Journal of Banking & Finance 26 (2002) 795–818
Journal of BANKING & FINANCE
www.elsevier.com/locate/econbase

The economic and statistical significance of spread forecasts: Evidence from the London Stock Exchange

Nicholas Taylor

Department of Accounting and Finance, Cardiff Business School, Cardiff University, Cardiff CF10 3EU, UK

Received 28 October 1999; accepted 29 January 2001

Abstract

This paper measures the economic and statistical significance of econometric forecasts of bid–ask spreads. The economic importance of these forecasts is assessed by considering the benefits of scheduling trades based on these forecasts. The unrestricted vector autoregression (VAR) model of Huang and Masulis [Rev. Financial Studies 12 (1999) 61] and the two-equation structural model of Huang and Stoll [Rev. Financial Studies 7 (1994) 179] are used to generate intraday *h*-step ahead forecasts of spreads for 50 stocks listed on the London Stock Exchange (LSE). The period corresponding to the minimum expected spread is then scheduled into the trading activity of the investor. The results indicate that when the unrestricted VAR model is used, the spreads incurred are around 35% lower than the spreads incurred by investors who do not schedule their trades. By contrast, spread discounts of only 5% are obtained when the two-equation structural model is used. The heterogeneity of the economic importance of the spread forecasts generated by the models is confirmed by tests of the statistical significance of the forecasts. © 2002 Elsevier Science B.V. All rights reserved.

JEL classification: G14; C32
Keywords: Bid–ask spreads; Economic significance; Forecasting

PII: S0378-4266(01)00173-X

1. Introduction

It is well known that the quality of any scientific theory can only be assessed by considering its predictive power. This often takes the form of a test of the statistical quality of the forecasts generated by the model in question. However, when considering financial models this testing framework should be extended to incorporate the economic significance of the forecasts generated by such models. In this paper, we consider the quality of two econometric models of bid–ask spreads. The economic significance of such models is assessed by examining the quality of a trading schedule based on the models' forecasts. Any spread reduction achieved through use of such a trading schedule is vitally important given the current emphasis, by organisers of financial markets, on transaction cost reduction.

One of the main objectives of stock exchanges is to reduce the costs associated with trading securities. To this end, institutions spend enormous amounts of time and money in an attempt to reduce such costs. For example, the London Stock Exchange (LSE) has recently introduced an electronic trading system that has led to a reduction in bid–ask spreads (Gemmill, 1998; Naik and Yadav, 1999). In this context, good forecasts of future spreads are likely to be valuable as they permit further transaction cost reduction through trade scheduling. Given this importance, it is surprising that no studies have been conducted that have assessed the quality of the forecasts generated by alternative models of spreads. This paper attempts to fill this gap in the literature by assessing the economic and statistical significance of spread forecasts generated by various econometric models.

Two models of spreads are considered in this paper. The first is an augmented version of the two-equation structural model of Huang and Stoll (1994). Using various microstructure theories they develop a two-equation model of mid-point quote and transaction price changes. [1] As the effective half-spread is the difference between the mid-point quote and the transaction price, then this model can be used to forecast future spread levels. The second model considered is an informationally enriched version of the unrestricted vector autoregressive (VAR) model of Huang and Masulis (1999). This model assumes that the touch half-spread is a function of past spread levels, past levels of inter-dealer competition, past levels of return volatility, past levels of trade volume and past levels of trade intensity. These two models are estimated using data obtained for 50 stocks listed on the LSE.

The quality of each model is assessed by considering the statistical and economic significance of the spread forecasts. Statistical significance is evaluated by consideration of the accuracy of the resultant forecasts with respect

[1] The mid-point quote is defined as the average of the bid and ask quotes.

to the actual outcomes. Economic significance is assessed by considering the profitability of a trading schedule based on the spread forecasts. In particular, the period corresponding to the minimum expected spread is scheduled into the trading activity of the investor. Furthermore, we assume that the investor making use of the proposed trading schedule is a passive portfolio manager interested in reducing the costs of buying and selling moderate amounts of stock over the course of a trading day. The performance of this manager is compared with the performance of an investment manager who trades at fixed times during the trading day. The results indicate that the former manager has a significant advantage over the latter manager in terms of the spreads faced and the prices paid and received for the stocks traded. Moreover, investors using the unrestricted VAR model enjoy superior profits to those enjoyed by investors using alternative spread models. These results confirm the superior accuracy of the unrestricted VAR model forecasts over the two-equation structural model forecasts.

This paper is organised as follows: Section 2 provides a description of recent developments in the organisational structure of the LSE. Section 3 describes two econometric models of spreads, Section 4 provides an account of the empirical findings and Section 5 concludes.

