ELSEVIER Journal of Empirical Finance 4 (1997) 187–212

Journal of EMPIRICAL FINANCE

Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model [1]

Robert F. Engle [a], Jeffrey R. Russell [b,*]

[a] Department of Economics 0508, University of California, San Diego, La Jolla, CA 92093-0508, USA
[b] University of Chicago, Graduate School of Business, 1101 East 58th St., Chicago, IL 60537, USA

Abstract

This paper applies the Autorecressive Conditional Duration model to Foreign Exchange quotes arriving on Reuters screens. The Autoregressive Conditional Duration model, developed in Engle and Russell (1995) [Engle, R., Russell, J., 1995. Autoregressive conditional duration; a new model for irregularly spaced time series data, University of California, San Diego, unpublished manuscript.], is a new statistical model for the analysis of data that does not arrive in equal time intervals. When Dollar/Deutschmark data are examined, it is clear that many of the price quotes are simply noisy repeats of the previous quote. By systematically thinning the sample, a measure of the time between price changes is developed. These price durations are modeled with the ACD to obtain estimates of the instantaneous intensity of price changes. This measure is related to standard measures of volatility but is formulated in a way that incorporates the information in the irregular sampling intervals. A simple market microstructure model implies that the bid–ask spread should have predictive power for the volatility which is supported by the data. A model of price leadership however, is not supported. © 1997 Elsevier Science B.V.

* Corresponding author.

[1] The authors would like to thank David Brillinger, Sir David Cox, Clive Granger, Andrew Karolyi, and Gennady Samorodnitsky for their comments and suggestions. Useful comments were also received from participants at the 1995 Olsen and Associates meetings on high frequency data analysis. The research was supported in part by National Science Foundation Grant SES-9122056. Jeffrey Russell is grateful for financial support from the Sloan Foundation and the UCSD Project in Econometric Analysis Fellowship.

PII S0927-5398(97)00006-6

1. Introduction

The foreign exchange market is a massive market with international participants trading billions of dollars 24 h a day. Transactions are carried out in split seconds between parties across the globe. As the market is a dealer market there are no systematic accounts of the movements in the market except for the advertised quotes which appear continuously on trading desks over the Reuters screens. These update traders on a second by second basis giving bid and ask prices for the major players in the market. Transactions then occur bilaterally and are not systematically recorded.

Often investigators are interested in studying the behavior of the exchange rate process. See, for example, Bollerslev and Domowitz (1993) or Goodhart et al. (1993). Price quote data inherently arrive in irregular time intervals. Since the majority of time series techniques are based on equally spaced intervals between arrivals of data, there is a natural inclination of the econometrician to choose some fixed interval and aggregate the information within each interval. The choice of an interval is potentially a very important part of the analysis. If too short of an interval is chosen then most of the data points will be zeros and heteroskedasticity of a particular form will be introduced. On the other hand if too lone, an interval is chosen, then features of the data will be smoothed and potentially hidden.

Engle and Russell (1995) developed a model for data that do not arrive in equal time intervals. Their paper proposed a new point process model for intertemporally correlated event arrival times. Rather than aggregate to some fixed interval, it is proposed to treat the arrival times of the data as a point process with an intensity defined conditional on past activity. Because the model formulation focuses on the intertemporal correlations of the durations (time interval between events), the model is called the autoregressive conditional duration (ACD) model.

Point processes frequently incorporate other information which modifies or 'marks' the arrival times. Engle and Russell (1995) describe an algorithm for modeling marks by focusing only on some types of events. For the FX data analyzed in this paper, it is apparent that many of the quotes are simply repeats of previous quotes. At some point, however, a quote will arrive that is different. This paper uses the thinning algorithm discussed in Engle and Russell to model the intensity of quote arrivals which signify price chances. The model is called the ACD model for price based durations.

This model predicts how long it will be until prices change. A trader might be interested in knowing this time interval as it could influence the speed with which he places an order. Each time the price changes, there is an interval of time during which he can trade at that price. If the market is slow, this interval may be quite long, while in an active market, the price may last much less than a minute. If the trader is using an automated trading, system, then the system itself must recognize that sometimes a delay will have no impact on the available prices while at others it may eliminate the opportunity.

By modeling the time it takes for prices to move a certain amount, the model is essentially a volatility model or more precisely the inverse of a volatility model. Because the model allows weakly exogenous and lagged dependent variables to influence the intensity of the price process, simple LM tests can be performed to investigate the determinants of price volatility. The hypothesis of price leadership can be examined since the sources of the quotes are recorded.

Each price duration can potentially be interpreted as the delayed response of the market makers quotes to an information event as in the Glosten and Milgrom (1985) model. Market makers infer the direction and size of a new piece of privately available information by examining the requests to buy or sell assets. As some of the agents may be privately informed, the market makers will set and adjust bid and ask quotes in response to excess demand or supply which serves as a signal of the news. Both the speed of adjustment and the bid–ask spread depend upon the fraction of the traders assumed to be informed and whether they are successful in disguising themselves as liquidity traders. Easley and O'Hara (1992), Easley et al. (1994, 1995) point out that market makers do not in general even know whether there has been an information event. Thus, slow trading can be interpreted as evidence that there has been no new information, and the price adjustment will be consequently slowed. Thus the time until prices change can be given a structural interpretation in terms of the rate at which information is released and the rate at which the market incorporates this information into prices.

Section 2 of the paper will present the ACD model with both the statistical underpinnings and the motivation. Section 3 describes the data and Section 4 gives the results. Section 5 gives further analysis of the model and tests some interesting hypotheses, Section 6 develops a relation between the price duration measure of volatility and standard measures of volatility and Section 7 concludes.

2. The ACD model

The statistical problem is to estimate the probability of an event such as a quote arrival at each point in time. This requires specifying the stochastic process of the arrival times, estimating the parameters and computing the probabilities of events. The instantaneous probability of an event is called the intensity of the process and in dependent processes such as the ones considered here is conditional on past information. Once the intensity is parameterized, the likelihood can be computed and parameters estimated and checked.

Engle and Russell (1995) proposed the ACD class of statistical models for arrival times which are dependent. The data are simply a list of times and possibly characteristics, or marks, associated with the arrival times. They considered the arrival rates of IBM transactions on the NYSE, while this paper will develop similar models for the arrival rates on Reuters screens of foreign exchange quotes on the Dollar Deutschmark.

Consider the stochastic process which is simply a sequence of times $\{t_1, t_2, \ldots, t_n, \ldots\}$. As these are points distributed in time, this is called a 'point process' and the times are called 'arrival times' of the point process. Corresponding to these arrival times is a counting process, $N(t)$ which is the number of events which have occurred by the time t. Obviously, $N(t)$ is a non-decreasing function of time with $N(t_0) = 0$.

Defining the 'conditional intensity' of a process as

$$\lambda(t; N(t), t_1, \ldots, t_{N(t)}) = \lim_{\Delta t \to 0} \frac{P(N(t+\Delta t) > N(t) | N(t), t_1, \ldots, t_{N(t)})}{\Delta t}, \tag{1}$$

the intensity over the next instant is potentially a function of the entire past history of the process. This expression generalizes the familiar Poisson process where λ is a constant, and the inhomogeneous Poisson process where it depends only upon t but not past arrivals. The Eq. (1) is used to characterize 'self exciting point processes' as described in Snyder and Miller (1991) and attributed variously to Rubin (1972), or Hawkes (1971). Self exciting processes allow the past events to influence the future evolution of the process. Such processes have been used to model dependent point processes such as earthquakes, electron emissions and nerve cell firings. The model nests a variety of point processes described in the literature. If the intensity depends only on $N(t)$ but not the timing of events, then the process is a 'birth process', and if it depends only on the m most recent arrival times, it is called a m-memory self exciting process.

