Irish Accounting Review (1999), Vol. 6, No.1, 23–58

AN ANALYSIS OF THE MICROSTRUCTURE OF THE IRISH GILT MARKET IN TWO TRADING REGIMES

Peter G. Dunne

The Queen's University of Belfast

ABSTRACT

This paper concerns the measurement of microstructural characteristics of the Irish Gilt Market in both an 'agency' regime and a market-making regime. A time-series method is developed which, when applied to a market-making system, reveals the effective spread, the information content of trade type, partial adjustment to trade communicated information and trade-type dependence. When applied to the agent-only case, it measures the expected cost of immediacy emanating from an implicit spread and reveals serial dependence in trades at a discount/premium. A surprising aspect of the empirical results is their similarity for the pre- and post-market-making periods.

INTRODUCTION

This paper provides an analysis of a range of 'microstructural' characteristics of the Irish Gilt Market in two trading regimes namely, (i) an 'agency' regime, where the broker acts on behalf of the two parties to a trade and takes no speculative position himself, and (ii) a market-making regime, where the market maker himself stands ready to be the immediate trading partner for whatever type of trader arrives. The difference between the two types of trading arrangement has for some time been of interest to large-scale investors concerned with the speed with which they can liquidate their investments in different markets, as well as to governments wishing to increase the attractiveness of their securities. The field of study which is most concerned with these issues is that of 'market microstructure'.

A subset of the microstructure literature has concentrated on measuring the cost of immediacy of trading (for example, see Roll, 1984). The ability to transact with immediacy can be improved by changing trading regulations, but this is usually achieved at some extra cost to the investor. For example, the market-making trading regime involves a gap between buying and selling prices, the bid-ask spread, which can be quite large for assets that are infrequently traded and may be larger than commission costs in an agency market. The presence of a bid-ask spread on average results in a positive compensation to the market maker for holding an unplanned portfolio position. Interestingly, the quoted bid-ask spread is not an accurate measure of the true cost of trading, or of the compensation to the market maker, in such a market. This is because the quoted prices may not be positioned symmetrically around the underlying value of the asset being traded, i.e. the market maker may want to reduce his inventories and to achieve this he may not charge a potential buyer as much as a potential seller of the asset. The position of the quotes changes from one trade to the next so that an investor who buys and then sells in quick succession may incur an expense which is different from the quoted spread (usually less than it because the previous trade was informative about the underlying value of the asset). Since the cost of immediacy is affected by microstructural characteristics such as the information content of trade and inventory control issues, the Roll approach to measurement of trading cost can provide a basis for the study of these more general microstructural characteristics. So this study generalises the approach used by Roll to measure more than just trading costs in a market-making regime.

Roll's (1984) contribution was to show how transaction price returns (rather than quotes themselves) can be used to measure the implicit or 'effective' bid-ask spread operating in a market-making environment. In particular Roll shows how the presence of a spread affects the covariance between successive transaction price returns (this was for the case where there is no complex microstructural behaviour). The estimated covariance can therefore be used to measure the spread. This measure might be applicable in the case of an agency market. In an agent-only market, the investor can sometimes achieve immediacy in the enactment of his trading wishes by accepting unfavourable prices. Thus he can attract a trading partner by 'giving a good deal'. This means that there is effectively an implicit bid-ask spread in operation. Since this is not quoted, it is a cost of immediacy that cannot be directly observed by investors. However, since the Roll measure of bid-ask spread does not rely on publicly available quotes, but uses transaction prices alone, the same approach could be applied to the measurement of the implicit spread of the agency market.

Despite the likely presence of an implicit spread, other aspects of the cost of immediacy in an agent-only market are hard to tie down. However, it is argued below that non-commission costs of an agency-only system can be discussed and analysed in terms of their relation to costs in a market-making system. This is convenient since most of the empirical microstructure techniques and literature have been concerned with measuring the cost of liquidity in the case of market making.

As mentioned above, in the market-making system the cost of immediacy is measured by what is called the 'effective spread' (Roll, 1984; Stoll, 1989; Jegadeesh and Titman, 1995; and many others). The effective spread involves the average difference between the buying price and selling price when these are separated by the time between trades. Comparison of the effective spread with the quoted spread can reveal information about the components of the spread as described by Stoll (1989) and discussed in detail below. If the quoted and effective spreads are different this must imply that there is predictable movement in the position of the quoted spread between trades. This would be explained by the inventory effect mentioned already. The other main reason for difference in the effective and quoted spread is due to the information content of trade. A large sell may imply adverse private information about the value of the asset and this likelihood is taken into account by the market maker when making the quote. Since it is impossible to know which trades are informed, all trades affect the market maker's belief about the value of the asset. In the absence of other costs (i.e. inventory holding costs and fixed order processing costs), a spread between buying and selling prices would still exist due to this information content of trade. The quoted bid (ask) price would, in this case, be set equal to the market maker's belief about the value of the asset conditional on a sell (buy) by a private investor. Thus, in this special case, if a sell (buy) occurs the new quotes would be set around the previous bid (ask) price since that price reflects the market maker's best guess of the underlying value of the asset given that a sell had occurred. Essentially, Stoll extended the model of Roll to allow for the presence of information trade and an inventory effect. Stoll's analysis remains focused on the use of covariance measures but unfortunately also requires the covariance of successive quote changes.

In the agent-only market the search costs incurred by the broker are covered by commission charged to both the buyer and the seller. However, the additional cost accepted on some occasions by the investor for the sake of immediacy in a non-liquid period should also constitute part of the effective spread implicit in these markets. Indeed, it is this cost rather than commission that causes some large traders to shy away from less developed markets endowing them with an unwanted liquidity premium.

In what follows, a time-series method is developed which is capable of measuring a number of microstructural characteristics that were of interest to Stoll (1989). Fortunately, unlike the analysis of Stoll, only transaction prices are needed in the new approach. Time-series models such as autoregressive moving average models (ARMA)[1] have a direct relation with the covariance properties of transaction price returns under different scenarios regarding trading behaviour. So one can think of the time-series approach as an alternative to covariance based measures of the effects of trading costs under different types of trading behaviour and regulatory regimes. In fact, this approach is most suited to a market-making structure. When applied to the case of a market-making system, the method proposed here reveals the effective spread, the information content of trade type, the partial adjustment to trade communicated information, and the trade-type dependence (which often emanates from regulation that restricts the size of market-maker price movements). When applied to the agent-only case, the method proposed measures the expected cost of immediacy emanating from an implicit spread and (in place of trade-type dependence) reveals the tendency for liquidity shortage to be one-sided (i.e. sales at a discount following sales at a discount rather than random switching between premium and discount). There is no theoretical model for the agent-only case which gives rise to as general a model as for the market-making case. However, the possibility that the general model has relevance to the agent-only case is not ruled out. This will be left for the data to reveal.

The remainder of the paper proceeds with a discussion of the important literature, followed by the development of the empirical model, its application and results. Two approaches to the modelling of microstructural characteristics have arisen: (1) a covariance based method, and (2) a 'time-series' econometric approach. Both methods are based on the use of transaction prices. Below, attention is paid to the translation of the Stoll (1989) covariance approach to a time-series approach. The resulting time-series model is a new extension to the existing time-series literature as it is applied to this subject and it encompasses aspects of several earlier papers in the literature.