2. The structure of the London Stock Exchange

A fundamental change to the London equity market occurred when the LSE introduced the Stock Exchange Electronic Trading Service (SETS) on 20 October, 1997. The most important aspect of SETS was the change from quote-driven to order-driven trading. Under the quote-driven system, market makers post firm quotes and quantities on the LSE's Stock Exchange Automatic Quotation (SEAQ) bulletin board and subsequent trades are conducted over the telephone. Also, of the prices quoted by the market makers, the best bid and ask prices appear in a section of the screen known as the yellow strip. Investors can trade at the yellow strip quotes with those market makers who are willing to do so. Competition for order flow is further increased by the provision that investors can negotiate better prices on larger deals. [2]

The market currently operates under a dual trading system. In this system, market makers can still post quotes but these are only indicative. In this respect, the system differs from other markets like NASDAQ. Competition for order flow is increased by allowing members to post orders on the electronic order book via a computer terminal. Under this order-driven system, members

[2] For excellent accounts of the institutional arrangements under the quote-driven system, see Menyah and Paudyal (1996) and Snell and Tonks (1999).

can post a variety of firm orders, including limit orders and 'at best' orders. The latter type of order allows members to trade at the best price available. Limit orders allow members to post orders with a specified price and volume. Only when a corresponding order is placed with a matching price will a trade occur. The placing of such orders is not confined to members trading on behalf of clients. Indeed, of the trades carried out under the new system around 98% involve counter-parties trading on their own account (i.e., market makers). Such statistics suggest that spreads under the new system are likely to be determined in the same way as spreads under the quote-driven system. This is important when considering theories of bid–ask spreads.

Early indications suggest that SETS has been successful in achieving its aims of greater choice, greater transparency, lower spreads and increased volumes. Recent editions of the LSE's newsletter *Market Analysis* claim that the number of orders entered has increased, spreads have fallen and now over 50% of all trades are conducted through the electronic order book. [3] These achievements have led to more stocks being made available through SETS. Originally, only FTSE100 stocks could be traded electronically. This coverage has gradually been extended to stocks outside the FTSE100 index but within the FTSE250 index.

3. Econometric models of spreads

As dealers are still largely responsible for the setting of spreads in the LSE, traditional market-maker based theories of spread are still valid. In this section, we consider two models of spreads whose validity depends on such an assumption.

3.1. The unrestricted VAR model

Most theories consider the impact of relevant variables upon the order-processing cost, inventory cost and adverse-selection cost components of spread. Of the variables considered, competition and volatility are the most important (Ho and Stoll, 1983; Biais, 1993; Glosten, 1994; Huang and Masulis, 1999). It has been shown, under a variety of market structures (for example, Ho and Stoll, 1983; Biais, 1993; Glosten, 1994; Huang and Masulis, 1999), that increased (decreased) dealer competition leads to a decrease (increase) in bid–ask spreads and that an increase (decrease) in return volatility leads to an increase (decrease) in bid–ask spreads.

[3] The statistics are based on the March and November 1999 editions of *Market Analysis*. This newsletter can be downloaded from the following website: www.sets.co.uk.

In the absence of informed traders, Ho and Stoll (1983) show that the spread offered by risk-averse dealers is positively related to transaction size, dealer risk-aversion, return volatility and is negatively related to the number of dealers actively trading. Biais (1993) also assumes that dealers are risk-averse but assumes that dealers know the distribution of competing dealers' inventory positions. As the number of dealers increases, dealers are forced to post more attractive prices and so leads to lower bid–ask spreads. Similar arguments are used by Glosten (1994) where an electronic order book is assumed to be in operation. He shows that if dealers can observe competitors' quotes and suffer adverse-selection costs when trading with informed traders, then bid–ask spreads are positively related to the number of dealers trading. Ho and Stoll (1983) and Biais (1993) also demonstrate the mechanism by which return volatility affects bid–ask spreads via the inventory cost component of spreads. Increases in return volatility lead to an increase in inventory risk and thus, dealers post less attractive prices in an attempt to avoid unexpected inventory accumulation.

Using the above arguments, Huang and Masulis (1999) use a trivariate VAR to model bid–ask spreads, competition and return volatility in foreign exchange markets. We apply this methodology to the UK equity market in order to generate spread forecasts. However, we augment their model by including additional variables. In particular, we make use of measures of liquidity such as trading volume and trade intensity. [4] It has been argued that liquidity affects all three of the cost components of spread. Moreover, the theoretical microstructure literature argues that liquidity has a mixed effect on bid–ask spreads. For instance, Stoll (1978a,b) argues that trading volume will have a negative impact on spreads via the inventory cost component whereas Glosten and Milgrom (1985) argue that trading volume will have a positive impact on spreads via the adverse-selection cost component. Despite the mixed nature of these predictions, we proceed by including these measures of liquidity in the model in the hope that the effects are not perfectly offsetting.

For reasons of parsimony, we restrict the information set to the variables described above, that is, spreads, dealer competition, return volatility, trading volume and trade intensity – the exact definitions of these variables are given below. Following Huang and Masulis, we assume that conditional expectations of these variables are linear functions of the information set. As we only consider the recent history of the variables and wish to construct the best linear projection of the variables given this restricted information set, the variables are best represented in the following unrestricted VAR framework:

[4] Huang and Masulis could not include these variables in their model because such data are unavailable in foreign exchange markets.