Once the local properties of the process are specified by a parameterization of Eq. (1), the probabilities of counts and waiting times are all fully specified. In particular the likelihood function of the observed times can be derived from the intensity process. It is convenient to write this likelihood function in terms of the durations which are the intervals between successive arrival times. Let $x_i = t_i - t_{i-1}$ for $i = 1, \ldots, N(T)$ be the data set. Then the log likelihood can be expressed in terms of the conditional density function of x_i as:

$$L(x_1, \ldots, x_{N(T)}; \theta) = \sum_{i=1}^{N(T)} \log f(x_i | x_1, \ldots, x_{i-1}; \theta) \tag{2}$$

The ACD model is a new and convenient class of specifications for the self exciting process. The model does not necessarily have limited memory and is easy to estimate by maximum likelihood. The simplest specifications allow analytic forecasts of waiting times. The crucial assumption for the ACD model is that the time dependence can be summarized by a function ψ which is the conditional expected duration given past information and has the property that x_i/ψ_i are independent and identically distributed. That is, the density of these 'standardized durations' satisfy:

$$\psi_i = E[x_i | x_{i-1}, \ldots, x_1; \theta] \tag{3}$$

and

$$x_i/\psi_i \text{ are i.i.d.} \tag{4}$$

This can also be written as

$$x_i = \psi_i \varepsilon_i \tag{5}$$

where ε is an i.i.d. series of disturbances with a distribution which must be specified. Associated with this distribution is a hazard function given by the probability density of ε divided by the survival function of ε which is simply one minus the cumulative distribution function. The hazard function of ε is often called the baseline hazard since it does not depend upon any conditioning information. If f is the density of ε, the baseline hazard is given by:

$$h_0(t) = \frac{f(t)}{\int_t^\infty f(u)\,\mathrm{d}u} \tag{6}$$

It is natural to call the model autoregressive conditional duration or ACD because it parameterizes the conditional duration in terms of the lagged durations.

From Eqs. (3) and (5) it is apparent that there is a vast set of ACD model specifications defined by different distributions of ε and specifications of ψ. For example, a natural starting point is to assume that durations are conditionally exponential. In this case, the conditional intensity is given by

$$\lambda\left(t \mid x_{N(t)}, \ldots, x_1\right) = \psi_{N(t)+1}^{-1} \tag{7}$$

which is therefore constant from one arrival to the next, although it has a step function at each arrival time. A simple m-memory specification of the intensity is given by:

$$\psi_i = \omega + \sum_{j-1}^{m} \alpha_j x_{i-j} \quad \text{for } \alpha_j \geq 0,\ \omega > 0,\ \forall i,\ i = 1, \ldots, N,\ j = 1, \ldots, m \tag{8}$$

A natural generalization of this model introduces infinite memory by including q lagged durations:

$$\psi_i = \omega + \sum_{j-1}^{m} \alpha_j x_{i-j} + \sum_{j=1}^{q} \beta_j \psi_{i-j} \quad \text{for } \alpha_j,\ \beta_j \geq 0,\ \omega > 0,\ \forall i,\ i = 1, \ldots, N,$$

$$j = 1, \ldots, m, q \tag{9}$$

This model called the ACD(m, q) is convenient because it allows various moments to be calculated by expectation. For example, the conditional mean of x_i is ψ_i, the conditional duration, but the unconditional mean is

$$E(x_i) = \frac{\omega}{1 - \sum_{j=1}^{m} \alpha_j - \sum_{j=1}^{q} \beta_j}. \tag{10}$$

This is most easily seen by taking expectations of both sides of Eq. (9) although the proof also requires that all roots of an associated difference equation lie outside the unit circle. In the ACD(1, 1) case this requires that $\alpha + \beta < 1$. Similarly, in the ACD(1, 1) case, the conditional variance of x is ψ_1^2 but the unconditional variance is given by σ^2 where

$$\sigma^2 = \mu^2(1 - \beta^2 - 2\alpha\beta)/(1 - \beta^2 - 2\alpha\beta - 2\alpha^2). \tag{11}$$

Since $\alpha > 0$, the unconditional standard deviation will exceed the mean exhibiting 'excess dispersion' as often noticed in duration data sets. From taking repeated expectations, multistep forecasts of durations can be computed directly. That is, the expected duration of the nth transaction can be computed directly from Eq. (9).

Readers who are familiar with the ARCH class of models will immediately recognize the relationship to models of conditional variance. The ACD(1, 1) is analogous to the GARCH(1, 1) and will have many of the same properties. Just as the GARCH(1, 1) is often a good starting point, the ACD(1, 1) seems like a natural starting point. However, as there are many alternative volatility models, there are many interesting possibilities here. For recent surveys on ARCH models and lists of different classes, see Bollerslev et al. (1994), Bollerslev et al. (1992) and Bera and Higgins (1992). The ARCH model was originally introduced by Engle (1982) and GARCH by Bollerslev (1986).

The specifications in Eqs. (7) and (9) can be generalized in many ways. The durations in Eq. (7) are assumed to be conditionally exponential but there are countless ways to relax this restriction. One popular suggestion is the Weibull. Engle and Russell (1995) found that a conditional Weibull density appeared to fit the data better than the conditional exponential density for most models estimated on the stock market data examined in that paper. For the Weibull distribution, the hazard is a slightly more complicated form

$$\lambda\left(t | x_{N(t)}, \ldots, x_1\right) = \left(\Gamma\left(1 + (1/\gamma)\psi_{N(t)+1}^{-1}\right)^{\gamma}\left(t - t_{N(t)}\right)^{\gamma-1}\gamma \tag{12}$$

where $\Gamma(\cdot)$ is the gamma function and γ is the Weibull parameter. The conditional intensity is now a two parameter family which can exhibit either increasing or decreasing hazard functions. This makes especially long durations more or less likely than for the exponential depending on whether γ is greater or less that unity respectively. When $\gamma = 1$, this simplifies to Eq. (7).

In practice, additional variables may well enter the ACD model. Of particular importance is time itself. There are systematic variations in the arrival rate of price quotes over the day and over the week. Thus the conditional expected duration typically depends upon t directly. Thus the specification in Eq. (3) may naturally be generalized to give

$$\psi_i(x_{i-1}, \ldots, x_1; \theta)\Phi(t_{i-1}) = E[x_i | x_{i-1}, \ldots, x_1; \theta] \tag{13}$$

where it is assumed that the seasonal function of time is separable multiplicatively from the stochastic component.

More interestingly, economic variables can enter the equation which determines the frequency of transactions. From this version of the model one can test hypothesis on economic determinants of the rates of transactions. It is here that the information loss due to thinning can be assessed and hypothesis about market microstructure such as price setting behavior with potentially asymmetric information and price leadership are examined.

3. Data

The foreign exchange data were provided by Olsen and Associates (O&A). The data set contains several variables extracted from the Reuters screens including bid–ask quotes and associated arrival times. The foreign exchange market operates around the clock 7 days a week and the complete data set is one year covering October 1, 1992 through September 30, 1993. The typical rate of quote arrivals differs dramatically on weekends and weekdays and between business hours in different countries and in different time zones. In order to utilize days with a common typical pattern, only Tuesdays, Wednesdays and Thursdays are analyzed. This subsample consists of 51 days and 303,408 observations on the Deutschmark–Dollar exchange rate for the months of May through August.

Although the primary focus of the paper is on price based durations (price durations), it is useful initially to describe the quote arrival rate and quote based durations (quote durations). For the O & A data set a typical weekday has almost 6000 quote arrivals. On average a quote arrives every 15 s. Seasonal patterns in the rate of arrival of quotes have been examined in several papers [2]. The first task is to model this rate. Assuming the separability of time function and stochastic function as in Eq. (13), the elimination of the time of day effect is comparable to seasonal adjustment. The strategy followed here is to regress the duration on a function purely of the time of day to obtain a consistent (but inefficient) estimator of the typical shape. Dividing, the durations by their estimated typical shape gives a 'seasonally adjusted' set of durations.