The time-series model proposed is an ARMA (2,2) with parameters that reflect speed of price adjustment to news, information asymmetry and trade type dependence effects (and their analogues, if they exist, in the agent-only case). The new model is estimated using transaction data for four gilts which were made available by a Dublin-based stockbroking firm. The results indicate that a less general model is applicable for both the pre- and post-market-making gilt market data. In fact, the data reveal very similar results for the two trading environments. The next section of the paper discusses the relevant literature. This is followed by a proposal for a theoretical extension and presents the various restrictions of the general model and how these are related to the existing literature. The results from the empirical implementation of the new model are then presented before conclusions are drawn.

LITERATURE AND EXTENSIONS

The theoretical microstructure literature has recently been surveyed by O'Hara (1995). Important recent econometric papers on microstructure include various papers by Hasbrouck (1991a, 1991b, 1993) as well as papers by Hausman, Lo and MacKinley (1995), Huang and Stoll (1994), de Jong, Nijman and Roell (1995) and the references therein. Of particular interest here is the adaptation of the Stoll (1989) autocovariance model to the time-series approach pioneered by Hasbrouck and Ho (1987). Roll (1984) presents a basic covariance analysis of transaction price returns, with no trade-type dependence or information content in trade type. The conclusion is that the existence of a spread gives rise to negative covariance of successive transaction price returns and the covariance can be used as a measure of the effective spread.

In the discussion to follow it is assumed that the reader is familiar with some tenets of the time-series econometric literature. In particular, it is important to know that the existence of covariance, or more correctly auto-covariance, between observations at various lags of a time series, such as transaction price returns, implies some time-series representation for the series, such as autoregressive up to lag p, AR (p), or moving average up to lag q, MA (q), or perhaps some mixed process called autoregressive moving average of order p q, ARMA (p,q). A simple example of this is the representation of a series which is first order autoregressive, in terms of an infinite order auto-covariance representation where the auto-covariances are geometrically declining with the lag length. An example which is more pertinent to the present analysis is the Roll (1984) case of a transaction returns series which has a first order auto-covariance structure (in the presence of a symmetric bid-ask spread) which can also be modelled as a moving average regression of order one. Other correspondences between ARMA representations and auto-covariance structures can easily be derived and indeed this is a subject which is covered in some detail in standard texts such as that by Pindyck and Rubinfeld (1991).

Choi, Salandro, and Shastri (1988) retain the focus on the auto-covariance structure but adapt Roll's (1984) model to a situation where there can be trade-type dependence. Stoll (1989) continues this generalisation of the covariance approach and allows the spread to consist of three components: order processing costs (a constant spread); inventory holding costs (implying asymmetry in the position of the spread relative to true price and giving rise to a difference between quoted and effective spread); and lastly, adverse information costs (this is the response to information in trade type and it is measured as a proportion of the quoted spread). George, Kaul, and Nimalendran (1991) further generalise the covariance approach of Stoll (1989) to adjust for movements in expected returns.

Time-series models have generally followed the approach of Hasbrouck and Ho (1987), who model returns as an autoregressive moving average (hereafter ARMA). The only parameter of this ARMA model that can be interpreted for its microstructural information is the AR (1) parameter (revealing trade-type dependence). Hsia, Fuller, and Kao (1994) link the time-series approach to the covariance approach of Roll (1984) and achieve a spread estimator which does not suffer from one of the defects of the Roll spread estimator, namely imaginary estimates. Madhavan and Smidt (1991) provide further extensions to the time-series approach by including Bayesian updating of beliefs and information signalling. They find weak evidence of an inventory effect and convincing evidence for the information effect of trade.

Some Generalisations of Covariance Models

The following notation will be used in the remainder of this section:

S	the quoted or effective spread depending on context.
cov_T	serial covariance of successive transaction price changes (usually in percentage terms).
cov_Q	serial covariance of successive bid or ask price changes (usually as percentage of mid-price).
P_t	transaction price at time t.
q_t	mid-price of quotes at time t.
w_t	underlying value of asset in period t.
δ	information effect of trade type (as proportion of spread) on the mid-price of quotes due to information asymmetry.
θ_i	moving average parameter for lag i.
L	Lag operator.
ε_t	Random binomial process, $\pm c$ (where $abs\{\varepsilon\} = .5S$) serves to shift the transaction price to the ask and bid prices.
π	probability of trade-type reversal. This enters directly into the conditional probabilities governing the series $\{\varepsilon_t\}$, e.g. $Prob(\varepsilon_t = +c)$ given $\varepsilon_{t-1} = +c$ is $(1-\pi)$.

Roll's (1984) effective spread estimator is based on the covariance of successive transaction price changes in the presence of a symmetrically positioned bid-ask spread with random trade-type occurrences. The Roll result is $Cov(\Delta P_t, \Delta P_{t-1}) = -(.25)S^2$. Choi, Salandro and Shastri (1988), referred to as CSS below, provide an adjustment to the Roll estimator to account for trade-type dependence. Their measure is not shown in detail at this point since the details will become clear when the relation between this measure and a time-series alternative is discussed below. Trade-type dependence is particularly interesting when viewed in tandem with information effects of trade. Stoll (1989) extended Roll's work further when he analysed these characteristics together. Stoll shows that, given the quoted spread and the measures of trade-type dependence and information effects, it is then possible to divide the quoted spread into components responsible for the various costs of trading.

The Stoll (1989) approach uses both transaction returns and quote returns. The covariances of these are shown by Stoll to be respectively:

$$S^2\left[\delta^2(1-2\pi)-\pi^2(1-2\delta)\right] \qquad (1)$$

and

$$S^2\left[\delta^2(1-2\pi)\right] \qquad (2)$$

The realised spread earned by the market maker is given by the expected price change following a dealer purchase less the expected price change following a dealer sale also shown by Stoll to be $2(\pi-\delta)S$.

The Stoll (1989) approach constructs estimates of the order processing, inventory holding, and adverse selection components of the spread from slope coefficients of the following simple regressions:

$$cov_T = a_0 + a_1 S^2 + u \qquad (3)$$

$$cov_Q = b_0 + b_1 S^2 + v \qquad (4)$$

The slope coefficients for these regressions are functions of the underlying microstructural parameters of interest as follows:

$$a_1 = \delta^2(1-2\pi)-\pi^2(1-2\delta) \qquad (5)$$

and

$$b_1 = \delta^2(1-2\pi) \qquad (6)$$

The parameters δ and π can be derived from the solving of these two simultaneous equations and can then be used to calculate the components of the quoted spread as:

adverse selection cost	=	$\left[1 - 2(\pi - \delta)\right]S$
inventory holding cost	=	$2\left[\pi - 0.5\right]S$
order processing cost	=	$\left[1 - 2\delta\right]S$

This was really the first serious attempt at measuring the various components of the cost of trading.

George, Kaul and Nimalendran (1991), referred to as GKN below, show that the Stoll (1989) approach is likely to suffer from biases due to time varying expected returns. However, simulations carried out by Afleck-Graves, Hedge and Miller (1994), referred to as AHM below, indicate that this bias in the estimates of the components is small (less than four per cent per component). AHM argue that the alternative method used by GKN is somewhat more useful than the Stoll approach when testing for differences between trading systems in a parametric framework. A drawback of the GKN alternative is that it requires the assumption of a zero inventory cost component.

The auto-covariance structure associated with the transaction return process can help to derive appropriate time-series representations under the various scenarios already analysed using the covariance approach. This idea was put into effect by Hsia, Fuller and Kao (1994), referred to as HFK below, who use the fact that the Roll (1984) covariance structure is equivalent to a moving average process of order one in the transaction returns. In doing so, they arrive at an effective spread estimator which avoids the problem of imaginary spread estimates which gave problems in the application of the Roll measure. As a demonstration and extension of their approach we can consider the conversion of the CSS model to an ARMA(1,1) in which it is possible to get an estimate of trade-type dependence directly from the AR(1) parameter.