$$\begin{bmatrix} s_{i,t} \\ I^{\mathrm{D}}_{i,t} \\ \sigma_{i,t} \\ V_{i,t} \\ I^{\mathrm{T}}_{i,t} \end{bmatrix} = \begin{bmatrix} \phi_1(L) & \phi_2(L) & \phi_3(L) & \phi_4(L) & \phi_5(L) \\ \phi_6(L) & \phi_7(L) & \phi_8(L) & \phi_9(L) & \phi_{10}(L) \\ \phi_{11}(L) & \phi_{12}(L) & \phi_{13}(L) & \phi_{14}(L) & \phi_{15}(L) \\ \phi_{16}(L) & \phi_{17}(L) & \phi_{18}(L) & \phi_{19}(L) & \phi_{20}(L) \\ \phi_{21}(L) & \phi_{22}(L) & \phi_{23}(L) & \phi_{24}(L) & \phi_{25}(L) \end{bmatrix} \begin{bmatrix} s_{i,t-1} \\ I^{\mathrm{D}}_{i,t-1} \\ \sigma_{i,t-1} \\ V_{i,t-1} \\ I^{\mathrm{T}}_{i,t-1} \end{bmatrix} + \begin{bmatrix} v_{1,i,t} \\ v_{2,i,t} \\ v_{3,i,t} \\ v_{4,i,t} \\ v_{5,i,t} \end{bmatrix}, \tag{1}$$

where $s_{i,t}$ is the (demeaned) touch half-spread on the ith stock, $I^{\mathrm{D}}_{i,t}$ denotes a (demeaned) measure of dealer competition on the ith stock, $\sigma_{i,t}$ denotes a (demeaned) measure of return volatility of the ith stock, $V_{i,t}$ denotes a (demeaned) measure of trading volume in the ith stock, $I^{\mathrm{T}}_{i,t}$ denotes a (demeaned) measure of trade intensity in the ith stock, $\phi_1(L)$, $\phi_2(L)$, etc. are lag polynomials each of order p, and $v_{1,i,t}$, $v_{2,i,t}$, etc. are error terms.

This VAR can be stacked into a first-order system and can be written as a first-order VAR,

$$\boldsymbol{x}_{i,t} = \boldsymbol{\Upsilon}\boldsymbol{x}_{i,t-1} + \boldsymbol{v}_{i,t}, \tag{2}$$

where $\boldsymbol{x}_{i,t}$ is a $5p \times 1$ vector of variables, $\boldsymbol{\Upsilon}$ is a $5p \times 5p$ matrix of coefficients and $\boldsymbol{v}_{i,t} \sim \mathsf{IN}_{5p}[\boldsymbol{0}, \boldsymbol{\Sigma}]$, with expectation $\mathsf{E}[\boldsymbol{v}_{i,t}] = \boldsymbol{0}$ and variance matrix $\mathsf{V}[\boldsymbol{v}_{i,t}] = \boldsymbol{\Sigma}$. Assuming that the forecasting model coincides with the DGP and the parameters are known then the h-step predictor is given by the conditional expectation,

$$\mathsf{E}(\boldsymbol{x}_{i,T+h} \,|\, \Omega_T) = \boldsymbol{\Upsilon}^h \boldsymbol{x}_{i,T}, \tag{3}$$

where Ω_T denotes the information set at time T.

3.2. The two-equation structural model

The motivation lying behind Huang and Stoll's (1994) model is the desire to predict the short-run behaviour of stock returns. In particular, they use various microstructure theories to derive reduce-form equations for mid-point quote changes and transaction price changes. The model can also be used to predict future spreads as the absolute value of the difference between the logarithm of the mid-point quote and the logarithm of the transaction price is the effective half-spread.

Huang and Stoll begin their analysis with a specification of the relationship between the logarithm of the mid-point quote on stock i, $q_{i,t}$, and the logarithm of the transaction price of stock i, $p_{i,t}$,

$$p_{i,t} = q_{i,t} + z_{i,t}, \tag{4}$$

where $z_{i,t}$ is the (signed) effective half-spread expressed as a proportion of the mid-point quote. Public purchases result in $z_{i,t} > 0$ and public sales result in $z_{i,t} < 0$. Taking first differences of (4) gives

$$\Delta p_{i,t} = \Delta q_{i,t} + z_{i,t} - z_{i,t-1}, \tag{5}$$

where $\Delta p_{i,t} = p_{i,t} - p_{i,t-1}$ and $\Delta q_{i,t} = q_{i,t} - q_{i,t-1}$. To give these equations an empirical content, Huang and Stoll develop a testable version by specifying the generating process for $\Delta q_{i,t}$ and $z_{i,t}$. These processes are described below.

The mid-point quote return can be decomposed as follows:

$$\Delta q_{i,t} = \mathsf{E}(\Delta q^*_{i,t} \mid \Omega_{t-1}) + f(\Delta I_{t-1}) + \epsilon_{i,t}, \tag{6}$$

where $\mathsf{E}(\Delta q^*_{i,t} \mid \Omega_{t-1})$ is the expectation of the consensus mid-point quote return conditional on the public information set Ω_{t-1} and $f(\cdot)$ is a function that captures the effect of inventory changes, ΔI_{t-1}, on the quote return. The conditional expectation is assumed to be a function of $z_{i,t-1}$ and the lagged return to an index futures contract Δf_{t-1}. The former variable measures the degree to which market makers adjust quotes on the basis of private information revealed through trading (Glosten and Milgrom, 1985). The latter variable measures the degree of quote adjustment in response to public information (Miller, 1990). Huang and Stoll assume that inventory effects, as captured by $f(\Delta I_{t-1})$, are functions of various measures of liquidity (defined below). Finally, linearity is imposed in Eq. (6) such that $\Delta q_{i,t}$ is a linear function of past information and inventory variables,