Fig. 1 presents the expected duration conditioned on time-of-day alone. This expectation was formed by regressing the observed duration on 96 time-of-day binary variables, each 15 min long. The day of week was neglected so Tuesdays, Wednesdays, and Thursdays are assumed to have the same seasonal component. Notice that this plot reveals the same patterns as previous studies of quote frequency, although it looks roughly like the inverse. In particular, instead of observing a trough for a lack of quote arrivals during the Japanese lunch hour as observed in Bollerslev and Domowitz (1993), we observe a corresponding spike for particularly long durations between quote updates. Analogous to the previous

[2] See Bollerslev and Domowitz (1993) and Muller et al. (1990) for example.

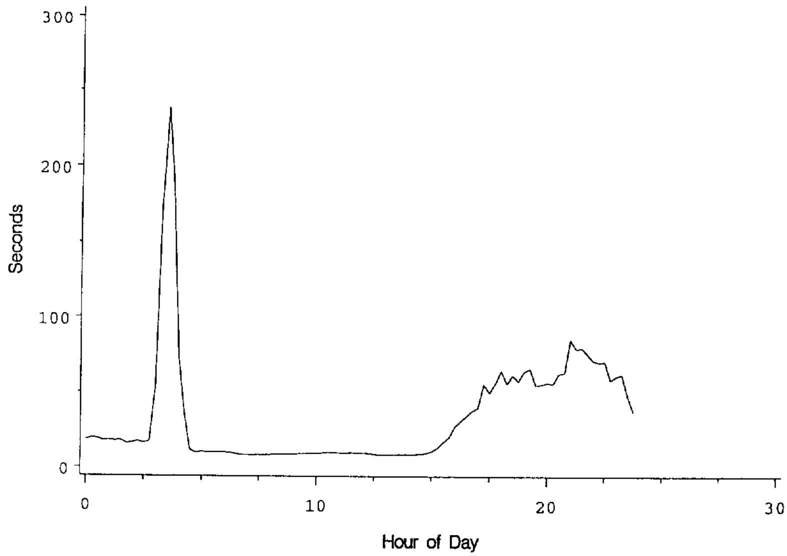

Fig. 1. Expected quote duration conditioned on time of day.

studies, we find the most quote activity in the middle of the day between hours 5 and 14. During this time quotes are arriving at a rate of just under 1 every 9 s on average. At the other extreme, during the Japanese lunch hour (hours 3–4 GMT), quotes arrive at a rate of around 1 every 4 min.

Fig. 2 presents a plot of the histogram of raw durations up to 300 s. Most of the quotes are recorded in intervals of multiples of 6 s because the recording device checks for new quotes every six seconds. In the few exceptions where quotes arrived off multiples of 6 s, the durations were rounded up to the nearest 6 s interval. The majority of the quotes (nearly 60% arrive within 6 s of the previous quote. Furthermore, approximately 97% of the quotes arrive within 24 s of the previous quote.

An interesting feature of the quote durations is the presence of autocorrelation even after partialling out the time-of-day effects. In particular, consider the 'seasonally adjusted' series

$$\tilde{x}_i = \frac{x_i}{\Phi(t_{i-1})} \quad \text{where } \Phi(t_{i-1}) \equiv E(x_i | t_{i-1}). \tag{14}$$

This adjusted series now has a mean of approximately 1. The autocorrelations and partial autocorrelations for the month of May are presented in Table 1 with the Ljung-Box statistic. The Ljung Box is calculated for the seasonally adjusted series using 15 lags. This will have the usual Chi-squared distribution with 15 degrees of freedom. There are 71,557 observations in this sample with a corresponding

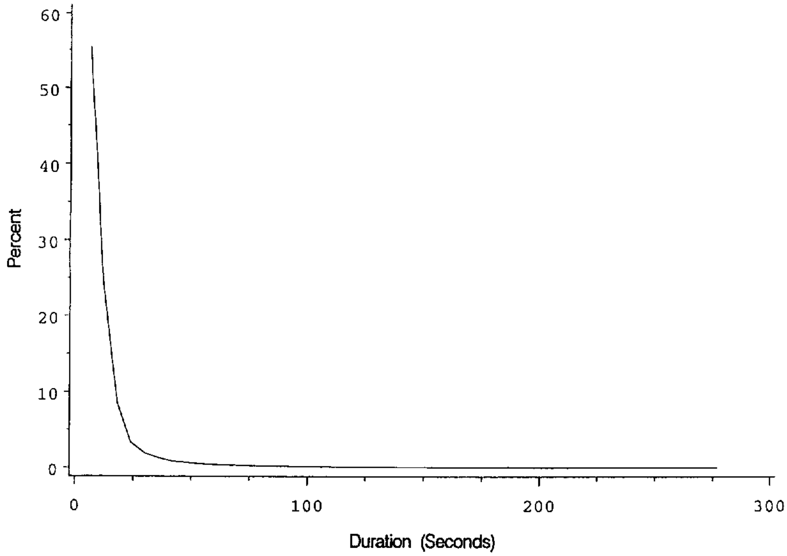

Fig. 2. Histogram of raw quote durations.

Ljung-Box statistic of 3019. The null hypothesis of white noise is easily rejected with the test statistic well above the critical value of 24.99 at the 5% level. The long sets of positive autocorrelations are indicative of duration clustering, suggest-

Table 1
Autocorrelations and partial autocorrelations for price durations

	Quote durations (May)		Price durations (Tues. Wed. Thurs. full sample)	
	acf	pacf	acf	pacf
lag 1	0.083	0.083	0.120	0.120
lag 2	0.076	0.070	0.109	0.096
lag 3	0.064	0.053	0.080	0.058
lag 4	0.053	0.039	0.066	0.042
lag 5	0.059	0.045	0.080	0.057
lag 6	0.048	0.031	0.075	0.049
lag 7	0.050	0.033	0.051	0.021
lag 8	0.038	0.020	0.063	0.036
lag 9	0.048	0.031	0.059	0.032
lag 10	0.040	0.022	0.060	0.031
lag 11	0.049	0.031	0.062	0.032
lag 12	0.043	0.023	0.056	0.025
lag 13	0.039	0.019	0.047	0.016
lag 14	0.040	0.020	0.059	0.030
lag 15	0.042	0.022	0.043	0.011
	Ljung-Box 3019.25		Ljung-Box 1307.58	
	Sample size 71557		Sample size 16277	

ing periods of relatively high quote activity and others of relatively low quote activity beyond a seasonal component [3].

In order to define price durations a further transformation of the data is necessary. A new series for the midpoint of the bid and ask quotes is constructed. Define

$$p_i = \frac{(\text{ask}_i + \text{bid}_i)}{2} \tag{15}$$

This will be referred to as the price series. Now, consider the time series data consisting of N arrival times and N associated prices indexed by i $\{(t_1, p_1), \ldots, (t_N, P_N)\}$. From these pairs the price based arrival times are defined by selectively deleting some of the pairs. It is the history and the current value of the price that determine if the point is deleted. For $i = 2, \ldots, N$, delete the current point if the price is 'close' to the last retained point [4]. In essence only the points at which the price has changed significantly since the occurrence of the last price change are kept. In order to minimize the effects of errant quotes two consecutive points were required to have changed significantly since the last price change. Hence one errant quote will not trigger a price change. A more formal definition of the retained series is as follows:

(i) Retain point 1.

(ii) Retain point $i > 1$ if $\text{abs}(p_i - p_j) > c$ and $\text{abs}(p_{i+1} - p_j) > c$ where j is the index of the most recent retained point, and c is a constant.

Reindexing the retained series by i' for $i' = 1, \ldots, N'$ for $N' < N$ the price durations are then defined as $x_{i'}^{\mathrm{p}}(c) = t_{i'} - t_{i'-1}$ for $i'(c) > 1$. For simplicity the dependence on c will be dropped. As discussed in Engle and Russell (1995) retaining some points of the point process and deleting others is known as thinning a point process. In this case, the thinning is a function of the price marks.

