that trading will be active when market openings overlap and will be much less active when markets are closed. This will be the case particularly over the weekend. These repeating quiet periods may cause a cyclical pattern within the data which could distort the speed of adjustment estimate. The impact of this, which has been termed non-synchronous trading, needs to be examined and is done so within this section of the paper preceding a description of the data.

Data Description:

Each currency within the study is quoted in U.S. dollar terms[6]. Observations have been collected from GMT 00:30 1 January 1996 to GMT 00:00 1 January 1997. The year 1996 was a typical year within currency markets. There were no major significant financial or currency crises during the sample period.

---

[6] US dollar terms refers to how many US dollars it takes to buy one unit of foreign currency.

The data set[7] is known as HFDF96[8] which originates from Olsen and Associates[9] who have collected the data using proprietary real-time data collection software. The data represents quotes as they are displayed on the Reuters Money Market Headline News screen. Table 3 provides a snapshot of the data as provided by Olsen and Associates. Each exchange rate observation includes the date, time, bid and ask rate, quoted at half hourly intervals for the entire year of 1996.

Table 3: Snapshot of HFD96 Data set for the Australian Dollar

date	time GMT	bid	ask
21.02.96	5:00:00	0.7542	0.7547
21.02.96	5:30:00	0.7536	0.7538
21.02.96	6:00:00	0.7534	0.7539
21.02.96	6:30:00	0.7535	0.7539
21.02.96	7:00:00	0.7528	0.7533

The bid and ask rates are quoted prices rather than prices at which actual transactions took place[10]. When estimating speed of adjustment coefficients, transaction prices would be preferred to quote prices, however, transaction prices are not available from the over-the-counter spot foreign exchange market.[11] Cheung and Wong (2000) point out, however, that the reputation of the trader is so essential to their role as market makers that it is unlikely they would quote prices at which they would not be willing to transact. This means that transaction prices are almost certainly contained within the bid-ask spread of the quoted price. This is confirmed by Petersen and Fialkowski (1994) and Goodhart *et al.* (1995) who find that the midpoint is a good approximation of the transaction price. This has been empirically confirmed by Dacorogna *et al.* (2001 p. 39).
Table 4 lists each currency used in this study, its currency code[12] and the currency system in operation over the sample period.

---

[7] Four outliers were removed from the data set: one observation for the Canadian dollar and three observations for the Malaysian ringgit. Outliers were replaced by the average of the previous and post observation.

[8] The data set contains 25 spot FX rates, 4 spot Metal rates, 6 position 1 Eurofutures contracts, 6 position 2 Eurofutures contracts, 2 stock indices.

[9] Olsen & Associates, Research Institute for Applied Economics. Seefeldstrasse 233, CH-8008 Zurich, Switzerland <http//:www.olsen.ch>

[10] Quoted prices are, however, binding for the market maker.

[11] A growing number of foreign exchange spot transactions go through automated, electronic order-matching systems such as the Electronic Brokering Services (EBS) and Reuters Dealing 2000. These markets deliver high frequency data which reflect real transaction prices and volume. A very limited number of researchers have been able to gain access to this data. See Lyons, R (1995), "Tests of Microstructure Hypotheses in the Foreign Exchange Market", *Journal of Financial Economics*, **39** (23), 321–352 and Goodhart, CAE and Payne, RG (1996), "Microstructural Dynamics in a Foreign Exchange Electronic Broking System", *Journal of International Money and Finance*, **15** (6), 829–852.

[12] The currency code used in this study is defined by the International Organisation for Standardization (ISO code 4217).

Table 4: The Sample of Eight Currencies Used in Study

Country	Currency	Code	Regime
Australia	Dollar	AUD	Independently Floating
Canada	Dollar	CAD	Independently Floating
United Kingdom	Pound	GBP	Independently Floating
Switzerland	Franc	CHF	Independently Floating
Germany	Mark	DEM	Cooperative Arrangement
France	Franc	FRF	Cooperative Arrangement
Italy	Lira	ITL	Independently Floating
Japan	Yen	JPY	Independently Floating

For each currency, the log of the midpoint between the bid and ask rate is calculated:

$$x_j = \log[(P_{bid} + P_{ask})/2] \qquad (16)$$

There is no distinction made for prices quoted during evenings, weekends or holidays; times at which trading activity would be at a minimum. This is because the intention of the current study is to examine the continuous nature of the market. There are 17,568 observations for each of the eight currencies which represent 48 daily half hourly observations over 365.98 days.

Figure 5 plots the log of the midpoint between the bid and ask rate for each currency within the sample at half hourly intervals for the entire year of 1996. Note that the plots of the data are similar to Figure 2 which was the simulated series used to examine properties of the ARMA (1,1) model.