$$\begin{aligned}\Delta q_{i,t} = \beta_0 + \beta_1 \Delta q_{i,t-1} + \beta_2 \Delta f_{t-1} + \beta_3 z_{i,t-1} + \beta_4 Q_{i,t-1} + \beta_5 L^{\mathrm{A}}_{i,t-1} \\ + \beta_6 L^{\mathrm{B}}_{i,t-1} + \beta_7 D_{i,t-1} + \epsilon_{i,t},\end{aligned} \tag{7}$$

where $Q_{i,t-1}$ is the signed cumulative volume of trading in stock i that occurs between $t-2$ and $t-1$, $D_{i,t-1}$ is the difference between the logarithm of the quoted volume at the ask (depth at the ask) and the logarithm of the quoted volume at the bid (depth at the bid), and the large trade indicator variables, $L^{\mathrm{A}}_{i,t-1}$ and $L^{\mathrm{B}}_{i,t-1}$, are defined as follows:

$$\begin{aligned} L^{\mathrm{A}}_{i,t-1} &= \begin{cases} 1 & \text{if } z_{i,t-1} > 0 \text{ and } V_{i,t-1} > 10{,}000, \\ 0 & \text{otherwise}, \end{cases} \\ L^{\mathrm{B}}_{i,t-1} &= \begin{cases} 1 & \text{if } z_{i,t-1} < 0 \text{ and } V_{i,t-1} > 10{,}000, \\ 0 & \text{otherwise}. \end{cases} \end{aligned} \tag{8}$$

In addition to these liquidity-type variables, Huang and Stoll also include the lagged quote revision, $\Delta q_{i,t-1}$, to allow for slow adjustment in quote revisions.

Having specified the process followed by $\Delta q_{i,t}$, Eq. (5) shows that an attempt at price prediction is possible providing that a process is specified for $z_{i,t}$. For this purpose Huang and Stoll assume that $z_{i,t}$ follows a simple autoregressive process,

$$z_{i,t} = \rho z_{i,t-1} + \eta_{i,t}, \tag{9}$$

where ρ is a parameter and $\eta_{i,t}$ is the order arrival shock. This assumption is compatible with the Easley and O'Hara (1987) asymmetric information model where trade bunching is observed ($\rho > 0$). Alternatively, trades that restore the inventory equilibrium result in $\rho < 0$. Substituting (7) and (9) into (5) gives

$$\Delta p_{i,t} = \beta_0 + \beta_1 \Delta q_{i,t-1} + \beta_2 \Delta f_{t-1} + \beta_3^* z_{i,t-1} + \beta_4 Q_{i,t-1} + \beta_5 L^{\mathrm{A}}_{i,t-1} + \beta_6 L^{\mathrm{B}}_{i,t-1} + \beta_7 D_{i,t-1} + v_{i,t}, \tag{10}$$

where $\beta_3^* = \beta_3 + \rho - 1$ and $v_{i,t} = \epsilon_{i,t} + \eta_{i,t}$. The jointly estimated versions of Eqs. (7) and (10) represent the equations used by Huang and Stoll for predicting short-run stock returns. The mid-point quote return equation (7) can be interpreted as the inventory-adjusted equilibrium return while the transaction return in (10) essentially adjusts because of induced order arrival and bid–ask bounce around the quoted return. Augmented versions of these equations are used in the current application to predict future effective half-spreads. [5] In particular, the mid-point quote return and the transaction return are predicted h periods into the future. These predicted returns are then integrated to give the predicted mid-point quote and transaction price. [6] The absolute value of the difference between these predicted series is then used as a measure of the predicted effective half-spread.

4. Empirical results

This section contains a description of the data used, an account of various summary statistics pertaining to these data, and develops and tests a trading schedule based on time-consistent forecasts of bid–ask spreads.

4.1. Data

Transaction data covering 50 stocks over the period 3 August 1998–30 July 1999 were obtained from the LSE Data Service. The stocks and the sectors to which they belong are listed in Table 5 in Appendix A. These stocks are FTSE100 companies selected from each of the major sectors and represent companies of varying liquidity. We use a frequency of 5 minutes which is sufficiently low to avoid use of stale quotes and high enough to capture short-run movements in spreads.

[5] The equations are augmented by allowing the explanatory variables to be lagged by more than one period.

[6] The initial values of the series that are integrated are the actual mid-point quote and the transaction price. This ensures that the integrated series is correctly centered.

The trading day starts at 9.00 a.m. and ends at 4.30 p.m. However, the way that the variables are constructed means that the variables observed at 9.00 a.m. are the same as the variables observed at 4.30 p.m. on the previous day. To avoid this duplication, we only consider variables observed between 9.05 a.m. and 4.30 p.m. during each trading day.