The CSS analysis gave rise to the following covariance structure for transaction price changes:

$$\mathrm{Var}(\Delta P_t) = \pi S^2 \qquad (7)$$

$$Cov(\Delta P_t \Delta P_{t+1}) = -\pi^2 S^2 \qquad (8)$$

Although CSS do not investigate higher order auotcovariance it is easy to show that it is described by the following function:

$$Cov(\Delta P_t \Delta P_{t+k}) = -(-1)^k (1-2\pi)^{k-1} \pi^2 S^2 \qquad (9)$$

or

$$Cov_k = (1-2\pi) Cov_{k+1} \qquad (10)$$

This auto-covariance structure corresponds with that which would occur for any ARMA (1,1) process. Consider the general ARMA (1,1), $X_t - \phi X_{t-1} = \varepsilon_t - \theta \varepsilon_{t-1}$. The covariance structure is given as:

$$\mathrm{Var}(X_t) = \frac{(1-2\phi\theta+\theta^2)}{(1-\phi^2)} \sigma_\varepsilon^2 \qquad (11)$$

$$Cov(X_t X_{t+1}) = \frac{(1-\phi\theta)(\phi-\theta)}{(1-\phi^2)} \sigma_\varepsilon^2 \qquad (12)$$

$$Cov(X_t X_{t+k}) = \phi Cov(X_t X_{t+k+1}) \quad k > 1 \qquad (13)$$

Following the same principle as HFK it is possible to solve for the microstructure parameters by equating the alternative covariance statements above and solving simultaneously. The most obvious result is that $\phi = (1 - 2\pi)$. Thus the MA parameter collapses to $\theta = 1$ when there is no trade-type dependence which is the same as the HFK MA parameter. The important point is that the trade-type dependence parameter can be obtained in an easier way than attempted by the CSS covariance approach using only the transaction price data. If a variable is available to proxy the underlying permanent changes in the value of the asset, or *NEWS*, and if there are no information asymmetry effects, then the following regression would yield the parameter of interest:

$$\Delta P_t = \beta \Delta NEWS_t + (1-2\pi) \Delta P_{t-1} + (1-\theta L) \varepsilon_t \qquad (14)$$

The effective spread estimator can once again be derived by solving the covariance equations with some additional information from the regression. The effective spread is now given as $\sqrt{(1/((1-\pi)\pi)} \, \sigma_\varepsilon$.

So this demonstration shows that the CSS assumptions about transaction price movements can lead to a time-series representation from which useful microstructural information can be obtained. In the next section it is shown that the Stoll (1989) approach can also be represented as an ARMA from which additional microstructural information can be obtained. The easiest way to do this is to make an amendment to the time-series model of transaction returns due to Hasbrouck and Ho (1987). Their model is based upon the presence of random public news shocks, lagged adjustment to this news, the presence of a constant spread and trade-type dependence. The assumptions give rise to an ARMA (2,2) where the AR (1) parameter reveals trade-type dependence. The other parameters reflect a combination of lagged adjustment to information surprises and the information content of trade.

THEORETICAL EXTENSION

A more detailed examination of the Hasbrouck and Ho (1987) model of observed returns shows that it can be adapted so that the news component arises through the actual trades themselves, as in Stoll (1989). This leads to a very interesting and tractable outcome which can be represented by an ARMA (2,2) and also contains several of the models in the existing literature as special cases. One of the drawbacks of the Stoll approach is that it relies on a regression of covariances for many stocks on their respective quoted spreads in order to arrive at the effective spreads and the components mentioned above. In this study there are not enough gilts to do such a cross section regression and this in itself provides some motivation for developing a time-series alternative. In any case, the small sample deficiencies of the Stoll spread component estimator have recently been highlighted by Brooks and Masson (1994).

Hasbrouck and Ho (1987) establish three points: (i) transaction based stock returns exhibit statistically significant positive autocorrelation at orders greater than one; (ii) market buy and sell orders do not arrive independently but are in fact characterised by positive dependence (buys tend to follow buys and sells follow sells); and, (iii) returns based on prices constructed as the midpoint of current bid and ask quotes exhibit positive autocorrelation similar to that found for the actual returns. They derive a model that incorporates lagged adjustment in quote-midpoint prices and serial dependence in market orders.

They define the 'true' price of the security that would be observed in a perfect market. They denote this as w_t and it follows a random walk.

$$w_t = w_{t-1} + u_t \qquad (15)$$

u_t is an i.i.d. stochastic process with $E(u_t) = 0$ and $\mathrm{var}(u_t) = \sigma_u^2$. They denote the mid-point of the current bid-ask spread as p_t. Limit orders follow changes in the true price with a delay.

$$p_t = p_{t-1} + \alpha(w_t - p_{t-1}) \qquad (16)$$

The last term implies partial adjustment. If $\alpha = 1$ this implies instantaneous adjustment and $w_t = p_t$. The transaction price which is actually observed will coincide with either the current bid or ask price. The actual transaction price is denoted as π_t.

$$\pi_t = p_t + \varepsilon_t \qquad (17)$$

ε_t is a binomial random variable serving to shift π_t to either the bid or ask.

π_t is below p by $-c$ if at the bid price.

π_t is above p by $+c$ if at the ask price.

The sign of ε_t is determined by whether the order is a buy or sell. Serial dependence in these events will lead to autocorrelation in the ε_t series. To model this trade-type dependence they assume that the bid-ask spread is $2c$.

The possible values for ε_t are $+c$ and $-c$.

prob($\varepsilon_t = +c$) given ($\varepsilon_{t-1} = +c$)

prob($\varepsilon_t = -c$) given($\varepsilon_{t-1} = -c$)

both equal to $\dfrac{1+\gamma}{2}$

If the most recent transaction was a market buy order there is a probability of $(1+\gamma)/2$ that the next transaction will be a market buy order.

Hasbrouck and Ho (1987) then proceed to model observed returns as the first difference of the transaction price π_t to give:

$$r_t = (1-L)\pi_t = (1-L)p_t + (1-L)\varepsilon_t$$

or[2]

$$r_t = (1-(1-\alpha)L)^{-1}\alpha u_t + (1-L)\varepsilon_t \qquad (18)$$

Now ε_t is correlated with ε_{t-1} and this would not be true of moving average errors from the implied regression. Let $\eta_t=(1-\gamma L)\varepsilon_t$ such that η_t is an independent process. This gives:

$$r_t = (1-(1-\alpha)L)^{-1}\alpha u_t + (1-\gamma L)^{-1}(1-L)\eta_t \qquad (19)$$

Rearranging this produces:

$$(1-(1-\alpha)L)(1-\gamma L)r_t = (1-\gamma L)\alpha u_t + (1-(1-\alpha)L)(1-L)\eta_t \qquad (20)$$

This is ARMA (2,2). The left hand side is AR (2) and can be written as:

$$r_t - \phi_1 r_{t-1} - \phi_2 r_{t-2}$$

where the autoregressive parameter estimates will have the following relation with the coefficients of the theoretical model, and can provide estimates of these theoretical coefficients.

$$\phi_1 = (1-\alpha) + \gamma$$
$$\phi_2 = -(1-\alpha)\gamma$$

The right hand side is MA (2) and can be written as

$$e_t - \theta_1 e_{t-1} - \theta_2 e_{t-2}$$

Unfortunately, the empirical estimates of the moving average parameters do not provide any additional information about the theoretical parameters.