Clearly the value of c is what characterizes a significant price change. If $c = 0$ then we would count every single movement in the midpoint as a price change. We might expect, however, that there will be some movements in the midpoint as a function of individual bank decisions that reflect idiosyncratic portfolio adjustment that should not be considered as a movement in the fundamental price at which individuals are exchanging. In order to better capture movements in the price at which transactions are occurring, we will choose a threshold value of c greater than 0. A histogram for the spread is presented in Fig. 3. Most spreads are 0.0005 accounting for 46% and 0.001 which accounts for 47%. Another 4% are just under 0.00075. To minimize the impact of asymmetric quote setting due to portfolio adjustment by individual banks we will set c = 0.0005 (five pips). This

[3] Although not reported, the autocorrelations were present in various subsamples of the day.

[4] Note that the first point is retained.

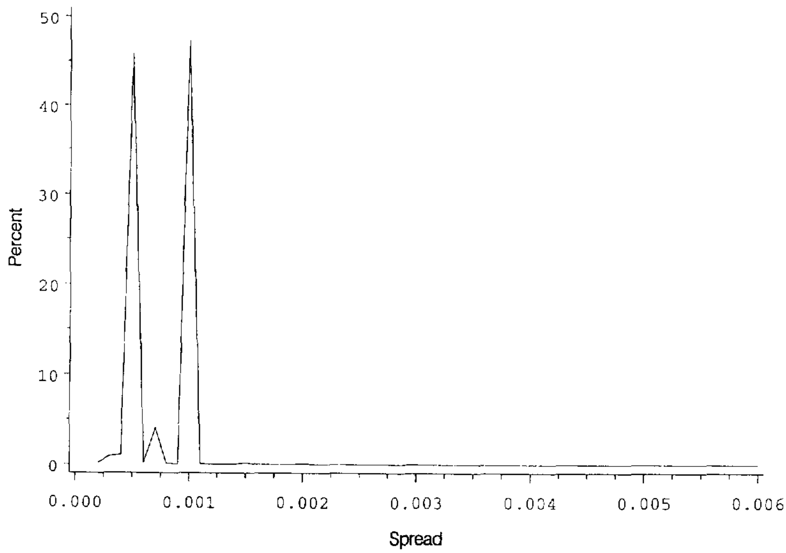

Fig. 3. Histogram of observed spreads.

means that a price change will be triggered only when the midpoint of the quote moves by half the largest observed spread (0.001). This choice of c yields a sample size of 16,277 or 5.36% of the original sample. The minimum price duration is 6 s, the maximum is 11,556 s (or just over 3 h), and the average for the sample is 258 s (or just over 4 min). Fig. 4 presents the histogram for the price durations. Not surprising is the lack of price durations at 6 s. The most common price duration is 18 s accounting for 5% of the of all price durations.

Fig. 5 presents the seasonal component for the price durations. This component was calculated in the same manner as for the quote durations. It is apparent that price changes occur most frequently when the American market overlaps with the European markets between the hours of 12:30 and 15:30 GMT. At this time the price changes on average once every 100 seconds or once every minute and 40 s. When the European market closes, price changes occur much less frequently. They become as infrequent as once every half hour around 22:00 GMT.

The seasonally adjusted series is again formed by partialling out the seasonal component using the same technique as was used for the quote durations. The mean of the new series is approximately unity. The standard deviation is 1.28. The standard deviation for the exponential distribution should be equal to the mean. Relative to the exponential, there is clear excess dispersion in the data. Table 1 presents the autocorrelations and partial autocorrelations for the new series. The long sets of positive autocorrelations are present again indicating price duration clustering. Prices tend to experience periods of rapid movement and periods of

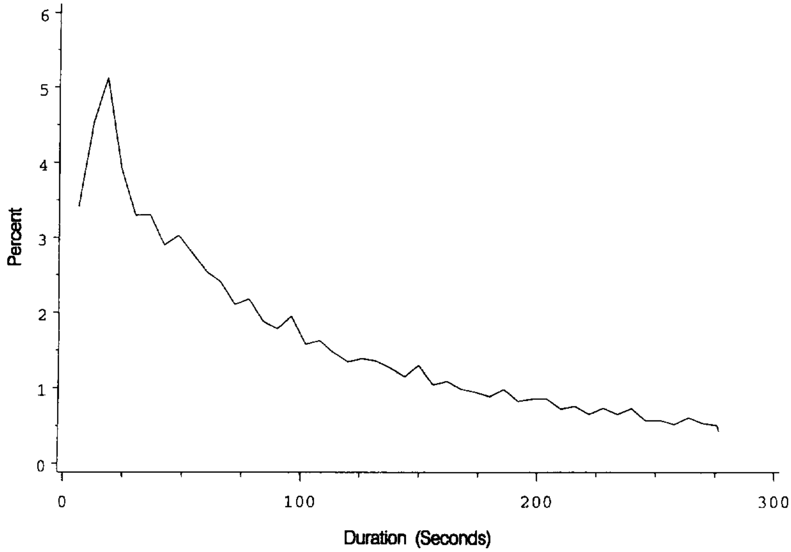

Fig. 4. Histogram of price durations.

slow movement. The autocorrelations are much larger than those for the quote durations (about 1.5 times as large). The Ljung-Box associated with 15 lags is 1044 indicating significant autocorrelations again.

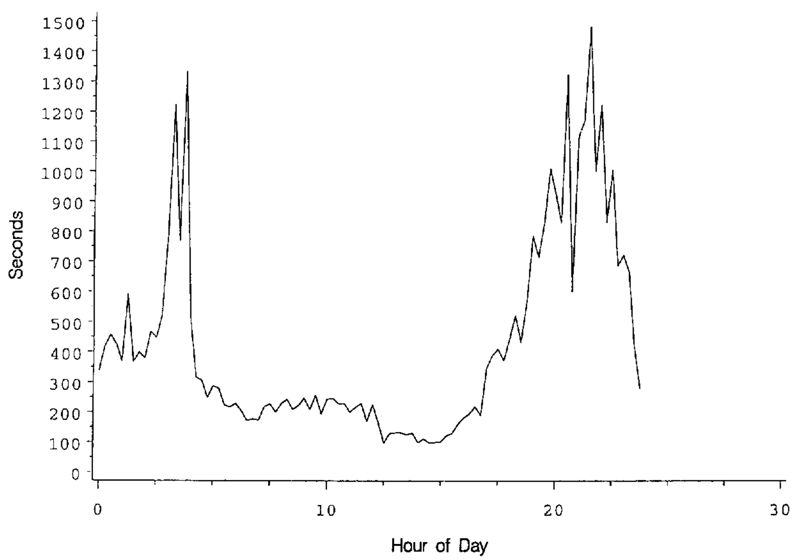

Fig. 5. Expected price duration conditioned on time of day.

The plot of the estimated seasonal component suggests that the price movements change characteristics as markets around the world open and close. A natural question is whether price duration clustering occurs only during certain parts of the day. That is, could a particular portion of the day be responsible for the large autocorrelations. Indeed we find evidence suggesting that the autocorrelations are driven by a particular time of day. The boundaries of when particular markets around the world are 'open' are rather vague. Unlike the NYSE, there is no official opening and closing of the market. Furthermore, active periods in markets around the world overlap. One possible way of dividing the day is to define two periods of the day; a high activity period when Europe and Asia are active from hours 5 through 14, and a low period from hours 17 through 24. This resulted in price duration sample sizes of 8638 and 1938 respectively. The autocorrelations and partial autocorrelations for these two periods, after partialing out the seasonal component, are presented in Table 2. It is apparent that the structure of the autocorrelations is different for these two subperiods of the day. In particular, the low activity period contains several negative autocorrelations corresponding to the 4th, 8th, 9th, 10th, 12th, and 15th lag. The null hypothesis of white noise is clearly not rejected with a test statistic of 15.0. It is apparent that the autocorrelations observed for the day as a whole were generated in the high activity period. The low activity period essentially says that there is nothing to model. Consequently, the model will be estimated on the high activity period only.