One measure of spread used in the analysis is based on the best available limit prices from the electronic order book. Therefore, should an investor wish to trade immediately, he or she can do so at the price posted. Using these prices the touch half-spread is constructed as follows:

$$s_{i,t} = \frac{(A_{i,t} - B_{i,t})/2}{M_{i,t}}, \tag{11}$$

where $s_{i,t}$ denotes the touch half-spread for the ith stock at time t, $A_{i,t}$ denotes the ask price, $B_{i,t}$ denotes the bid price and $M_{i,t}$ denotes the price mid-way between the ask and bid prices. This measure of spread is used in the unrestricted VAR model. By contrast, the two-equation structural model uses the effective half-spread as defined in (4).

Space limitations prevent us from presenting results pertaining to all 50 stocks used in the analysis. Rather, stocks are grouped according to their level of liquidity. This enables an examination of the affect of liquidity on the various forecasting performance metrics. We assume that the mean level of the touch half-spread measured over the sample period gives a reasonable measure of liquidity. Stocks are then assigned to one of the 10 liquidity deciles ranging from highly liquid (low touch half-spread), denoted $L1$, to highly illiquid (high touch half-spread), denoted $L10$. Table 5 in Appendix A shows the allocation of the stocks to each of the liquidity deciles.

4.2. Preliminary results

The means of the variables used in this paper are presented in Table 1. The mean touch half-spread ranges from 18.84 basis points ($L1$) to 57.64 basis points ($L10$). The mean of these spreads across all stocks is 33.90 basis points. This compares to a mean spread on all FTSE100 companies during 1998 of 26.9 basis points (Naik and Yadav, 1999). One measure of dealer competition is the number of different dealers placing orders on SETS between $t-1$ and t. The results indicate that, on average, approximately three dealers are placing orders for each stock during each 5-minute interval. Moreover, the three measures of dealer competition ($I^{\mathrm{D}}_{i,t}$, $I^{\mathrm{T}}_{i,t}$, and $I^{\mathrm{O}}_{i,t}$) increase with the liquidity of the stock. The results in Table 1 also show that large block trades are more likely to occur at the bid ($\overline{L}^{\mathrm{B}}_{..} > \overline{L}^{\mathrm{A}}_{..}$) and depth at the ask in greater than depth at the bid ($\overline{D}^{\mathrm{A}}_{..} > \overline{D}^{\mathrm{B}}_{..}$). Finally, there appears to be heterogeneity in the volume and intensity of trading across the liquidity deciles. All these variables were also subjected to various moment tests. The results indicate that all variables

Table 1
Summary statistics[a]

Variable	$L1$	$L2$	$L3$	$L4$	$L5$	$L6$	$L7$	$L8$	$L9$	$L10$	All
$s_{i,t}$	18.84	24.99	27.26	28.52	30.24	32.78	36.12	39.83	42.83	57.64	33.90
$p_{i,t}$	6.99	6.80	6.45	6.39	6.57	5.99	6.02	6.20	6.03	6.43	6.39
$V_{i,t}$	9.03	7.21	10.90	8.97	7.73	17.81	16.33	10.71	13.32	7.67	10.97
$I_{i,t}^{\mathrm{D}}$	5.47	3.23	2.79	3.12	2.75	2.63	2.65	2.42	2.50	2.16	2.97
$I_{i,t}^{\mathrm{T}}$	12.40	6.55	4.16	5.43	5.85	3.73	5.15	3.38	3.44	2.06	5.21
$I_{i,t}^{\mathrm{O}}$	8.13	4.27	3.53	4.14	3.34	3.17	3.24	2.84	2.93	2.36	3.79
$\sigma_{i,t}$	50.71	48.97	34.90	43.49	43.67	47.23	43.29	45.71	47.60	54.86	46.04
$L_{i,t}^{\mathrm{A}}$	0.04	0.04	0.06	0.04	0.04	0.07	0.07	0.06	0.06	0.04	0.05
$L_{i,t}^{\mathrm{B}}$	0.09	0.08	0.10	0.07	0.06	0.11	0.11	0.10	0.10	0.07	0.09
$D_{i,t}^{\mathrm{A}}$	13.88	12.11	15.25	13.54	11.29	20.93	15.81	12.35	19.59	8.31	14.31
$D_{i,t}^{\mathrm{B}}$	13.38	11.91	14.96	12.93	11.07	19.47	16.00	12.47	19.23	8.01	13.94

[a] This table gives the mean values of 5-minute frequency touch half-spreads ($s_{i,t}$), trading volumes ($V_{i,t}$), dealer intensities ($I_{i,t}^{\mathrm{D}}$), trade intensities ($I_{i,t}^{\mathrm{T}}$), order intensities ($I_{i,t}^{\mathrm{O}}$), return volatilities ($\sigma_{i,t}$), proportions of block trades at the ask ($L_{i,t}^{\mathrm{A}}$), proportions of block trades at the bid ($L_{i,t}^{\mathrm{B}}$), depths at the ask ($D_{i,t}^{\mathrm{A}}$) and depths at the bid ($D_{i,t}^{\mathrm{B}}$). Spreads are measured in basis points and prices are given in pounds sterling. Trade volumes and order depths are measured in terms of thousands of shares. The liquidity deciles are denoted $L1$ for the most liquid stocks to $L10$ for the least liquid stocks.

are skewed, highly leptokurtic and have non-normal distributions. [7] As such, statistical inference conducted in the subsequent sections cannot be based on standard statistical techniques.