The adaptation of this model to account for the information content of order flow is now considered. In contrast to the above model, this gives rise to an ARMA(2,2) which does lead to identification of the theoretical parameters of interest. This is essentially an extension of the work of Stoll (1989) to a time-series context. Stoll investigates the probability of an information asymmetry effect emanating from trade type along the lines of Glosten and Milgrom (1985) who develop the idea of regret-free prices given the previous trade type. One problem with the Stoll (1989) model is that it does not allow for a lagged adjustment to trade type, i.e. the information content of trade type. Such a lagged adjustment is plausible however. It might be more appropriate to react to the general balance of trade type over a longer period of time than the gap between each transaction.

Consider a special case where news only emanates from private information through trades. In this case, regret-free prices without lags in adjustment would imply a price change due to the information in trade type just as in the Stoll (1989) model. This would imply a different news component in the Hasbrouck and Ho (1987) model as follows:

$$w_t = w_{t-1} + \beta\varepsilon_{t-1} \tag{21}$$

Recall that w_t is the 'true' price, or underlying fundamental value, of the security. β is the proportion of the half-spread representing the information effect (β can be thought of as being equal to 2δ in Stoll's notation above). As before, ε_t is the binomial $+c$ or $-c$ which shifts the mid-price to either the bid or ask. The sign of ε_t is then positive ($+c$) for a public buy at the ask price and negative ($-c$) for a public sell at the bid. A buy (sell) would imply good (bad) news and so one would expect β to be a positive value. In the case of an implicit spread rather than quoted spread the error ε_t shifting the price to the implicit ask (bid) for an urgent public buy (sell) is more likely to be normally distributed than binomially distributed. This implies that conventional econometric analysis and inference can be used to implement the model below (for a discussion and analysis of the discrete price situation see Glosten, 1987, and Harris, 1990).

From here the model proceeds as in Hasbrouck and Ho (1987), but the outcome is much more tractable. Recall that p_t is the midpoint of the spread and α is the partial adjustment factor. Hasbrouck and Ho's equation for lagged adjustment is:

$$\begin{aligned} p_t &= p_{t-1} + \alpha(w_t - p_{t-1}) \\ p_t &= [1-(1-\alpha)L]^{-1}\alpha w_t \\ (1-L)p_t &= (1-L)[1-(1-\alpha)L]^{-1}\alpha w_t \\ (1-L)p_t &= [1-(1-\alpha)L]^{-1}\alpha\beta\varepsilon_{t-1} \end{aligned} \qquad (22)$$

The last line comes from substituting in the alternative news component. The return generating process is thus[3]:

$$[1-(1-\alpha)L]r_t = \{1-[2-\alpha(1+\beta)]L+(1-\alpha)L^2\}\varepsilon_t \qquad (23)$$

Since the error contains trade-type dependence we can represent the price changes as AR (1) on this account. This is done in the same way as above by substituting in $\eta_t = (1-\gamma)\varepsilon_t$:

$$(1-\gamma L)[1-(1-\alpha)L]r_t = \{1-[2-\alpha(1-\beta)]L+(1-\alpha)L^2\}\eta_t \qquad (24)$$

This is ARMA (2,2), with the left hand side being the same as in Hasbrouck and Ho (1987), while the right hand side is different. For clarity it is worth displaying this in its full form:

$$r_t = (1-\alpha+\gamma)r_{t-1} + \gamma(\alpha-1)r_{t-2} + \eta_t + [\alpha(1+\beta)-2]\eta_{t-1} + (1-\alpha)\eta_{t-2} \qquad (25)$$

This result is particularly useful as we now have a much more tractable model than Hasbrouck and Ho (1987) and in this case the MA part of the model is interpretable.

There are five special cases of the above equation (assuming that α is trivially equal to 1 when $\beta = 0$, since there is no trade communicated news to react to, and that $\beta \neq 0$ if $\alpha < 1$). The special cases are as follows:

- instantaneous adjustment, $\alpha = 1$, $\beta \neq 0$ and $\gamma \neq 0$. This implies the return generating process is ARMA (1,1) as follows:

$$r_t = \gamma r_{t-1} + \eta_t - (1-\beta)\eta_{t-1} \quad (26)$$

- instantaneous adjustment and no trade-type dependence, $\alpha = 1$, $\beta \neq 0$ and $\gamma = 0$. This implies MA (1) for the return generating process:

$$r_t = \eta_t - (1-\beta)\eta_{t-1} \quad (27)$$

- no information content in trade, $\alpha = 1$, $\beta = 0$ and $\gamma \neq 0$. This again implies ARMA (1,1) for the return generating process but with the MA parameter restricted to -1. This is the Choi, Salandro and Shastri (1988) model where the AR parameter reveals the extent of trade-type dependence.
- $\alpha = 1$, $\beta = 0$ and $\gamma \neq 0$. This is the simple MA (1) with a restricted coefficient of -1 and is exactly the Hsia, Fuller and Kao (1994) alternative to Roll's (1984) simple spread estimator.
- $\alpha < 1$, $\beta \neq 0$ and $\gamma \neq 0$ gives the general model so that the only remaining possibility is $\alpha < 1$, $\beta \neq 0$ and $\gamma = 0$. This gives an ARMA (1,2) as follows:

$$r_t = (1-\alpha)r_{t-1} + \eta_t + [\alpha(1+\beta) - 2]\eta_{t-1} + (1-\alpha)\eta_{t-2} \quad (28)$$

This section has shown that a simple theoretical model which includes a number of microstructural components can be represented by way of a simple time-series regression. The parameter estimates from the implied empirical model can be used to provide estimates of the various behavioural parameters of the theoretical model. The empirical section to follow investigates whether there is support for the theoretical model in the two trading regimes. Since the most general theoretical model may not be distinguishable from some of the special cases mentioned above, the empirical analysis will investigate various alternative empirical representations to see which representation is most suitable and therefore which theoretical model is most plausible.

EMPIRICAL IMPLEMENTATION

Empirical estimation of the ARMA models with a view to recording the microstructural characteristics which have an interpretation if there is an implicit or quoted spread in existence is now proposed. Since the theoretical foundations for the various effects in the 'agent-only' situation are not well developed, caution is urged in interpreting the results too emphatically for that case. The model parameters have well known interpretations, outlined above, in the case of a market-making situation.

The Data

The data set consisted of transactions records spanning eight weeks in the pre-market-making period and 16 weeks for the post-market-making period for the same three gilts. The data included the time, price, amount traded, and day number for each deal for the three gilts, (two medium and one long). **Figures 1** to **3** show, for the pre-change data, the transaction price change in chronological order conditional on the time gap between the trades being less than one hour. **Figures 4** to **6** show this for the post-market-making case. Casual observation indicates that in these short intervals the price does indeed bounce to and fro as though there were an implicit spread in the pre-market-making case. Only one of the gilts (a long gilt) was selected for more rigorous econometric analysis and this was due mainly to the larger absolute number of trades for this gilt and the higher proportion of smaller time gaps between trades in this case.

Figure 1: Pre-Market-Making. Gilt 1 – Price change conditional on time gap <1 hour.

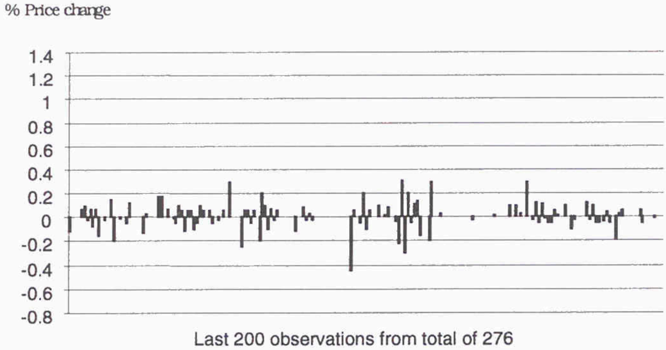

Figure 2: Pre-Market-Making. Gilt 2 – Price change conditional on time gap <1 hour.