Figure 5
Log of Midpoint between Bid and Ask Rate for Eight Currencies 1996

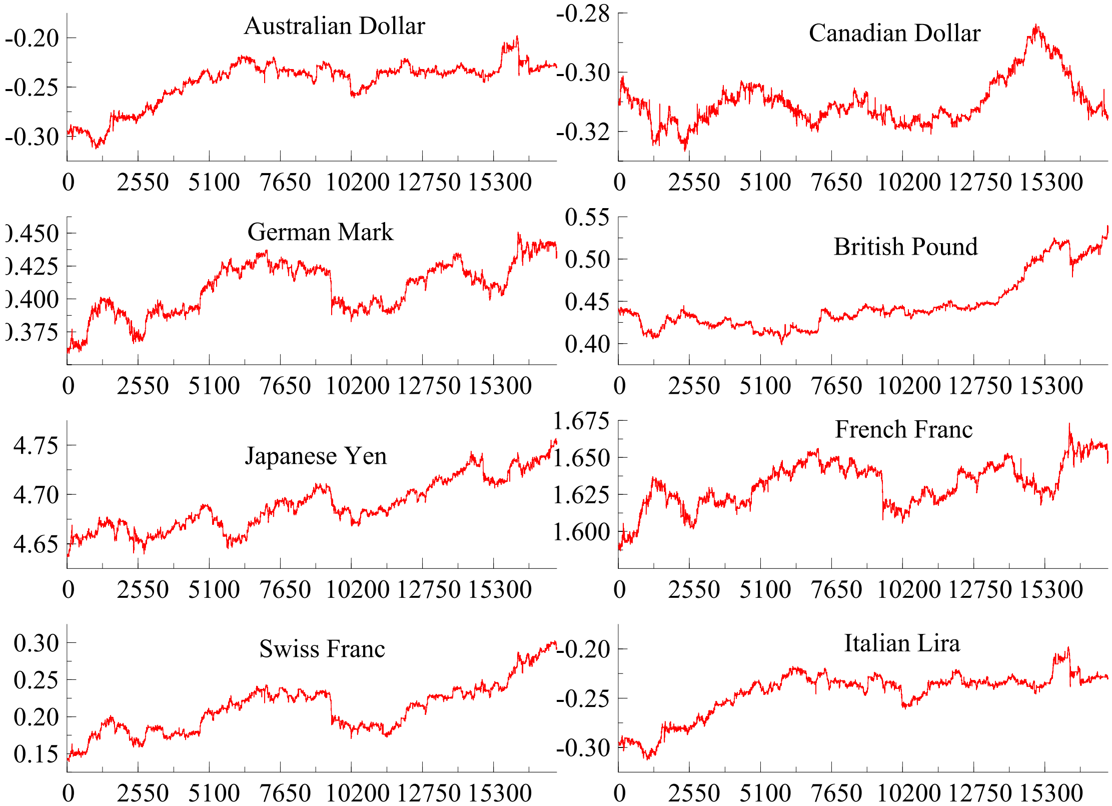

Summary statistics of the half hourly log exchange rate changes are presented in Table 5. Mean returns are close to zero, standard deviations are similar across currencies ranging from a high of 0.09 percent for the Italian lira to a low of 0.04 percent for the Canadian dollar. As is typical of high frequency financial time series data, kurtosis values for each currency are very high. There are no significant differences in the characteristics of these series that would suggest that they should not be compared to one another.

Table 5 Summary Statistics of the Eight Currencies Used in Study

	Mean	*Std. Dev.*	*Skew*	*Kurt*	*ρ(1)*	*ρ(2)*	*ρ(3)*	*Min*	*Max*
Australian dollar	0.0367	0.080759	-0.19	54.35	-0.19	-0.04	0.00	-1.50	1.66
Canadian dollar	-0.0030	0.043754	0.04	52.63	-0.24	-0.04	0.00	-0.89	0.90
British pound	0.0522	0.070965	-0.24	31.01	-0.19	-0.03	-0.01	-1.22	1.20
Swiss franc	0.0840	0.093087	-0.21	29.86	-0.17	-0.01	-0.01	-1.59	1.62
German mark	0.0408	0.073180	-0.08	24.77	-0.13	-0.03	-0.01	-0.93	0.97
French franc	0.0320	0.081155	-0.07	32.89	-0.22	-0.02	-0.03	-1.39	1.31
Italian lira	-0.0199	0.092004	0.07	14.79	-0.30	-0.06	0.02	-0.88	0.87
Japanese yen	0.0641	0.080015	-0.07	13.14	-0.17	-0.02	0.00	-0.90	0.92

Note: the table presents statistics for log exchange rate changes for half hourly observations. Mean and standard deviation are multiplied by 100. ρ (i) records autocorrelation coefficient at lag i.

Min and max record the smallest and largest half hourly percentage changes over the same period.

When estimating adjustment coefficients, increasing non-overlapping, equally spaced time intervals are used to estimate returns. As discussed in Section 3, non-overlapping return intervals need to be used when estimating speed of adjustment coefficients. This is because of the serial correlation that is present within overlapping returns. The presence of serial correlation is shown in Figure 6, where overlapping returns[13] are generated for three differencing intervals for the German mark. When non-overlapping returns are plotted for the same differencing intervals, Figure 6 demonstrates that autocorrelation disappears for each series. This is the case for all of the eight currencies used within the study and is consistent with what was demonstrated within the simulation exercise that was discussed in Section 3.