The intraday periodicity of selected variables can be seen in panels A, B, C, and D of Fig. 1. This shows the intraday mean values of touch half-spreads, trading volumes, dealer competition and return volatilities averaged over all stocks. All variables possess a strong periodic component. Spreads are on average highest during the first hour of trading. [8] Return volatilities appear to be high during the opening and closing of the trading day while trading volumes are extremely high during the closing period. Dealer intensities have an interesting intraday pattern. They appear to increase rapidly around 2.30 p.m. when US equity markets open. All the variables were also subjected to a more formal test of periodicity based on a simple *F*-test. [9] These results show that the majority of variables, including those presented in Fig. 1, are characterised by significant periodicity. This finding suggests that failure to incorporate a periodic component in both models could result in less accurate forecasts.

4.3. Model estimation

Following Huang and Stoll (1994) and Huang and Masulis (1999), the VAR given in (1) and Eqs. (7) and (10) are estimated using Hansen (1982) generalised method of moments (GMM) estimation methodology. As there exists a significant periodic component in the variables, a lag length equal to one trading day is imposed on the model. The remaining lag order is selected using the

---

[7] These results are available upon request.

[8] Naik and Yadav (1999) and Shah (1999) find a similar result using all stocks in the FTSE100.

[9] To formally test for intraday periodicity we estimate the following regression by ordinary least squares:

$$x_{i,t} = \sum_{j=1}^{15} \omega_j D_{j,t} + \epsilon_{i,t},$$

where $x_{i,t}$ denotes the (demeaned) variable of interest for the *i*th stock, $D_{j,t}$ is a dummy variable that equals unity during the *j*th non-overlapping half-hour period during the day and zero otherwise, ω_j is a coefficient to be estimated and $\epsilon_{i,t}$ is an error term. The length of the trading day means that we consider fifteen different intraday half-hour intervals. As the variables are non-normally distributed statistical inference is carried out using the bootstrap technique. This technique involves shuffling the dataset (with replacement) and re-estimating the coefficients, *t*-statistics and *F*-statistics. This process is repeated 999 times with the statistics being recorded after each replication. Once a set of statistics has been compiled, the original statistics are compared with the bootstrap statistics and *P*-values are calculated. When a (bootstrap-based) *F*-test is carried out to test for the significance of all the dummy variables, the results indicate that the majority of variables have a significant deterministic periodic component. Details of these results are available upon request.

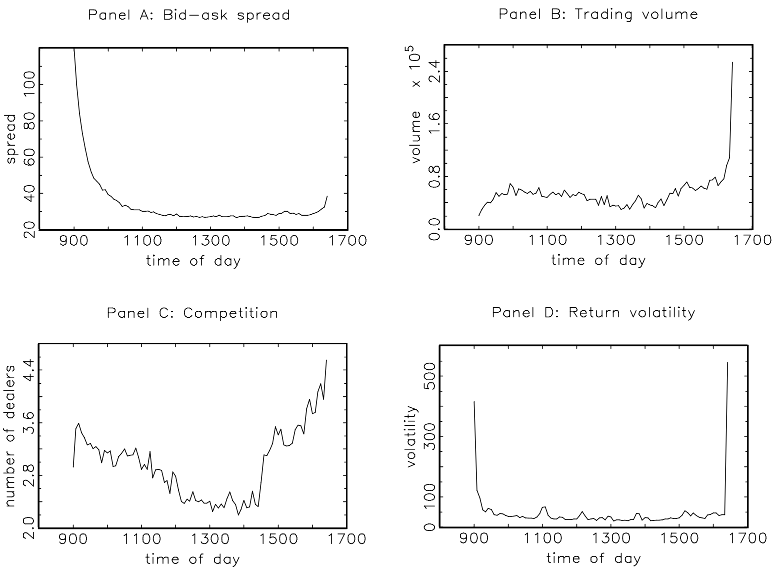

Fig. 1. Intraday periodicity.

Akaike information criterion (AIC) and the Schwarz information criterion (SIC). These criteria select the optimal lag length from lag lengths ranging from 1–12. However, for reasons of space we only report the results pertaining to the AIC-determined lag lengths. [10]

To generate numerous h-step ahead out-of-sample forecasts the amount of in-sample data used is varied in two different ways. First, beginning with one week of in-sample data the estimation process is recursively pushed through the sample while maintaining a fixed start point of the data. Second, again beginning with one week of in-sample data the estimation is rolled through the sample while maintaining exactly one week of in-sample data. [11]

Upon completion of each estimation of the models, 1-step ahead (5 minutes ahead), 2-step ahead (10 minutes ahead), 3-step ahead (15 minutes ahead), 4-step ahead (20 minutes ahead) and 5-step ahead (25 minutes ahead) forecasts are generated. In the case of the two-equation structural model, the quote and

[10] The SIC-based results are similar in nature and can be obtained upon request.