Table 2
Autocorrelations and partial autocorrelations for price durations

	Hours 5–14		Hours 17–24	
	acf	pacf	acf	pacf
lag 1	0.123	0.123	0.038	0.038
lag 2	0.098	0.084	0.026	0.025
lag 3	0.098	0.078	0.012	0.010
lag 4	0.068	0.042	−0.017	−0.018
lag 5	0.088	0.063	0.020	0.021
lag 6	0.086	0.056	0.017	0.016
lag 7	0.053	0.020	0.011	0.009
lag 8	0.064	0.033	−0.009	−0.011
lag 9	0.072	0.042	−0.016	−0.016
lag 10	0.075	0.043	−0.020	−0.018
lag 11	0.057	0.021	0.029	0.031
lag 12	0.052	0.018	−0.019	−0.021
lag 13	0.078	0.047	0.045	0.045
lag 14	0.067	0.031	0.004	0.001
lag 15	0.033	−0.007	−0.021	−0.020
	Ljung-Box 772.94		Ljung-Box 15.07	
	Sample size 8638		Sample size 1938	

4. Results

This section will present results for ACD model estimates for the 4 month subsample of price based durations over the active period of the day containing hours 5 through 14. Due to the non linearity of the model, the BHHH algorithm was used with numerical derivatives. The algorithm has no trouble converging for the sample and the results appear robust to initial values imposed.

Parameter estimates for the Weibull ACD (1, 1) (WACD(1, 1)) are presented in Table 3. The top of the figure presents the model estimated for hours 5–14 price durations data with the seasonal portion partialed out. For hours 5–14 the seasonal component appears to be relatively less important than for the day as a whole so estimates for the raw price durations are presented at the bottom of Table 3. With the expected exception of the constant term the models appear very similar. Furthermore, all coefficients for both sets of data are significant according to the t-ratios. The t-ratio for the parameter γ is presented for the null hypothesis of $\gamma = 1$ which would be the case if the true model were exponential ACD. The null hypothesis of an exponential is easily rejected with a t-ratio against this null of 1 of 12.92 and 15.61 for the normalized and raw price durations respectively. $\gamma < 1$ implies that the hazard is decreasing in t. Equivalently, the longer the observed duration of no price change, the less likely a trade will occur at that time.

For a constant unconditional mean to exist lemma 1 of Engle and Russell (1995) requires $(\alpha + \beta) < 1$. For these estimates the sum is approximately 0.9762 for the first sample and 0.9696 for the second suggesting that the process is not integrated. The implied unconditional mean obtained from lemma 1 is 1.01 and 198.42 for the first and second data sets respectively. The actual unconditional

Table 3
Parameter estimates for price durations hours 5–14 GMT.
$\psi_{i'} = \omega + \alpha X_{i'-1} + \beta\psi_{i'-1}$ and γ is the Weibull parameter

	Estimate	Std. error	t-ratio
De-seasonalized price durations			
ω	0.02420	0.00351	6.8789
α	0.07315	0.00487	15.002
β	0.90313	0.00692	130.42
γ	0.90782	0.00713	12.923 *
Raw price durations			
ω	6.0140	0.74590	8.0627
α	0.09807	0.00557	17.584
β	0.87161	0.00766	113.68
γ	0.89137	0.00695	15.618 *

* $H_0: \gamma = 1$.

mean for the second data set is 191.45 with standard deviation of 264. It is useful to examine the standardized series:

$$\hat{\varepsilon} = \left(\frac{\tilde{x}_{i'}^{\mathrm{p}}}{\psi_{i'}} \right)^{\gamma} \tag{16}$$

Correct specification implies that this series should be i.i.d. exponential with $\lambda = 1$ for both data sets. First order conditions imply a mean of one for this series. However, the standard deviation of the series should be unity as well. Although the standard deviation has been reduced from 1.31 to 1.12 there is still evidence of excess dispersion in the first data set. Similarly, for the second data set the standard deviation of the standardized series suggests excess dispersion with a value of 1.15. Correct specification also implies that these standardized series are i.i.d. The Ljung-Box statistics associated with 15 lags are 20.3 and 16.84. This has been reduced from 772.9 and 935 for the first and second data set respectively. In both cases, the null hypothesis of white noise is not rejected at the 5% level. The WACD(1, 1) captures the autocorrelations observed in both data sets. The implied unconditional means are very close to those observed so the models have desirable forecasting properties.

The one-step forecast of the expected price duration is given by

$$E(x_{i'+1}^{\mathrm{p}} | I_{i'}) = E(\tilde{x}_{i'+1}^{\mathrm{p}} | I_{i'}) \Phi(t_{i'}) = \psi_{i'+1} \Phi(t_{i'}). \tag{17}$$

Fig. 6 presents the one-step forecast of the WACD(1, 1) model estimated with a seasonal component. The day was chosen arbitrarily. The dashed line is the

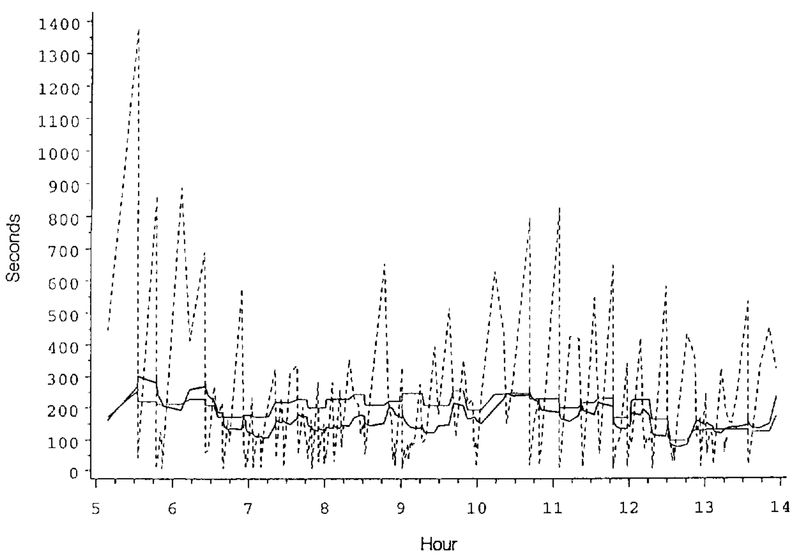

Fig. 6. One-step forecast of price durations, observed durations and seasonal component for WACD(1, 1) for an arbitrary day.

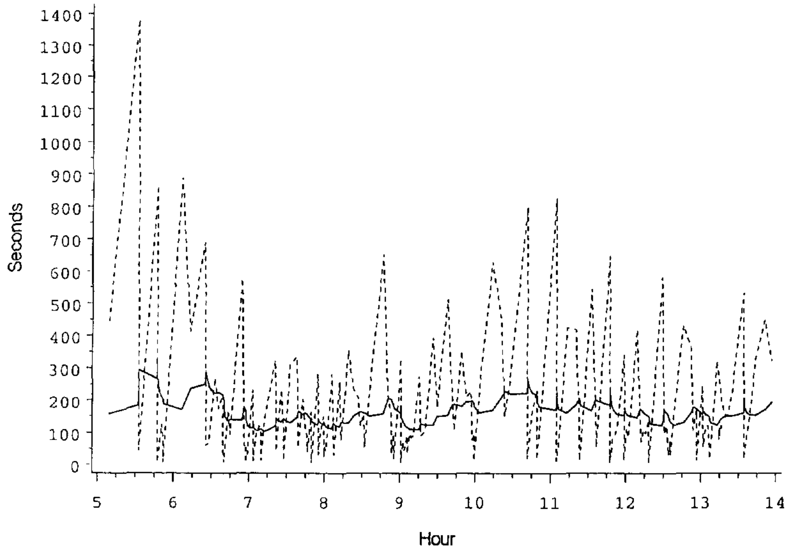

Fig. 7. One-step forecast of price durations and observed durations for WACD(1, 1) for an arbitrary day.

observed durations. The two solid lines are the seasonal spline and the one-step forecast. It is apparent from the graph that the one-step forecast of $\tilde{x}_i^p$ is substantially more variable than the seasonal component. At times the forecast is nearly half the expected value conditioned on time of day alone. In particular, hours 7 through 9:30 exhibit a conditional expectation around 125 s while the seasonal component alone suggests around 225 s.