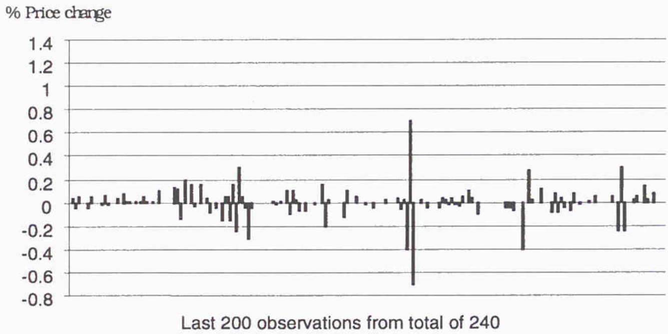

Figure 3: Pre-Market-Making. Gilt 3 – Price change conditional on time gap <1 hour.

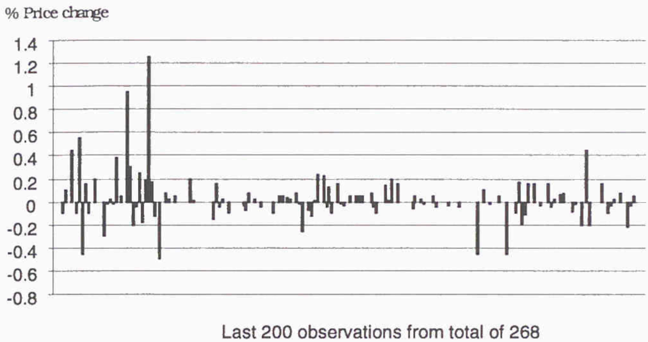

Figure 4: Post-Market-Making. Gilt 1 – Price change conditional on time gap <1 hour.

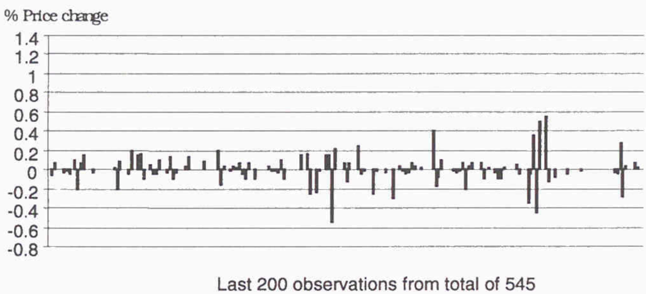

Figure 5: Post-Market-Making. Gilt 2 – Price change conditional on time gap < 1 hour.

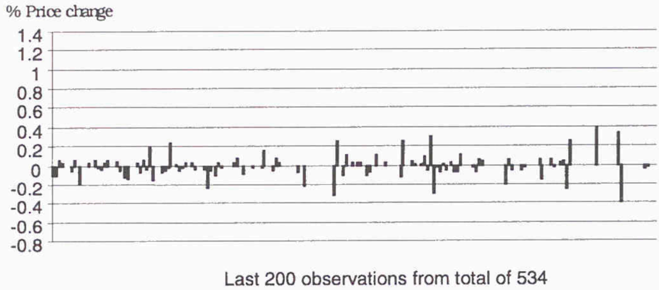

Figure 6: Post-Market-Making. Gilt 3 – Price change conditional on time gap < 1 hour.

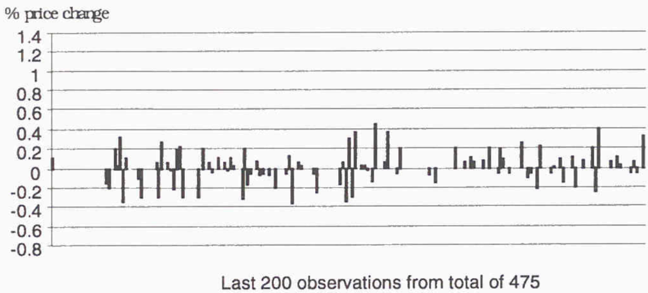

Because of the agency system that existed in the Dublin market, both sides of each transaction were included in the original data set. The second of these deals was therefore omitted. Very small deals, i.e. those of less than £10,000, were also excluded. The price change from one transaction to the next was then calculated and this was the series subjected to empirical analysis. A maximum likelihood estimation procedure was used in which it was possible to model restrictions on the parameters. Initial analysis of the data is presented in the form of the autocorrelogram and partial autocorrelogram with 95 per cent confidence bands (see **Figures 7** and **8** for the pre- and post-market-making cases respectively). These do not indicate the presence of autoregressive or moving-average models with a large number of lags. However, it was decided to begin the empirical modelling with an ARMA (2,2) which was consistent with the most general theoretical model outlined above. This was then restricted to models in which there were fewer autoregressive and moving-average lags. This general-to-specific approach combined with the restrictions implied by the theoretical models will tend to bias results in favour of the model. The results, however, still retain some merit as an exploratory analysis but to keep things in perspective both the unrestricted and restricted results are presented. The results show that the most appropriate empirical representation of the pre-change data is an moving average of order one (meaning a moving average model which contains one lag of the error). This empirical representation has an interpretation as a restriction on the theoretical model proposed above. This interpretation implies that trade does indeed convey information, and that the information effect is a large proportion (0.766) of what can be considered the implicit spread.

Figure 7: Gilt 4 – Pre-Market-Making Autocorrelations and Partial Autocorrelations

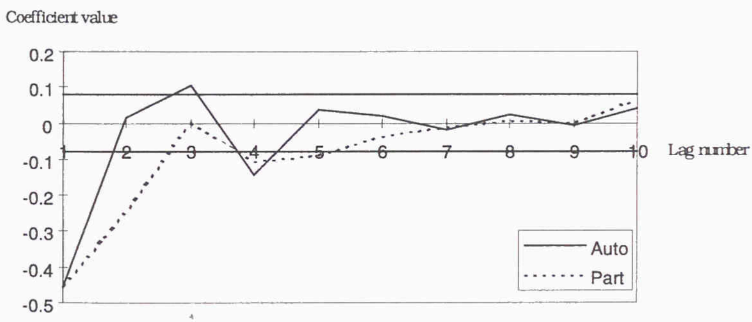

Figure 8: Gilt 4 – Post-Market-Making Autocorrelations and Partial Autocorrelations.

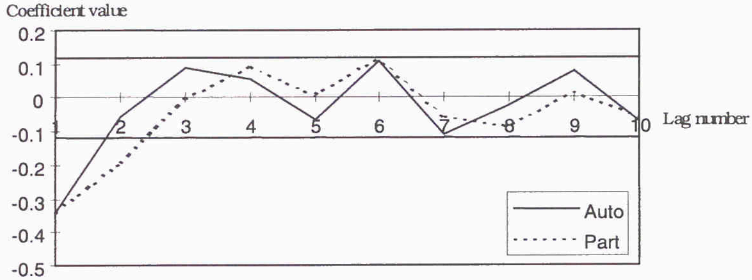

Results

The pre-change data results are presented in **Tables 1** and **2**, while those for the post-change data are presented in **Tables 3** and **4**. Diagnostics on autocorrelation of the residuals is in the form of Ljung-Box statistics. Other tests, not presented, indicated that the residuals were more leptokurtic than in the normal distribution but there was no clear trace of autoregressive conditional heteroscedasticity (ARCH) effects in the data. Also investigated was the possibility that variance was a function of the time gap between trades. This did not have any sizeable effect on the mean equation results and for this reason the simple regression model with constant variance was retained.