[13] In this case, Deutschmark returns are used as an example. All currencies used in this study follow the same pattern. When non-overlapping returns are generated, the autocorrelation in the return series disappears.

Figure 6 Comparison of Correlogram of German Mark Returns Overlapping and Non-overlapping

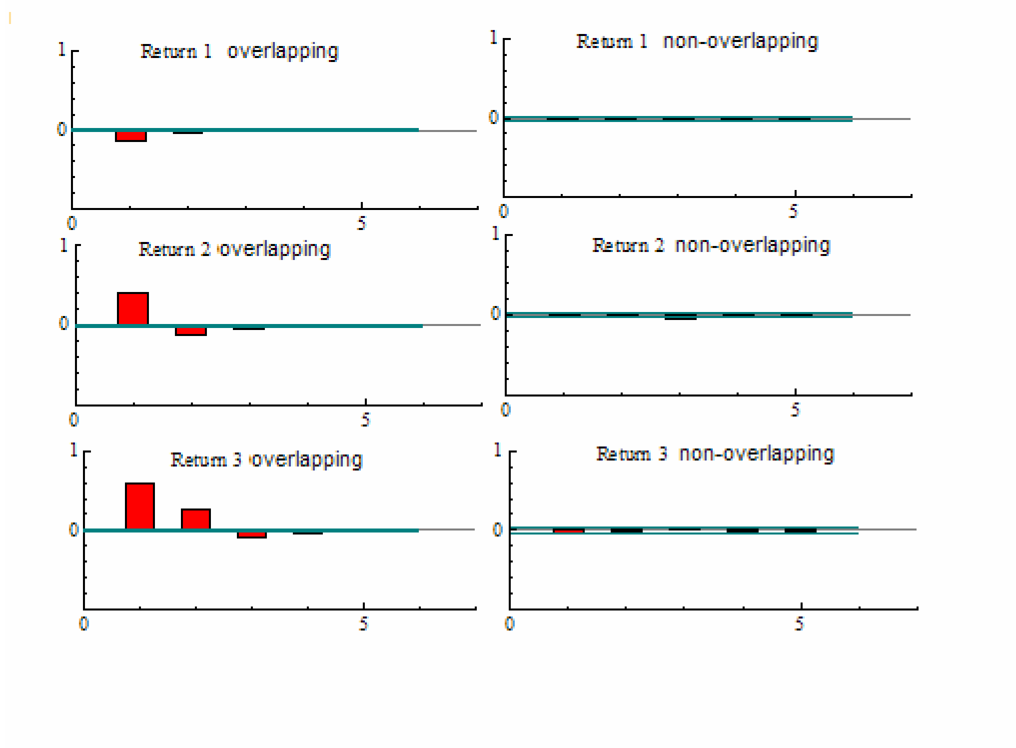

Five different time intervals of non-overlapping returns are calculated for each of the eight currencies. For example, a half hourly non-overlapping return interval is calculated as $[(P_{t2}-P_{t1}),(P_{t3}-P_{t2}),(P_{t4}-P_{t3})......]$, whereas an hourly return interval is calculated as $[(P_{t3}-P_{t1}),(P_{5}-P_{t3}),(P_{t7}-P_{t5})...]$. Table 6 demonstrates that as the non-overlapping return time interval increases, the number of data points decreases.

Table 6: Non-overlapping Return Interval and Sample Size

Series Name	Non-overlapping Time Interval (Hour)	Number of Data Points
R0	0	17,567
R1	.5	8,784
R2	1	5,856
R3	1.5	4,392
R4	2	3,513
R5	2.5	2,928

<u>The Effect of Non-synchronous Trading:</u>

An additional potential source of serial dependence in return series is linked to non-synchronous (infrequent) trading. This occurs when there are successive quotes at the same rate, which causes runs of zero in returns. This can occur when the major currency markets are closed. If the return series has a non-zero mean, this will induce reversion to the mean, which will distort the speed of adjustment estimate. Table 5 indicates that each of the eight exchange rates have means which are very close to zero but which are not equal to zero. This suggests that the effect of infrequent trading would not be significant unless there are a large number of persistent runs of zeros. If this is the case, non-synchronous induced serial dependence will distort speed of adjustment estimates. In other words, the estimate will be less a function of under or over-reaction and more a function of persistent runs of zero returns. It is therefore worth examining the impact that infrequent trading has for each currency within the sample.

An important, associated issue is that the sets of return series used within the current study represent quote data and not transaction data. It is a stylised fact that quote data adjusts more often than do transaction prices.[14] This means that quote data will be less prone to extended runs of zero returns. It remains unclear, however, to what extent non-synchronous trading generates serial dependence within exchange rate quote data.