Fig. 7 presents the one-step forecast plots for the same day but for the data set with no seasonal component. We still observe significant variability in the one-step forecast and the forecasting characteristics are very similar for this period of the day.

5. Testing market microstructure hypotheses

There are several interesting questions that can now be addressed by introducing additional predetermined variables which theoretically could contain information useful in forecasting the price durations. Such hypotheses can easily be tested using LM tests on various variables suggested by theory. The three hypotheses to be examined all relate to the method by which the market participants set prices. One hypothesis is that there is a price leader who commonly signals changes in prices which other bidders follow. A second is that large numbers of quotes between price events signal rapid market adjustment and shorter durations. Third is that market makers concerned about informed traders set the bid ask spread widest when the probability of informed trading is greatest and this is also when they adjust quotes most rapidly leading to short price durations.

It is well known that the LM statistic can be obtained from $N * R^2$ where R^2 is the R-squared from regressing a vector of ones on the scores evaluated under the null hypothesis. The R-squared can be obtained from the first iteration of the BHHH algorithm taken from the maximum likelihood estimates under the null. The test statistic will have the usual chi-squared distribution.

One way of examining the price leadership question is to test if the quote activity of a large player in the market helps to forecast price durations. The quote activity of the largest quote contributor, Deutsche Bank-Asia, is highly autocorrelated for hours 5–14. For this firm we find periods of high quote activity (frequent quote postings) and periods of low activity. Do prices tend to change more rapidly following frequent quote updates from this bank? Deutsche Bank-Asia posted 17,149 quotes between the hours of 5 and 14 in the 4 month sample. The autocorrelations, partial autocorrelations, and the Ljung-Box statistic for the duration between quote updates for Deutsche Bank-Asia are presented in Table 4. The set of large positive autocorrelations yields a Ljung-Box statistic of 3420. The estimated parameters of the quote duration model for Deutsche Bank-Asia are presented in Table 5. For this model the sum of $(\alpha + \beta$ is 0.9233. Excess dispersion is still present in the standardized series having been reduced from 1.51 to 1.30. Although the null hypothesis for the standardized series of white noise is rejected, the Ljung-Box has been greatly reduced from 3420 for the demeaned

Table 4
Autocorrelations and partial autocorrelations for Deutsche Bank-Asia quote durations

Deutsche Bank-Asia quote durations		
	acf	pacf
lag 1	0.275	0.275
lag 2	0.234	0.171
lag 3	0.178	0.087
lag 4	0.068	−0.032
lag 5	0.062	0.007
lag 6	0.039	0.006
lag 7	0.044	0.026
lag 8	0.064	0.043
lag 9	0.073	0.041
lag 10	0.042	−0.007
lag 11	0.063	0.028
lag 12	0.034	−0.005
lag 13	0.048	0.025
lag 14	0.067	0.041
lag 15	0.047	0.010
	Ljung-Box 3420.36	
	Sample size 17149	

Table 5
Parameter estimates for quote durations – Deutsche Bank-Asia
$\psi_i = \omega + \alpha X_{i-1} + \beta\psi_{i-1}$ and γ is Weibull parameter

	Coefficient	Std. error	t-ratio
ω	0.0769424	0.0045514	16.905
α	0.1021910	0.0004780	25.498
β	0.8213549	0.0074811	109.79
γ	1.100348	0.0041622	24.107 *

* $H_0: \gamma = 1$.

series to 47.76 for the standardized series. Plots of the one-step forecast for these durations are presented with the seasonal component in Fig. 8.

Two tests are performed to examine the predictive power of the quote rate of Deutsche Bank-Asia. The first examines the predictive power for the most recent quote duration for Deutsche Bank-Asia prior to the current price duration for the market as a whole. The second test examines if the most recent one-step forecast of Deutsche Bank-Asia's quote activity has predictive ability for the future price durations against the null of a WACD(1, 1).

$x_{i'-1}^{\mathrm{DBA}}$ is the most recent quote duration prior to time $t_{i'}$, for Deutsche Bank-Asia and

$\psi_{i'-1}^{\mathrm{DBA}}$ is the most recent one-step quote duration forecast prior to time $t_{i'}$, for Deutsche Bank-Asia.

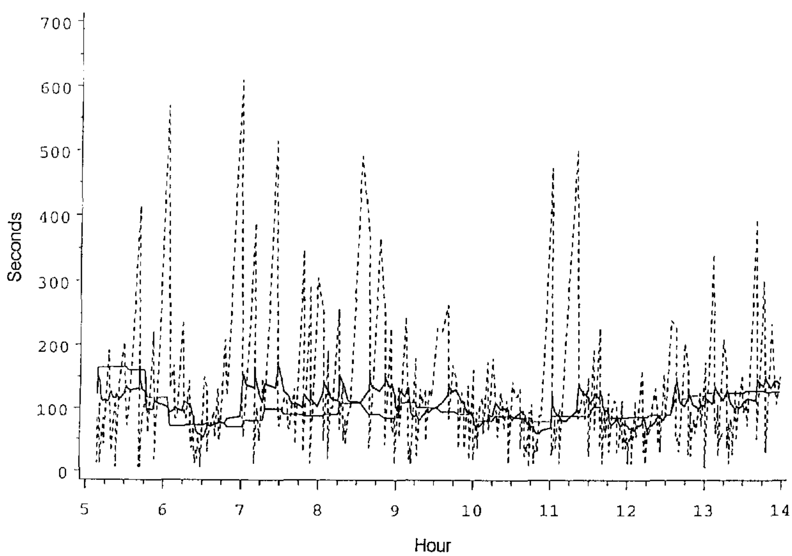

Fig. 8. One-step forecast of quote durations, observed durations and seasonal component for Deutsche Bank-Asia.

Table 6
LM test statistics for predictability of price durations

	$\psi_{i'-1}^{DBA}$	$x_{i'-1}^{DBA}$	$\bar{x}_{i'-1}$	$\bar{s}_{i'-1}$
Test statistic	1.29	0.34	0.881	10.59
Direction	(−)	(−)	(+)	(−)

$\psi_{i'-1}^{DBA}$ is the one-step forecast for the Deutsche Bank-Asia quote duration immediately preceding the price duration.
$x_{i'-1}^{DBA}$ is the Deutsche Bank-Asia quote duration immediately preceding the price duration.
$\bar{x}_{i'-1}$ is the average rate of arrival of quotes over the previous price duration $i'-1$.
$\bar{s}_{i'-1}$ is the average spread observed over price duration $i'-1$.
$\chi^2_{5\%}(1) = 3.84$.

The LM statistics for this section are presented in Table 6. The LM test fails to reject the null hypothesis of no linear relationship for both variables. The test statistics are 1.29 and 0.34 for the duration and forecast respectively. Both statistics are far below the 5% critical value of 3.84. Hence, these tests provide no evidence in favor of price leadership.