Pre-change Results

The results for the regressions **without** microstructure parameter restrictions are presented in **Table 1**. The ARMA(2,2) and ARMA(1,2) both suffer from an exceedingly high MA(1) parameter which is due to unrealistic restrictions. The ARMA(2,1), ARMA(1,1) and MA(1) are virtually indistinguishable from each other in terms of maximised value of the likelihood function and residual diagnostics. Of this group, the ARMA(1,1) and MA(1) are superior on the grounds of a low t-statistic for the MA(2) parameter in ARMA(2,1) specification. If parsimony is given a high priority then the MA(1) model looks like a reasonable specification to choose out of all the unrestricted regressions.

Table 1: Pre-change Regressions without Microstructure Parameter Restrictions (t-statistics are shown in parentheses; L-Box indicates the Ljung-Box statistic).

ARMA(2,2):

$$r_t = \underset{(1.950)}{0.700}\, r_{t-1} + \underset{(0.699)}{0.076}\, r_{t-2} + \eta_t \underset{(-3.599)}{-1.274}\, \eta_{t-1} \underset{(2.129)}{+0.392}\, \eta_{t-2}$$

$Log\ Likelihood = 1404.423.$ $\hat{\sigma}^2 = 0.000001846(18.444).$ $\overline{R}^2 = 0.227283$

$L-Box(1) = 0.114$ $L-Box(2) = 0.276$ $L-Box(3) = 0.497$ $L-Box(6) = 6.936$

ARMA(2,1):

$$r_t = \underset{(-2.582)}{-.757}\, r_{t-1} \underset{(-2.308)}{-0.321}\, r_{t-2} + \eta_t \underset{(0.701)}{+0.212}\, \eta_{t-1}$$

$Log\ Likelihood = 1404.514.$ $\hat{\sigma}^2 = 0.000001844(18.427).$ $\overline{R}^2 = 0.228174$

$L-Box(1) = 0.1298$ $L-Box(2) = 0.8737$ $L-Box(3) = 3.2836$ $L-Box(6) = 6.313$

ARMA(1,2):

$$r_t = \underset{(4.033)}{0.865}\, r_{t-1} + \eta_t \underset{(-6.750)}{-1.422}\, \eta_{t-1} \underset{(4.827)}{+0.499}\, \eta_{t-2}$$

$Log\ Likelihood = 1404.207.$ $\hat{\sigma}^2 = 0.00000185(18.737).$ $\overline{R}^2 = 0.226081$

$L-Box(1) = 0.332$ $L-Box(2) = 0.376$ $L-Box(3) = 0.894$ $L-Box(6) = 7.042$

ARMA(1,1):

$$r_t = \underset{(-1.081)}{-0.1184}\, r_{t-1} + \eta_t \underset{(-4.166)}{-0.449}\, \eta_{t-1}$$

$Log\ \ Likelihood = 1403.93.$ $\hat{\sigma}^2 = 0.00000185(18.83).$ $\overline{R}^2 = 0.230391$

$L-Box(1) = 0.1912$ $L-Box(2) = 0.4139$ $L-Box(3) = 0.7886$ $L-Box(6) = 6.528$

AR (1):

$$r_t = \underset{(-11.825)}{-0.45}\, r_{t-1} + \eta_t$$

$Log\ \ Likelihood = 1396.720.$ $\hat{\sigma}^2 = 0.00000195(18.336).$ $\overline{R}^2 = 0.191198$

$L-Box(1) = 3.944$ $L-Box(2) = 13.3818$ $L-Box(3) = 15.196$ $L-Box(6) = 21.626$

MA (1):

$$r_t = \eta_t \underset{(-12.769)}{-0.533}\, \eta_{t-1}$$

$Log\ \ Likelihood = 1403.381.$ $\hat{\sigma}^2 = 0.00000186(18.730).$ $\overline{R}^2 = 0.230106$

$L-Box(1) = 1.096$ $L-Box(2) = 1.140$ $L-Box(3) = 2.1036$ $L-Box(6) = 7.061$

In **Table 2** the regressions **with** microstructure parameter restriction are presented and it is interesting to note that the restricted ARMA(2,2) fits as well as the unrestricted regression although the gamma parameter is not supported as different from zero. A plausible result is obtained for the partial adjustment parameter (i.e. instantaneous adjustment is supported). In the remaining cases the partial adjustment parameter is restricted to unity. With this restriction and the beta parameter restricted to zero the last two regressions give very unsatisfactory results both in terms of goodness of fit and diagnostics on the residuals.

Thus, one is left with a choice between the ARMA(1,1).a or MA(1).a regressions, both of which are identical to their unrestricted counterparts. Since the t-statistic indicates the absence of trade-type dependence, we can conclude that the most appropriate model is the MA(1).a. Although this is consistent with an information effect through the beta parameter, it is actually an unrestricted MA(1). Since beta is equal to a proportion of half the implicit spread it is clear that the information content of the spread is half beta or (0.7665). As described above, this implies that trades convey information and that this is reflected in the implicit effective spread. This is much larger that the information proportion found using Stoll's covariance technique in the market-making context.

Table 2: Pre-change Regressions with Microstructure Parameter Restrictions (t-statistics are shown in parentheses; L-Box indicates the Ljung-Box statistic).

Recall that the basic model is as follows:

$$r_t = (1-\alpha+\gamma)r_{t-1} + \gamma(\alpha-1)r_{t-2} + \eta_t + [\alpha(1+\beta)-2]\eta_{t-1} + (1-\alpha)\eta_{t-2}$$

ARMA (2,2): $\alpha = 1.474(29.538)$, $\beta = 0.158(1.222)$, $\gamma = 1.1778(6.307)$.

$$\Rightarrow r_t = -0.011\ r_{t-1} - 0.0033\ r_{t-2} + \eta_t + 0.5563\ \eta_{t-1} - 0.063\ \eta_{t-2}$$

$Log\ Likelihood = 1403.975$. $\hat{\sigma}^2 = 0.00000185(12.2889)$. $\overline{R}^2 = 0.190978$

$L-Box(1) = 0.161$ $L-Box(2) = 0.361$ $L-Box(3) = 0.5787$ $L-Box(6) = 6.5865$

ARMA (1,1).a: $\alpha \equiv 1$ $\beta = 1.449(13.44)$, $\gamma = 0.1184(1.084)$.

This is the same as the unrestricted ARMA (1,1) in table 1 above.

MA (1).a: $\alpha \equiv 1$ $\beta = 1.533(36.718)$, $\gamma \equiv 0$

This is the same as the unrestricted MA (1) in table 1 above.

ARMA (1,1).b: $\alpha \equiv 1$ $\beta \equiv 0$ $\gamma = 1.0079(181.47)$.