Lo and MacKinlay (1990) derive a methodology to estimate the extent of serial correlation generated from infrequent trading. The methodology can be used to examine either transaction or quote data. The central idea is that if a return series is a random walk, then non-synchronous serial dependence will be a function of the sample mean, variance and the probability of no trade. The proportion of zero return observations within the series is thus a natural proxy for the probability of no trade. If returns are generated by:

$$r_t = \mu + \varepsilon_t \quad (17)$$

where ε_t is white noise and π is the probability of no trade, then non-synchronous induced correlation $\rho(i)$ in returns r_t at lag i can be written as:

$$\rho(i) = \frac{-\mu^2 \pi^i}{\sigma^2 + \frac{2\pi}{1-\pi}\mu^2} \quad (18)$$

Table 7 reports the serial dependence in returns caused by non-synchronous trading as a percentage of the autocorrelation at the first non-overlapping difference interval.

[14] There will be more quotes than real trades. Quotes will therefore adjust to information quickly while transaction data will accumulate the information arrival between one transaction and the next.

Table 7: Percentage Contribution of Non-Synchronous Trading to First Order Serial Correlation in Non-overlapping Returns

Currency	ρ(1)	ρ*(1)	% Contribution of ρ*(1)
Australian dollar	-0.19	-6.5E-06	0.003403
Canadian dollar	-0.24	-1.6E-07	6.7E-05
British pound	-0.19	-1.3E-05	0.0071
Swiss franc	-0.17	-2.1E-05	0.012601
German mark	-0.13	-7.1E-06	0.005454
French franc	-0.22	-4.8E-06	0.002162
Italian lira	-0.3	-1.2E-06	0.00041
Japanese yen	-0.17	-1.5E-05	0.009088

ρ(1) is the autocorrelation coefficient and ρ*(1) is the contribution that non-synchronous trading makes to the first order autocorrelation coefficient[15]

The results reported in Table 7 clearly indicate that the contribution that non-synchronous trading makes to the first order autocorrelation coefficient for each of the eight currencies is less than 0.01%. This confirms that infrequent trading has little impact on the size of the first order autocorrelation in non-overlapping return series of the eight exchange rates. This in turn implies that the size of the AR (1) coefficient is primarily a function of the market's over or under-reaction to information arrival. Whether this is indeed the case is examined within the results section which follows.

5. Results

Using half hourly return intervals, speed of adjustment estimates are calculated for a set of eight currencies using three speed of adjustment methodologies. Estimates from each methodology are then interpreted and compared. To enable aggregate comparisons, the set of eight currencies are separated into two subsets. The first set of four currencies represents the world's most actively traded currencies.[16] For the purposes of this study they are defined as Tier 1 currencies. They include the German mark, Japanese yen, British pound and French franc. Tier 2 currencies are also actively traded but trade less actively than Tier 1 currencies. They include the Canadian dollar, Australian dollar, Swiss franc and Italian lira.

Comparisons are also made among individual currencies. For example, the Australian dollar is compared to the speed of adjustment estimates of the other seven currencies as a way of determining whether AUD adjustment is consistent with what would be expected given the speed at which the others adjust.

[15] Equation 6.16 was estimated over each of the 20 non-overlapping return intervals for the Australian dollar. The percentage contribution that ρ*(1) made remained marginal regardless of the time interval used.
[16] This information is drawn from table F-3, Bank for International Settlement's Triennial Central Bank Survey and Derivatives Market Activity for 1995.

Comparisons are also made between currencies that are independently floating and those that are members of a cooperative arrangement (eg. the European Monetary System).
Table 8 reports average speed of adjustment coefficients for Tier 1 and Tier 2 currencies using three estimation procedures at the first two half hourly return intervals. Analysis and comparison of these estimates provides a number of interesting insights.

Table 8: Average Speed of Adjustment By Estimator

Type of Currency	Intrinsic Variance		Auto-covariance		ARMA (1,1)	
	½ hr	*1 hr*	*½ hr*	*1 hr*	*½ hr*	1 hr
Tier 1 Currencies	.8935	.9693	.8503	.8240	.7728	.7855
Tier 2 Currencies	.5920	.8296	.8413	.8269	.7852	.7377
All Currencies	.7595	.9072	.8463	.8252	.7783	.7617

All estimators report a degree of under-reaction at the shortest (half hour) differencing interval. On average, Tier 1 and 2 currencies do not fully adjust within the first thirty-minute interval. Table 8 does indicate, however, that on average the bulk of adjustment, (approximately 75 percent), for all currencies occurs within thirty minutes, and that there is very little additional adjustment taking place in subsequent periods. This attests to the fact that currency adjustment is “quick”. It is consistent with what would be expected within an efficient market.