The second hypothesis is that quote arrivals carry information about the intensity of the market which has been ignored in the thinned price duration process. These quote arrivals are observed by all agents including the market makers and therefore may reveal the intensity of activity in the market or the rate of information flow [5]. To test this hypothesis, a random variable was constructed as the average arrival rate of quotes over the previous price duration. That is,

$$\bar{x}_{i'} = \frac{1}{n_{i'}} \sum_{k=i_{\min}}^{i_{\max}} x_k \tag{18}$$

where $i_{\min} \equiv i \in \{1, 2, \dots, N(T)\} | t_i > t_{i'-1}$, $i_{\max} \equiv i \in \{2, 3, \dots, N(T)\} | t_i = t_{i'}$ and $n = (i_{\max} - i_{\min})$.

The LM statistic is 0.881 suggesting that lagged durations have no predictive power for price durations. Thus suggests that thinning the quotes to focus only on the price changes loses little information. The deleted quotes appear not to influence economic price setting behavior.

The third hypothesis tests whether bid–ask spreads help predict price durations. There is a vast literature associated with the determination of bid–ask spreads. See for example the excellent discussion in O'Hara (1995). The two general classes of models of price setting are inventory and information models. The former visualize the market makers as risk averse agents who set prices to control their inventories and set spreads to cover transaction costs plus whatever market power they can obtain. Since all the quotes are from potential market makers, the market

[5] This hypothesis has recessed recent attention in stock market literature. See, for example, Jones et al. (1994).

power of each should be minimal. Possibly when fewer quotes are being issued, the market power could be greater and thus spreads could increase. Unless transaction costs vary systematically with volatility, there need be no relation between spread and volatility so at best, the inventory model would predict a weak positive relation between spread and price durations.

The information models such as Glosten and Milgrom (1985) or Easley and O'Hara (1992) on the other hand suppose that there are agents with superior information to that of the risk neutral market makers. Thus the market makers set bid and ask quotes to reflect the possibility of trading with an informed agent. The greater this possibility, the greater the spread. Similarly, if there are many informed traders, the market maker will rapidly adjust both bid and ask prices in response to observed buy or sell orders. Thus the prices will move rapidly and the price durations will be short whenever the market maker has knowledge that he faces high proportions of informed traders. The market makers may well have access to this information through mechanisms such as eloquently discussed by Lyons (1993). Thinking of a price duration as an information event, which is only gradually incorporated into prices, the length of time for price adjustment and the initial bid ask spread should be negatively related.

Yet another strand of literature is the relation between volume and volatility. It is widely found that volume and volatility are closely contemporaneously related. See for example Tauchen and Pitts (1993) Lamoureux and Lastrapes (1990) and Jain and Joh (1988). The larger the volume, the greater the liquidity in the market and consequently, the smaller the required bid ask spread. It is easily seen that high volume stocks have relatively small spreads; presumably this is due either to increasing returns in the transaction technology or increasing numbers of liquidity traders. In either case, this association would predict that bigger spreads would lead to longer price durations.

To test these theories, data on the spread must be constructed. Associated with each price duration is a set of deleted quotes; the measure of bid and ask spread used is the average percent bid–ask spread over all the quote arrivals associated with each particular price duration. More formally,

$$\bar{S}_{i'} = \frac{100}{n_{i'}} \sum_{k=i_{\min}}^{i_{\max}} \log(\text{ask}_k) - \log(\text{bid}_k) \tag{19}$$

where $i_{\max}$, $i_{\min}$ and $n_{i'}$ are defined as above.

The minimum average percent spread is 0.00000118, the maximum average percent spread is 0.24783, with an average of 0.04684 and a standard deviation of 0.01905. The LM test statistic for $\bar{S}_{i'-1}$, is 10.59 which easily rejects the null hypothesis suggesting that prices change more rapidly following larger bid–ask spreads [6]. To follow up on this, a WACD(1, 1) model with the lagged spread was

[6] For the purposes of the test the spread series was multiplied by 10^5 for numerical accuracy.

Table 7
Parameter estimates for price durations with lagged spread
$\psi_{i'} + \alpha X_{i'-1} + \beta\psi_{i'-1} + \eta \bar{S}_{i'-1}$ and γ is Weibull parameter

	Coefficient	Std. error	t-ratio
ω	0.1294009	0.0045514	8.21571
α	0.0852118	0.0004780	18.0392
β	0.8782600	0.0074811	131.027
η	−174.7711	31.09269	−5.62097
γ	0.8920969	0.00692483	15.5820 *

* $H_0 : \gamma = 1$.

estimated. These results are shown in Table 7. The model predicts that a tenth of a percent increase in the average spread will reduce the expectation of the following duration by 17.4 s [7]. This effect is consistent with the asymmetric information model but not with the other two models.

6. A relationship between price durations and volatility

The price duration is a measure of the time per unit price change. Intuitively, this is related to measures of volatility. Rather than measure the expected price change per unit time, the ACD price duration model measures the expected time per unit price change. Clearly the motivation for applying the ACD price duration model to transaction type data is in the irregular spacing of the data. However, it might be interesting to relate the price durations to standard measures of volatility. This section will develop a relationship between the price duration based volatility and standard measures of volatility.

One way of identifying a relationship between price durations and standard measures of volatility is to assume that p_i follows a diffusion process. We can then think of the quote arrivals as a snapshot of the current location of p_i. From this viewpoint the ACD price duration model is a model for the expected time for the diffusion process, given an initial point p_j to escape some symmetric boundary $2c$ units wide. This type of problem has been examined in the context of continuous-time stochastic processes and is called a crossing time problem. If we make the more restrictive assumption that the diffusion parameter is constant within each price duration, but possibly time varying across price durations, then

[7] Some caution has to be used in interpreting this result. There might be a mechanical reason to observe this relationship. If a large spread is observed because the specialist moved either just the bid down or just the ask up then a price change (as measured by the midpoint) will be more likely to occur at this time.

Appendix A shows that the expected crossing time is c^2/σ^2. Matching this to ψ gives an implied diffusion variance of:

$$\tilde{\sigma}_i^2 = \frac{c^2}{\psi_i} \tag{21}$$

Hence can be interpreted as the instantaneous volatility measured in seconds.

An alternative interpretation of Eq. (21) is useful. Instead of assuming that the underlying price process is a diffusion process which is constant until a trade occurs and then shifts to a new rate, it is more natural to think of the prices as being only defined when there is a price setting transaction. Thus the underlying process is a binomial process with increments of $\pm c$ which takes expected time Ψ. Thus the expected variance per unit time is also simply

$$\hat{\sigma}_i^2 = \frac{c^2}{\psi_i} \tag{22}$$

These implied instantaneous volatility measures were calculated for the raw price durations using the estimated conditional durations. It is convenient to adjust the unintuitive units of squared price change per second to the more common percent annual standard deviation by multiplying the $\hat{\sigma}_i$ by the square root of the number of weekday trading seconds in a year. This is approximately the square root of (252 weekdays/year)(24 h/day)(60 min/h)(60 s/min). Essentially, this aggregates each diffusion parameter up to what the annual percent standard

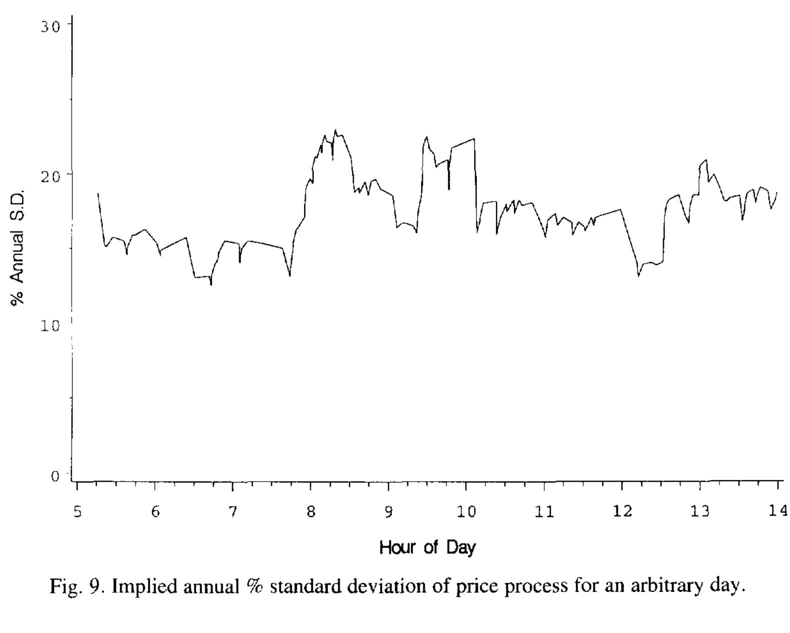

Fig. 9. Implied annual % standard deviation of price process for an arbitrary day.

deviation would be for the entire year if the diffusion parameter were constant for trading on weekdays for the whole year. The volatility over the weekend is ignored in this calculation. Fig. 9 presents the one-step forecast of the implied annualized volatility for an arbitrary day.