$Log\ Likelihood = 1362.045$. $\hat{\sigma}^2 = 0.00000252(18.864)$. $\overline{R}^2 \cong 0$

$L-Box(1) = 61.172$ $L-Box(2) = 61.463$ $L-Box(3) = 63.32$ $L-Box(6) = 68.244$

MA (1).b: $\alpha \equiv 1$ $\beta \equiv 0$ $\gamma \equiv 0$

$Log\ Likelihood = 663.07$. $\hat{\sigma}^2 = 0.0004388(11.58)$. $\overline{R}^2 \cong 0$

$L-Box(1) = 271.62$ $L-Box(2) = 536.529$ $L-Box(3) = 794.74$ $L-Box(6) = 1533.25$

Post-change Results

The unrestricted results are presented in **Table 3**. The ARMA (2,2) has an MA (1) parameter which exceeds unity in absolute value and this has no sensible interpretation. As in the pre-change data results, one would have to conclude that this stems from an over restrictive regression. In the pre-change results the ARMA (1,2) regression also suffered from this problem but this is not repeated for the post-change ARMA (1,2). The ARMA (2,1) has all autoregressive parameters insignificantly different from zero which implies that a reduction of the order of autoregression is warranted.

The AR (1) regression does not give satisfactory diagnostics for the residuals and has a maximised likelihood of 4.438 below its nearest rival specification. Based on likelihood ratio tests, this specification cannot be accepted over any of the other alternatives at the 95 per cent level of significance. The likelihood ratio statistic for the ARMA (1,2) against the MA (1) is 3.296 and is below the 95 per cent critical value of the chi-square (2) distribution. Thus, the more general model would not normally be accepted over the more parsimonious alternative. In fact, the choice between the ARMA (1,2), ARMA (1,1) and MA (1) can only be made on the basis of a parsimony argument. Thus the choice of the MA (1) in this case is not as clear-cut as for the pre-change data.

Unfortunately, the results from the post-change restricted regressions, as can be seen in **Table 4**, do not provide extra support for the presence of complex microstructural effects in the market-making data. The first specification, ARMA (2,2), gives rise to insignificant t-statistics for all parameters. The last two specifications, ARMA (1,1).b and MA (1).b, can be rejected on the grounds of the poor residual diagnostics and low maximised log-likelihood. The remaining regressions are not significantly different from each other but the MA (1).a is more appropriate both on parsimony grounds and the fact that the gamma parameter is not clearly significant in the alternative regression. The results can be interpreted as implying that there is a similar information effect from trade in both trading environments.

Table 3: Post-change Regressions without Microstructure Parameter Restrictions (t-statistics are shown in parentheses; L-Box indicates the Ljung-Box statistic).

ARMA (2,2):

$$r_t = \underset{(4.779)}{0.782}\, r_{t-1} + \underset{(1.604)}{0.140}\, r_{t-2} + \eta_t \underset{(-7.706)}{-1.235}\, \eta_{t-1} \underset{(2.347)}{+0.269}\, \eta_{t-2}$$

Log Likelihood = 2602.174. $\hat{\sigma}^2 = 0.000001918(19.894)$. $\overline{R}^2 = 0.1701$

$L-Box(1) = 0.0067$ $L-Box(2) = 0.076$ $L-Box(3) = 0.176$ $L-Box(6) = 7.569$

ARMA(2,1):

$$r_t = \underset{(0.350)}{0.0746}\, r_{t-1} \underset{(0.786)}{+0.0805}\, r_{t-2} + \eta_t \underset{(-2.521)}{-0.521}\, \eta_{t-1}$$

Log Likelihood = 2601.425. $\hat{\sigma}^2 = 0.000001924(19.95)$. $\overline{R}^2 = 0.1675$

$L-Box(1) = 0.0000$ $L-Box(2) = 0.0133$ $L-Box(3) = 0.0171$ $L-Box(6) = 8.111$

ARMA(1,2):

$$r_t = \underset{(-6.017)}{-0.875}\, r_{t-1} + \eta_t \underset{(2.732)}{+0.429}\, \eta_{t-1} \underset{(-3.931)}{-0.335}\, \eta_{t-2}$$

Log Likelihood = 2602.556. $\hat{\sigma}^2 = 0.000001925\ (19.254)$. $\overline{R}^2 = 0.1671$

$L-Box(1) = 0.0105$ $L-Box(2) = 0.1262$ $L-Box(3) = 0.2523$ $L-Box(6) = 5.509$

ARMA (1,1):

$$r_t = -0.0913\ r_{t-1} + \eta_t - 0.356\ \eta_{t-1}$$
$$(-1.036) \qquad (-4.292)$$

$Log\ Likelihood = 2601.26.$ $\hat{\sigma}^2 = 0.000001925(13.23).$ $\overline{R}^2 = 0.1671$

$L - Box(1) = 0.0003$ $L - Box(2) = 0.0124$ $L - Box(3) = 0.2993$ $L - Box(6) = 8.407$

AR (1):

$$r_t = -0.39\ r_{t-1} + \eta_t$$
$$(-11.461)$$

$Log\ Likelihood = 2596.47.$ $\hat{\sigma}^2 = 0.000001962(13.27).$ $\overline{R}^2 = 0.1506$

$L - Box(1) = 1.075$ $L - Box(2) = 10.256$ $L - Box(3) = 10.256$ $L - Box(6) = 19.375$

MA (1):

$$r_t = \eta_t - 0.432\ \eta_{t-1}$$
$$(-13.09)$$

$Log\ Likelihood = 2600.908.$ $\hat{\sigma}^2 = 0.000001928(13.39).$ $\overline{R}^2 = 0.1658$

$L - Box(1) = 0.1215$ $L - Box(2) = 0.942$ $L - Box(3) = 1.0945$ $L - Box(6) = 9.250$

Table 4: Post-change Regressions with Microstructure Parameter Restrictions (t-statistics are shown in parentheses; L-Box indicates the Ljung-Box statistic).

Recall that the basic model is as follows;

$$r_t = (1-\alpha+\gamma)r_{t-1} + \gamma(\alpha-1)r_{t-2} + \eta_t + [\alpha(1+\beta)-2]\eta_{t-1} + (1-\alpha)\eta_{t-2}$$

ARMA(2,2): $\alpha = 0.9771(0.071)$, $\beta = 1.346(0.1995)$, $\gamma = -0.177(-0.003)$.

$$\Rightarrow r_t = -0.154\, r_{t-1} + 0.0041\, r_{t-2} + \eta_t + 0.2922\, \eta_{t-1} + 0.023\, \eta_{t-2}$$

$Log\ Likelihood = 2601.269$. $\hat{\sigma}^2 = 0.00000193(13.111)$. $\overline{R}^2 = 0.1649$

$L-Box(1) = 0.000$ $L-Box(2) = 0.008$ $L-Box(3) = 0.285$ $L-Box(6) = 8.391$

ARMA(1,1).a: $\alpha \equiv 1$ $\beta = 1.3558(16.34)$, $\gamma = -0.0912(-1.036)$.

This is the same as the unrestricted ARMA(1,1) in table 3 above.

MA (1).a: $\alpha \equiv 1$ $\beta = 1.432(43.40)$, $\gamma \equiv 0$

This is the same as the unrestricted MA (1) in table 3 above.

ARMA (1,1).b: $\alpha \equiv 1$ $\beta \equiv 0$ $\gamma = -0.999(-254.28)$.