In addition, Table 8 demonstrates that there are differences in the degree of adjustment between Tier 1 and 2 currencies, and that the size of this difference varies depending on which estimator is used. For example, the Damodaran (1993) intrinsic variance estimator reports a significant difference between the degree of adjustment within the first half hour for Tier 1 (.8935) and Tier 2 (.5920) currencies. The other two estimators report a marginal difference. A possible explanation why adjustment differs between Tier 1 and Tier 2 may be linked to liquidity. Tier 1 currencies are more liquid than Tier 2 and thus may adjust faster because they are easier to trade. This explanation is consistent with what is found within the equities market. Jegadeesh and Titman (1995), for example, find that large capitalisation stocks which are very liquid adjust quicker than do small capitalisation stocks which are much less liquid.

Table 8 also demonstrates that the Auto-covariance ratio and the ARMA (1,1) estimator report insignificant differences in the degree of adjustment between Tier 1 and Tier 2 currencies. This is true at either the half hour interval or one hour interval. The question remains as to why a large difference exists depending on which estimator is used. The answer may simply be that the difference is a reflection of the poor calibration of one of the estimators. Theobald and Yallup (2004, p. 86) find this to be the case.

They find that when comparing the three methodologies using simulation and historical data the Damodaran (1993) intrinsic variance estimator, “is particularly prone to being swamped by noise at shorter differencing intervals, thereby reducing the precision and reliability of the results generated…”. It is certainly the case that half hourly differencing intervals, as used within the current study, are short differencing intervals and thus have the potential to swamp the intrinsic variance estimator. In fact, half hourly differencing is shorter than any differencing interval that has been used to date in any previous studies that have used this estimator. This being the case, it is likely that the intrinsic variance estimator is distorted by noise at this very short interval. For this reason, the Auto-covariance ratio and ARMA (1,1) estimators will be the preferred measurement methodologies for making subsequent comparisons of speed of adjustment coefficients. The Auto-covariance ratio and ARMA (1,1) estimators indicate an insignificant difference between the degree of adjustment within the first differencing interval for Tier 1 and Tier 2 currencies. This suggests that there is sufficient liquidity within the currency markets for both Tier 1 and Tier 2 currencies. Indeed, it is as easy to trade the German mark as it is to trade the Australian dollar. This is because there are a sufficient number of market makers who are ready to trade twenty-four hours a day, 365 days a year.

Since liquidity is not an issue, the speed at which the two sets of currencies adjust should also be similar. This is demonstrated in Table 8. There is less than a 2 percent difference in the degree of adjustment that occurs within the first half hour between Tier 1 and Tier 2 currencies. Irrespective of the fact that the size of the difference is marginal, a 2 percent difference remains between them. This difference may simply reflect the fact that Tier 1 currencies are traded more actively than Tier 2; the market will adjust marginally quicker as a result. Whether this is indeed the case is explored in Table 9.

Table 9 is a comparison of the individual ranking of exchange rates by average daily turnover against the ranking based on the degree of adjustment that takes place within the first hour. The higher the turnover of the currency, the higher the ranking it receives. Additionally, the greater the degree of adjustment of the currency, the higher the ranking it receives. As anticipated, rankings are relatively consistent. For example the German mark, which is ranked number one in terms of average daily turnover, is also ranked number one in terms of the degree of adjustment that takes place within the first hour. The rankings are consistent for the Japanese yen, British pound, Swiss franc, Canadian dollar and Italian lira. These rankings confirm that actively traded currencies adjust quicker irrespective of the fact that the degree of difference in adjustment is marginal. The only currency that does not to line-up is the Australian dollar; its speed of adjustment is much quicker than its turnover would suggest.

Table 9: Ranking of Sample Currencies by Market Turnover and Degree of Adjustment Within First Hour using the ARMA Estimator

Currency[1]	Percentage of average daily turnover[2]	Ranking by Turnover 1995	Speed of Adjustment Coefficient	Ranking by Degree of Adjustment
German mark	37	1	1	1
Japanese yen	24	2	.8292	3
British pound	10	3	.7895	4
French franc	8	4	.7379	7
Swiss franc	7	5	.7622	5
Canadian dollar	3	6	.7409	6
Australian dollar	3	7	.8362	2
Italian lira	1	8	.6118	8

[1] Note that the U.S. dollar is excluded; this is because all currencies are quoted direct against U.S. dollar. [2]Source: Bank of International Settlement, (1998). "Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity". Percentage share is adjusted for double counting.

It is worth examining whether the intrinsic variance estimator generates a similar consistency in rankings. Table 10 compares the ranking of currencies based on degree of adjustment and average daily turnover using the intrinsic variance estimator. Comparison of the rankings indicates consistency with the rankings reported in Table 9, even though the size of the adjustment coefficients is much smaller. This highlights that the Damodaran (1993) estimator produces rankings which are similar to those generated from the ARMA (1,1) estimator even though it is an inferior estimator.

Table 10: Ranking of Sample Currencies by Market Turnover and Degree of Adjustment within First Hour using Intrinsic Variance Estimator

Currency[1]	Percentage share of average daily turnover[2]	Ranking by Turnover 1995	Intrinsic Variance Estimator	Ranking by Degree of Adjustment
German mark	37	1	1.0146	1
Japanese yen	24	2	.9910	2
British pound	10	3	.9252	5
French franc	8	4	.9464	3
Swiss franc	7	5	.9385	4
Canadian dollar	3	6	.7018	8
Australian dollar	3	7	.8502	6
Italian lira	1	8	.8279	7

[1] Note that the U.S. dollar is excluded; this is because all currencies are quoted direct against U.S. dollar. [2]Source: Bank of International Settlement, (1998). "Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity". Percentage share is adjusted for double counting.