7. Conclusions

A measure and forecasts for the intensity of price changes as measured by the midpoint of the bid–ask spread for the Dollar/Deutschmark rate was proposed. Significant autocorrelations of the price durations were observed in this data set which the ACD model of Engle and Russell (1995) was able to successfully model. Both seasonal time of day effects and stochastic effects were observed and modeled. Using a Weibull density for the hazard proved superior to the exponential. This model provides a framework in which the instantaneous probability of events can be forecast.

When the durations are measured in terms of price change events, the ACD model becomes a volatility model. A simple analogy with brownian crossing times, and with a binomial process which has random time intervals, gives a formula for the volatility per unit time. The annualized instantaneous volatility of the FX rates can then be observed and measured over the day and during particularly interesting events.

The model is easily generalized to allow for other exogenous and lagged dependent variables to enter the model formulation. LM tests and model estimates with the significant variables shed new light on the performance of the FX market as revealed by the Reuters quotes. Various theoretical models of price setting behavior in this market suggest particular variables to introduce into the model.

The first model tests the hypothesis that one bank takes the role of price leadership so that their quotes are more influential in determining price movements than others. Lagged quote durations and expected quote durations of Deutsche Bank-Asia, were not significant for the ACD model for the price quotes of the entire market judging from LM tests.

The second hypothesis is that the quote arrival rate carries information useful in addition to the price duration in measuring volatility. The average quote duration during the previous price duration was tested for significance in the ACD. The t-statistic was under 1 indicating no evidence that thinning the process to remove price uninformative quotes sacrificed any efficiency either statistically or economically.

The third hypothesis examines the impact of the past bid ask spread on future volatility. Evidence was uncovered, using LM tests and then by re-estimating the model, suggesting that the lagged average spread, defined by averaging over all quotes received in the previous price duration, helps predict price durations. In particular, a tenth of a percent increase in the average spread reduces the expectation of the following price duration by 17.4 s. This is relative to a sample average near 100 s and a standard deviation of the average percent spread of around 0.02. This result is interpreted as supporting the asymmetric information model of price setting; market makers can assess the probability of facing informed traders and when doing so set big spreads and adjust prices rapidly.

Appendix A

This section develop the relationship between the price durations and volatility used in Section 6.

Following Karatzas and Shreve (1991) let W_t be Brownian motion defined for $0 < t < \infty$ with drift $\mu = 0$ and constant diffusion parameter assumed to be normalized to 1. Karatzas and Shreve propose the problem for Brownian motion defined on the state space [0, a] where a is positive and finite. Let T_a is denote the first passage time defined as

$$T_a = \inf\{t \geq 0;\ W_t = a\}.$$

Let T_0 denote the first passage time for 0. Let $t_{\min}$ denote min(T_0, T_a}. Then for some initial value $W_t = x$ $0 \leq x \leq a$ and $t > 0$, the distribution of $t_{\min}$ is shown to be (Karatzas and Shreve (1991), p. 99)

$$P^x[t_{\min} \in \mathrm{d}t;\ a,\ x] = \frac{1}{\sqrt{2\pi t^3}} \sum_{n=-\infty}^{\infty} \left[(2na + x)\exp\left\{ \frac{(-2na + x)^2}{2t} \right\} + (2na + a - x)\exp\left\{ \frac{(-2na + a - x)^2}{2t} \right\} \right] \mathrm{d}t \quad \text{(A.1)}$$

The Laplace transform of Eq. (A.1) is shown (p. 100) to be

$$E^x[\exp(-\alpha(t_{\min}))] = \frac{\cosh((x - (a/2))\sqrt{2\alpha})}{\cosh((a/2)\sqrt{2\alpha})} \quad \text{(A.2)}$$

Differentiating with respect to α and evaluating at $\alpha = 0$ yields

$$E(t_{\min}) = x(a - x). \quad \text{(A.3)}$$

Setting $a = 2c$ and $x = c$ yields and dividing a and x by the standard deviation yields

$$E(t_{\min}) = \frac{c^2}{\sigma^2} \tag{A.4}$$

Then, at time t_i, $E(t_{\min}) = E(x_i) = \psi_i$. Hence,we get the desired result

$$\sigma_i^2 = \frac{c^2}{\psi_i}. \tag{A.5}$$

References

Bera, A., Higgins, M., 1992. Survey of ARCH Models: Properties, Estimation and Testing, Department of Economics, University of Illinois, unpublished manuscript.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307–327.

Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. The Journal of Finance 48, 1421–1443.

Bollerslev, T., Engle, R., Nelson, D., 1994. ARCH models, in: Engle, R., McFadden, D. (Eds.), Handbook of Econometrics IV, North Holland.

Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH Modeling in Finance. Journal of Econometrics 52, 5–49.

Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. The Journal of Finance 19, 69–90.

Easley, D., Kiefer, N., O'Hara, M., 1994.

Easley, D., Kiefer, N., O'Hara, M., 1995.

Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. Econometrica 50, 987–1008.

Engle, R., Russel, J., 1995. Autoregressive Conditional Duration; A New Model for Irregularly Spaced Time Series Data, University of California, San Diego unpublished manuscript.

Glosten, L.R., Milgrom, P., 1985. Bid ask and transaction prices in a specialist market with heterogenously informed agents. Journal of Financial Economics 14, 71–100.

Goodhart, C.A.E., Hall, S., Henry, S., Pesaran, B., 1993. News effects in a high-frequency model of the Sterling–Dollar exchange rate. Journal of Applied Econometrics 8 (1), 1–13.

Hawkes, A.G., 1971. Spectra of some self-exciting and mutually exciting point processes. Biometrika 58, 83–90.

Jain, P.J., Joh, G., 1988. The dependence between hourly prices and trading volume. Journal of Financial and Quantitative Analysis 23, 269–284.

Jones, J.J., Kaul, G., Lipson, M.L., 1994. Transactions, volume, and volatility. The Review of Financial Studies 7, 631–651.

Karatzas, Shreve, 1991. Brownian Motion and Stochastic Calculus, 2nd Ed., Springer-Verlag, New York.

Lamoureux, C.G., Lastrapes, W.D., 1990. Heteroskedasticity in stock return data: Volume versus GARCH effects. Journal of Finance 45, 221–229.

Lyons, R., 1993. Tests of Microstructure Hypotheses Using Foreign Exchange Markets, Working Paper, Universitv of California at Berkeley.

Muller, U.A., Dacorogna, M.M., Olsen, R.B., Pictet, O.V., Schwarz, M., Morgenegg, C., 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.
O'Hara, M., 1995. Market Microstructure Theory, Basil Blackwell.
Rubin, I., 1972. Regular point processes and their detection. IEEE Trans. on Information Theory 18, 547–557.
Snyder, D.L., Miller, M.I., 1991. Random Point Processes in Time and Space, 2nd Ed., Springer-Verlag.
Tauchen, G., Pitts, M., 1993. The price variability-volume relationship on speculative markets. Econometrica 51, 485–505.