[11] For example, the whole sample period starts at 9.00 a.m. on 3 August and ends at 4.30 p.m. on 30 October. Therefore, the first in-sample period used (in both cases) starts at 9.05 a.m. on 3 August and ends at 4.30 p.m. on 7 August. In both cases an additional 30 min of data is used during each estimation. Therefore, the second recursive in-sample period will start at 9.05 a.m. on 3 August and will end at 9.30 a.m. on 10 August. In contrast, the second rolling in-sample period will begin at 9.35 a.m. on 3 August and will end at 9.30 a.m. on 10 August.

transaction returns are integrated to give the predicted levels. The absolute values of the difference of these series is used as a measure of the expected effective half-spread.

4.4. Assessing forecast quality

The quality of the VAR-based forecasts is compared with the quality of forecasts generated by a simple random walk (M1) model. In particular, we are assuming that cumulative spreads follow a random walk while the spread itself is a white noise process with positive mean. Therefore, M1 generates forecasts equal to the mean of the in-sample period spread. The quality of the M1, the unrestricted VAR (M2) and the two-equation structural model (M3) forecasts is assessed using the mean squared forecast error (MSFE) and the mean absolute forecast error (MAFE). These are defined as

$$\mathrm{MSFE} = \frac{1}{T}\sum_{t=1}^{T}(y_{t+h} - \hat{y}_{t+h})^2, \tag{12}$$

$$\mathrm{MAFE} = \frac{1}{T}\sum_{t=1}^{T}|y_{t+h} - \hat{y}_{t+h}|, \tag{13}$$

where y_{t+h} is the realisation of the series at time $t+h$, $\hat{y}_{t+h}$ is the h-step ahead forecast of the series using data observed up to and including time t and T is the number of h-step ahead forecasts considered. In the case of M2, the predicted series is the touch half-spread, while the M3 predicted series are quote returns, transaction returns and effective half-spreads.

To formally test the comparative accuracy of the M1, M2 and M3-based forecasts, we make use of the asymptotic test introduced by Diebold and Mariano (1995). This test allows use of an arbitrary loss function instead of the usual squared forecast error loss, non-zero mean forecast errors, non-normally distributed forecast errors and serially correlated forecast errors. In the current application, it is the robustness of the Diebold and Mariano statistic to the non-normality assumption that is most attractive. Indeed, when the forecast errors are tested for normality the null is rejected at the 1% significance level on every occasion. [12]

Diebold and Mariano show that the following test statistic is (asymptotically) normally distributed with zero mean and unit variance:

$$S_1 = \frac{\bar{d}}{\sqrt{(2\pi\hat{f}_{\mathrm{d}}(0))/T}}, \tag{14}$$

[12] These results are available upon request.

where

$$\bar{d} = \frac{1}{T}\sum_{t=1}^{T}[g(e_t) - g(e'_t)] \tag{15}$$

is the sample loss differential and $\hat{f}_{\mathrm{d}}(0)$ is a consistent estimate of the spectral density of the loss differential at frequency zero,

$$f_{\mathrm{d}}(0) = \frac{1}{2\pi}\sum_{\tau=-\infty}^{\infty}\gamma_{\mathrm{d}}(\tau), \tag{16}$$

where e_t is the M1 forecast error, e'_t is the M2 or M3 forecast error, $g(\cdot)$ is the loss function, $\gamma_{\mathrm{d}}(\tau) = E[(d_t - \mu)(d_{t-\tau} - \mu)]$ is the autocovariance of the loss differential at displacement τ and μ is the population mean loss differential. Following Diebold and Mariano, we use a uniform lag window of size $h - 1$ to estimate $f_{\mathrm{d}}(0)$. The loss functions used are the squared function (MSFE) and the absolute function (MAFE).

A summary of the results obtained when the Diebold and Mariano test is performed on all available forecasts is given in Table 2. For reasons of space, we report the proportion of stocks for which the M2 or M3 forecasts are significantly better (at the 5% level) than the M1 forecasts. [13] The results indicate that the M2 forecasts are almost always significantly better than the M1 forecasts. By contrast, the M3 forecasts are less impressive. However, a sizable proportion of these forecasts are significantly better than the M1 forecasts. This superiority is most apparent when 1-step ahead forecasts are considered and the MAFE loss function is used. Thus, the predictive power of the two-equation structural model is concentrated in the immediate future.

The strong intraday periodic pattern found in spreads may affect the quality of the forecasts. On most days between 9.00 a.m. and 10.00 a.m., spreads are high and falling (see Fig. 1). In these periods, M1 is likely to do very well particularly when h is large. This is because the spread will approach the mean of the process and the forecasts will appear to be accurate. To avoid such spurious accuracy, we remove all data observed before 10.00 a.m. and the Diebold and Mariano test is repeated. The results indicate that the M2 forecasts, and to a lesser extent the M3 forecasts, are still significantly more accurate than the M1 forecasts in the majority of cases. Therefore, the M2 and M3 forecasts are more accurate than the M1 forecasts regardless of the time of day in which forecasts are generated.

[13] Details of the test results for each stock can be obtained upon request.