Table 11 compares the degree of adjustment for currencies which are free floating against those that operate within managed systems. The size of the difference in degrees of adjustment for the two sets of estimates are much larger (approximately 10%) than are the differences that exist between actively traded versus less actively traded (approximately 3%).

Table 11: Average Speed of Adjustment Coefficient at First Half Hour

	Auto-covariance Ratio	ARMA (1,1)
Free floating Exchange Rates	.8506	.7987
Managed Exchange Rates	.8315	.7202
All Currencies	.8462	.7783

This larger difference highlights that the distinction between actively traded and less actively traded currencies is much less important than the distinction between managed and free floating exchange rates. This result is consistent with the argument that government intervention distorts adjustment to fundamental value. The ARMA (1,1) estimator finds that 80 percent of adjustment is complete within the first half hour for free floating rates, as opposed to only 72 percent for managed rates. This is consistent with what would be expected given the differences in the degree of government intervention that would occur within the two types of systems.

Table 12 reports speed of adjustment estimates for individual currencies over a two and half hour period. From the perspective of the efficient market hypothesis, the speed at which individual exchange rates adjust to intrinsic value and whether that adjustment persistently overshoots is of fundamental importance. Also from a "relative efficiency" perspective it is important to examine whether individual currencies are as efficient at absorbing information as their counterparts. These issues are examined within Table 12.

Table 12: Estimators of Exchange Rate Speed of Adjustment

	Damodaran					Theobald and Yallup									
	Intrinsic Variance Estimator					Auto-covariance Estimator					ARMA (1,1) Estimator				
	Non-overlapping Half Hourly Return Interval														
Tier 1 Currencies	**R1** ½ hr	**R2** 1 hr	**R3** 1 ½ hr	**R4** 2 hr	**R5** 2 1/2 hr	**R1** ½ hr	**R2** 1 hr	**R3** 1 ½ hr	**R4** 2 hr	**R5** 2 1/2 hr	**R1** ½ hr	**R2** 1 hr	**R3** 1 ½ hr	**R4** 2 hr	**R5** 2 1/2 hr
German mark	.9769	1.0146	1.0032	1.02	1	.7708	.9465	.7631	.4907	-.0377	.7290	.9417*	.3732	.4997	.2729
Japanese yen	.9329	.9910	1.0105	1.0385	1	.8956	.8464	.8014	.5211	.7652	.8655	.8292	.6689	.5645	.3875
British pound	.8218	.9252	.9759	.9434	1	.8424	.8174	.9987	.8911	.7746	.7853	.7895	.9523*	.8867*	.6268
French franc	.8424	.9464	1.0314	1.0111	1	.8922	.6855	.8862	.9411	.7246	.7114	.7379	.8397	.9430*	.8026*
Tier 2 Currencies															
Canadian dollar	.2954	.7018	.7042	.8484	1	.8546	.8616	.6680	.7293	1.1626	.7785	.7409	.6773	.7014	1.1290
Australian dollar	.6106	.8502	.8716	.9262	1	.7899	1.0187	.1671	.6685	.7035	.7459	.8362	.5658	.6117	.7503
Swiss franc	.8930	.9385	1.0324	1.0119	1	.9309	.8604	.7993	.5415	.5352	.8707	.7622	.7494	.7396	.5097
Italian lira	.5688	.8279	.9165	.9309	1	.7899	.5667	1.0519	.7375	1.0367	.7460	.6118	.8628*	.7641	.9601*

* Denotes statistically insignificant parameter estimate at 5 percent confidence level.

Table 12 indicates that the speed of adjustment at the half hourly interval for each currency is less than complete. For example, the degree of Australian dollar adjustment within the first half hour is approximately 75 percent. This degree of adjustment is consistent with the average level of adjustment for both Tier 1 (76%) and Tier 2 currencies (78%). Based on these comparisons, the Australian dollar is deemed to be relatively efficient in relation to its counterparts.

When examining the estimates for the Auto-covariance estimator and ARMA (1,1) estimator, for all currencies the greatest degree of adjustment occurs within the half hour to hour interval regardless of the estimator. As the return interval widens, adjustment tends to deteriorate for each currency. For example, the speed of adjustment for the Japanese yen is highest within the first half hour (.8655) but then begins to decline as the return interval widens (.8292, .6689). There are two interesting questions associated with these results. The first is why less than full adjustment occurs initially and second, why does adjustment deteriorate as the interval widens.