$Log\ Likelihood = 2555.027$. $\hat{\sigma}^2 = 0.00000231(14.95)$. $\overline{R}^2 = 0.00048$

$L-Box(1) = 77.78$ $L-Box(2) = 79.17$ $L-Box(3) = 80.67$ $L-Box(6) = 94.55$

MA (1).b: $\alpha \equiv 1$ $\beta \equiv 0$ $\gamma \equiv 0$

$Log\ Likelihood = 1280.525$. $\hat{\sigma}^2 = 0.0003635(13.80)$. $\overline{R}^2 = -156.283$

$L-Box(1) = 505.69$ $L-Box(2) = 1004.58$ $L-Box(3) = 1496.40$ $L-Box(6) = 2923.55$

CONCLUSION

A time-series representation of transaction returns was proposed which was based on three microstructural parameters. In the market-making context these parameters represent trade-type dependence, information content of trade type and partial adjustment to news communicated through trade type. It is suggested that in the context of an 'agent-only' system with an implicit spread for immediacy, a subset of these parameters can be interpreted in a similar way to that of the market-making context (i.e. as representing the information content of trade, the tendency for trades at a discount (premium) to follows trades at a discount (premium), and partial adjustment to news in order flow). The model proposed was ARMA (2,2) with some special cases such as ARMA (2,1) representing the case of no trade-type dependence and MA (1) representing instantaneous adjustment to trade-communicated news and (in the agent-only case) equal likelihood of trade at a premium (discount) following a trade at a premium (discount).

The model chosen as most appropriate for both the pre- and post-market-making data was MA (1) with no trade-type dependence, but containing significant information effects from trade. Caution should be exercised when interpreting the beta parameter as the theoretical framework for its role in the agent-only environment is not yet well defined. However, it does indicate that a large proportion of the implicit spread can be explained by some sort of trade impact effect. Thus, the price movement is not simply a random bounce between symmetrically positioned bid and ask prices. Rather, there is evidence that trades are, on average, followed immediately by quote revisions that affect the next trade price.

The most surprising aspect of the empirical results is the similarity between those before and those after the introduction of market making. The two sets of results are very like what would be expected from a market-making arrangement. This must imply that there was a definite spread in operation in the pre-market-making trading environment. The market did in that period operate on the basis of soft-quotes and this is probably the source of the large and significantly negative MA (1) parameter in the unrestricted MA regressions.

NOTES

[1] In the time-series literature an observed series is assumed to be the result of some data generating process (DGP), and the statistical analysis of the observed series is designed to reveal what the generating process is. The main building block of the DGP is uncorrelated innovations. The observed series may have resulted from a weighted average of a number of consecutive uncorrelated innovations. In this case the econometrician has some chance of finding the DGP by fitting a moving average model (MA) to the observed series. If the DGP involved the weighted average of three consecutive innovations then an MA (3) would most likely fit the observed data best. The observed series could have been produced by the past value of the observed series plus a random innovation. In this case an autoregression (AR) with one lag of the observed series would probably be the best fitting empirical model. The time-series approach generally concerns the appropriate selection of the number of lags of either moving average or autoregressive terms (or both ARMA) in an empirical modelling of the observed series.

[2] Recall that $p_t = p_{t-1} + \alpha(w_t - p_{t-1}) = (1-(1-\alpha)L)^{-1}\alpha w_t$ so that $r_t =$ $(1-(1-\alpha)L)^{-1}\alpha(1-L)w_t + (1-L)\varepsilon_t$. Furthermore, since $w_t = w_{t-1} + u_t$ or using the lag operator $(1-L)w_t = u_t$ the returns can be written as equation (18).

[3] Since $r_t = [1-(1-\alpha)L]^{-1}\alpha\beta\varepsilon_{t-1} + (1-L)\varepsilon_t$,
then a slight rearrangement gives
$[1-(1-\alpha)L]r_t = \alpha\beta\varepsilon_{t-1} + [1-(1-\alpha)L](1-L)\varepsilon_t$
and the right hand side of this in turn can be rearranged as
$\{\alpha\beta L + [1-(1-\alpha)L](1-L)\}\varepsilon_t$ or $\{1-[2-\alpha(1+\beta)]L + (1-\alpha)L^2\}\varepsilon_t$.

REFERENCES

Afleck-Graves, J., Hedge, S.P. and Miller R.E. (1994). 'Trading Mechanisms and the Components of the Bid-Ask Spread', *Journal of Finance*, Vol. XLIX, No. 4, pp. 1471-1488.

Brooks, R. and Masson, J. (1994). 'Small Sample Properties of Stoll's Spread Components Estimator', *University of Ottawa, Working paper No. 94-48.*

Choi, J.Y., Salandro, D. and Shastri, K. (1988). 'On the Estimation of Bid-Ask Spreads: Theory and Evidence', *Journal of Financial and Quantitative Analysis*, Vol. 23, No. 2, pp. 219-230.

de Jong., F., Nijman, T. and Roell, A. (1995). 'Price Effects of Trading and Components of the Bid-Ask Spread on the Paris Bourse', *LSE Financial Markets Group, Discussion paper No. 207.*

George, T.J., Kaul, G. and Nimalendran, M. (1991). 'Estimation of the Bid-Ask Spread and Its Components: A New Approach', *The Review of Financial Studies*, Vol. 4, No. 4, pp. 623-656.

Glosten, L. (1987). 'Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices', *Journal of Finance*, Vol. XLII, No. 5, pp. 1293-1307.

Glosten, L.R. and Milgrom, P.R. (1985). 'Bid, Ask and Transaction Prices in a Specialist Market With Heterogeneously Informed Traders', *Journal of Financial Economics*, Vol. 14, pp. 71-100.

Harris, L. (1990). 'Statistical Properties of the Roll Serial Covariance Bid-Ask Spread Estimator', *Journal of Finance*, Vol. XLV, No. 2, pp. 579-590.

Hasbrouck, J. (1991a). 'Measuring the Information Content of Stock Trades', *Journal of Finance*, Vol. 46, No. 1, pp. 179-207.

Hasbrouck, J. (1991b). 'The Summary Informativeness of Stock Trades: An Econometric Analysis', *The Review of Financial Studies*, Vol. 4, pp. 571-595.

Hasbrouck, J. (1993). 'Assessing the Quality of a Security Market: A New Approach to Transaction Cost Measurement', *The Review of Financial Studies*, Vol. 6, No. 1, pp. 191-212.

Hasbrouck, J., and Ho, T.S.Y. (1987). 'Order Arrival, Quote Behaviour, and the Return-Generating Process', *Journal of Finance*, Vol. XLII, No. 4, pp. 1035-48.

Hausman, J., Lo, A. and MacKinley A.C. (1992). 'An Ordered Probit Analysis of Transaction Stock Prices', *Journal of Financial Economics*, Vol. 31, pp. 319-380.

Hsia, C., Fuller, B.R. and Kao, G.W. (1994). 'A Modified Method for Inferring the Effective Bid-Ask Spread form Security Returns', *Journal of Business Finance and Accounting*, Vol. 21, No. 2, pp. 243-253.

Huang, R.D. and Stoll, H.R. (1994). 'Market Microstructure and Stock Return Predictions', *The Review of Financial Studies*, Vol. 7, No. 1, pp. 179-213.

Jegadeesh, N. and Titman, S. (1995). 'Short-Horizon Return Reversals and the Bid-Ask Spread', *Journal of Financial Intermediation*, Vol. 4, pp. 116-132.

Madhavan, A. and Smidt, S. (1991). 'A Bayesian Model of Intraday Specialist Pricing', *Journal of Financial Economics*, Vol. 30, pp. 99-134.

O'Hara, M. (1995), *Market Microstructure Theory*, Cambridge, Mass.

Pindyck, R.S. and Rubinfeld, D.L. (1991). *Econometric Models and Economic Forecasts*, 3rd ed. New York, McGraw-Hill.

Roll, R. (1984). 'A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market', *Journal of Finance*, Vol. XXXIX, No. 4, pp. 1127-1139.

Stoll, H.R. (1989). 'Inferring the Components of the Bib-Ask Spread: Theory and Empirical Tests', *Journal of Finance*, Vol. XLIV, No. 1, pp. 115-134.