One explanation is that full adjustment occurs in between the half hourly intervals. For example, it is possible that for the German mark, Japanese yen, Canadian dollar, Swiss franc and Italian lira, full adjustment occurs after 30 minutes but is complete sometime before 60 minutes, and that for the remaining currencies (British pound, French franc and Australian dollar) adjustment occurs after 60 minutes but within 90 minutes. This does not, however, explain why adjustment remains less than complete after this interval and why it deteriorates as the return interval is extended.

A possible explanation for the deterioration is that the impact of noise[17] within the widening non-overlapping return interval becomes increasingly prevalent. In other words, as the high frequency return interval is widened, a greater degree of noise is captured, which results in deteriorating speed of adjustment estimates.

Noise may also be the explanation why adjustment coefficients do not converge to one. It may be the case that complete adjustment never occurs because of persistent noise caused in part by government intervention, herd behaviour, technical trading, etc. The fact that the bulk of adjustment (75%) does occur “quickly” is, however, sufficient to conclude that the currency market is a relatively efficient processor of information albeit not perfect.

Less than full adjustment suggests there may be historical information that could be used to predict future returns. Indeed, this appears to be the case. Table 13 reports that as the non-overlapping return interval increases, the size of the autoregressive coefficient within the ARMA (1,1) estimation also has a tendency to increase.

[17] Noise which causes distortion would include that which is induced by government intervention, uninformed trading and trading that is based purely on technical trading rules and strategies; all of which are prevalent within foreign exchange markets.

Table 6.13 Estimates of the AR(1) Coefficients

Tier 1 Currencies	R1	R2	R3	R4	R5
German mark	.2710 (.046420)	.0583* (.077992)	.6269 (.072019)	.5003 (.092121)	.7272 (.075617)
Japanese yen	.1345 (.040928)	.1708 (.077493)	.3311 (.098481)	.4356 (.099723)	.6126 (.110801)
British pound	.2147 (.032148)	.2105 (.056677)	.0477* (.078756)	.1133* (.119929)	.3733 (.110229)
French franc	.2886 (.026000)	.2621 (.043456)	.1603 (.053751)	.0570* (.073012)	.1974* (.101872)
Tier 2 Currencies					
Canadian dollar	.2215 (.024107)	.2591 (.036810)	.3227 (.063571)	.2986 (.067891)	-.1290 (.063032)
Australian dollar	.2541 (.031026)	.1638 (.051350)	.4342 (.061393)	.3884 (.079784)	.2497 (.081247)
Swiss franc	.1295 (.041514)	.2378 (.081805)	.2506 (.080294)	.2604 (.107264)	.4903 (.125817)
Italian lira	.2541 (.031026)	.3882 (.045671)	.1372* (.072380)	.2360 (.084527)	.0399* (.087216)
* Indicates insignificance at the 95 percent confidence interval. Numbers in brackets are standard errors.					

The existence of statistically significant AR (1) estimates, particularly at increasing return intervals, is evidence that there is information in previous observations that is related to current observations. This is inconsistent with what the efficient market hypothesis (EMH) would predict. The EMH predicts there should be no past or current information that can be used to predict future prices. This suggests that according to Fama's (1970) definition of market efficiency, the foreign exchange market is inefficient because full adjustment to fundamental value does not occur with information arrival. Indeed, as the return interval lengthens, the strength of the AR (1) coefficient increases, causing the speed of adjustment coefficient to deteriorate. An increasing AR (1) coefficient with time indicates there is some form of past information that can be used to predict future returns.

This finding of inefficiency is consistent with other studies which examine the foreign exchange market. For example there continues to be no explanation for the failure of the speculative efficiency hypothesis nor is there a reason to explain the existence of the forward rate bias. The results of the current study contribute to this body of literature by finding the market to be inefficient in that full adjustment to fundamental value does not occur.

6. Conclusion

Efficient market theory predicts that prices should adjust "quickly" with the arrival of information. This paper has measured speed of adjustment to determine if currency markets are efficient. This has been done by using three unique measurement methodologies, for eight currencies, using half hourly continuous non-overlapping return intervals over the entire year of 1996. The estimators are grounded within the partial adjustment price model of Amihud and Mendelson (1987). Based on these measurement methodologies, foreign exchange markets are deemed to be relatively efficient in that the bulk of exchange rate adjustment (75%) occurs within thirty minutes.

This study is unique in that it measures speed of price adjustment using high frequency data and methodologies that have not previously been used to examine the foreign exchange market. The paper is therefore able to provide a number of important insights. It has been demonstrated that for all currencies there is an initial under-reaction at the half hour differencing interval. This is evidence indicating that exchange rates do not persistently overshoot fundamental value with information arrival. Results also confirm that there is little difference between the speediness of adjustment for currencies which are actively traded (Tier 1 currencies) as compared to those that are traded less actively (Tier 2). Why this should be so may in part be explained by the fact that the foreign exchange market is sufficiently liquid. Irrespective of the fact that the size of the difference is marginal (3%), a difference does exist. It is therefore the case that actively traded currencies adjust marginally quicker than less actively traded currencies. This is consistent with what would be expected.

Furthermore, a substantial difference in the degree of adjustment could be seen between exchange rates which operate within managed systems and those that are free floating. Free floating rates were found to adjust to a greater extent and thus were quicker (by approximately 10%) within the first half hour than regulated exchange rates. The reason for this may be that managed currencies are distorted by government intervention than are their free floating counterparts.

This paper has also been able to demonstrate that adjustment deteriorates as the differencing interval increases, and that full adjustment does not occur. There are two possible explanations for this. The first is that there is a persistent level of noise in the market which distorts price from fundamental value. Because it is persistent, full adjustment never occurs. The second is that, as the difference interval increases, it captures more information. As information builds, in other words, when the implication of past news is offset by the implication of recent news, then the speed of adjustment coefficient is unable to detect clear signals.

The fact that the AR (1) coefficient strengthens as the differencing interval widens indicates there is information in past observations that can be used to predict future values. This is inconsistent with what the efficient market hypothesis would predict.

The significant contribution of this study has confirmed that foreign exchange markets are relatively efficient processors of information given the constraints imposed by government intervention and noise. Exchange rates adjust “quickly” to the arrival of information although not completely.

References

Amihud, Y. and Mendelson, H. (1987), "Trading Mechanisms and Stock Returns: An Empirical Investigation/Discussion." Journal of Finance, 42 (3), 533-555.

Ariff, M. and Chan, D. (2002), "Speed of Share Price Adjustment to Information", Managerial Finance, 28 (8), 44-65.

Barberis, N., Shleifer, A., and Vishny, R. (1998), "A Model of Investor Sentiment", Journal of Financial Economics, 49 (3), 307.

Brisley, N. and Theobald, M. (1996), "A Simple Measure of Price Adjustment Coefficients: A Correction", Journal of Finance, 51 (1), 381-382.

Cheung, Y.-W. and Wong, C. Y.P. (2000), "A Survey of Market Practitioners' Views on Exchange Rate Dynamics", *Journal of International Economics*, **51** (2), 401-419.

Dacorogna, M. M., Gencay, R., Muller, U., Olsen, R. and Olivier, P. (2001), *An Introduction to High-Frequency Finance*, Academic Press, London.

Damodaran, A. (1991), " Information and Price Adjustment Processes: A Comparison of U.S. and Japanese Stocks", In S.J. Khoury (Editor), Recent Developments in International Banking and Finance, Elsevier Science, North-Holland, Amsterdam.

Damodaran, A. (1993), "A simple measure of price adjustment coefficients", Journal of Finance, 48 (1), 387-400.

Dornbusch, R. (1976), "Expectations and Exchange Rate Dynamics", Journal of Political conomy, 84 (6), 1161-1176.

Ederington, L. and Lee, J. (1993), "How Markets Process Information: News Releases and Volatility." The Journal of Finance, 48 (4), 1161-1191.

Eilifsen, A., Knivsfla, K., and Saettem, F. (2001), "Earnings Announcements and the Variability of Stock Returns", Scandanavian Journal of Economics, 17 (2), 187-200.

Goodhart, C., Ito, T. and Payne, R. (1995), "*One Day in June, 1993: A Study of the Working of Reuters 2000-2 Electronic Foreign Exchange Trading System*", NBER Working Paper #t0179.

Goodhart, C. A. E. and Payne, R. G. (1996), "Microstructural Dynamics in a Foreign Exchange Electronic Broking System", Journal of International Money and Finance, 15 (6), 829-852.

Jegadeesh, N. and Titman, S. (1995), "Over-reaction, Delayed Reaction, and Contrarian Profits", Review of Financial Studies, 8 (4), 973-993.

Lo, A. W. and MacKinlay, A. C. (1990), "An Econometric Analysis of Nonsynchronous Trading", Journal of Econometrics, 45 (1-2), 181-211.

Lyons, R. (1995), "Tests of Microstructure Hypotheses in the Foreign Exchange Market", Journal of Financial Economics, 39 (23), 321-352.

Miller, M., Muthuswamy, J. and Whaley, R. (1994), "Mean Reversion of Standard and Poor 500 Index Basis Changes: Arbitrage Induced or Statistical Illusion?", Journal of Finance, 49 (2), 479-514.

Muller, U. A. (1993), "Statistics of Variables Observed Over Overlapping Intervals", Olsen Associates Working Paper #UAM.1991-10-14.

Petersen, M. A. and Fialkowski, D. (1994), "Posted versus Effective Spreads: Good Prices or Bad Quotes?", *Journal of Financial Economics*, **35** (3), 269.

Richardson, M. and Smith, T. (1991), "Tests of Financial Models in the Presence of Overlapping Observations", Review of Financial Studies, 4 (2), 227-254.

Safvenblad, P. (1997), "On the Damodaran Estimator of Price Adjustment Coefficients". Stockholm: Stockholm School of Economics, Department of Finance.

Theobold, M. and Yallup, P. (2004), "Determining Security Speed of Adjustment Coefficients", Journal of Financial Markets 7 (1) 75-96.