Table 2
Comparing forecast quality[a]

Comparison	Variable	Metric	h 1	2	3	4	5
Panel A: Full sample							
M1 v. M2	$s_{i,t}$	MSFE	0.94	0.98	0.98	0.92	0.74
		MAFE	0.98	0.98	0.96	0.88	0.72
M1 v. M3	$\Delta p_{i,t}$	MSFE	0.52	0.02	0.02	0.00	0.00
		MAFE	0.18	0.00	0.00	0.00	0.00
	$\Delta q_{i,t}$	MSFE	0.64	0.10	0.06	0.06	0.02
		MAFE	0.90	0.10	0.06	0.06	0.04
	$\lvert z_{i,t}\rvert$	MSFE	0.06	0.00	0.02	0.02	0.00
		MAFE	0.42	0.36	0.20	0.14	0.16
Panel B: After 10 a.m.							
M1 v. M2	$s_{i,t}$	MSFE	0.94	1.00	0.96	0.84	0.72
		MAFE	0.98	1.00	0.98	0.88	0.74
M1 v. M3	$\Delta p_{i,t}$	MSFE	0.70	0.00	0.02	0.02	0.06
		MAFE	0.18	0.02	0.00	0.00	0.02
	$\Delta q_{i,t}$	MSFE	0.80	0.08	0.10	0.04	0.04
		MAFE	0.92	0.12	0.10	0.06	0.08
	$\lvert z_{i,t}\rvert$	MSFE	0.30	0.12	0.14	0.14	0.00
		MAFE	0.60	0.50	0.38	0.26	0.26

[a] This table gives the proportion of stocks for which the forecasts generated by M2 or M3 are significantly better than the forecasts generated by M1. The null hypothesis that M1 forecasts are of the same quality as M2 or M3 forecasts is tested using the Diebold and Mariano (1995) asymptotic test using a 5% significance level. The alternative hypothesis adopted is that M2 or M3 forecasts are better than M1 forecasts. $s_{i,t}$ is the (demeaned) touch half-spread on the ith stock at time t, $p_{i,t}$ is the logarithm of the transaction price, $q_{i,t}$ is the logarithm of the midpoint quote and $z_{i,t}$ is the (signed) half-spread expressed as a proportion of the mid-point quote. This test is implemented using h-step ahead forecasts with the mean squared forecast error (MSFE) or the mean absolute forecast error (MAFE).

4.5. A simple trading schedule

Having found that it is possible to generate accurate forecasts of touch half-spreads (via M2) and effective half-spreads (via M3), we now consider whether a trading rule can be devised that yields lower spreads to investors. Throughout this section, we assume that the investor is a passive portfolio manager whose objective is to purchase shares at anytime during the trading day. This manager does not wish to purchase mis-priced shares. Rather, the motivation for buying and selling shares is portfolio re-balancing.

The model-based methodologies generate forecasts of spreads every k minutes. Using these forecasts, we assume that an investment manager selects a time point within the next k minutes in which trading will take place. If the current spread is lower than all h-step ahead forecasts, then trading will take place immediately. Otherwise, trading will take place when the spread is expected to be smallest. This trading schedule is compared with a naive trading schedule where trades take place immediately during the k-minute trade horizon. [14] This naive trading schedule is compatible with the assumptions of M1.

We consider k-minute trade horizons of 15, 30 and 60 minutes. However, for reasons of space only the 30-minute horizon results are presented. [15] The results of using such trading schedules are presented in Table 3. Without exception the mean spreads faced by an investment manager using the M2-based schedule are lower than those faced by an investment manager using a naive schedule. That is, the mean percentage spread discount ranges from 24.85% ($L10$) to 45.31% ($L1$). This particular result indicates that higher spread discounts are achieved when highly liquid stocks are traded. Averaging over all stocks, a mean percentage spread discount of 34.39% is obtained when an investor uses the M2-based schedule. The spread discounts obtained by an investment manager using the M3-based schedule are less impressive. Most importantly, the spread discounts vary considerably over the liquidity deciles, ranging from −31.32% ($L5$) to 19.65% ($L9$). Also, the average spread discount across all stocks is lower than the average spread discount enjoyed by an M2-based investment manager (3.61% vs. 34.39%).

These discounts are not evenly distributed over the trading day. Fig. 2 gives the mean percentage spread discounts averaged over all stocks during each trade horizon within the trading day. As we are using a 30-minute trade horizon, then there are 15 such trade horizons during the trading day. The results indicate that the highest discounts occur at the beginning of each trading day. This is because spreads are high and falling at this time of day. As such, because the M2 and M3 forecasts approach the mean of the process when h is large, there is a natural tendency to schedule trades when h is large. At this point in time the actual spread is smaller than at other times in the trade horizon. By contrast, a (very) naive investment manager trading immediately during the trade horizon invariably incurs higher spreads. Outside of this time period, consistent touch half-spread discounts of around 20% (via M2) and effective half-spread discounts of around 5% (via M3) can be achieved.

4.6. Intraday spread differences

Having found that mean spread discounts above zero occur, it is natural to consider whether the spreads incurred by a model-based investment manager are significantly different from the spreads incurred by a naive investment

[14] The market impact costs associated with these schedules are likely to be the same. As such, we do not attempt to measure these costs. For a practitioner's account of market impact costs, see Mizzi (1999).

[15] Results obtained using the other horizons are similar in nature and can be obtained upon request.

Table 3
Spread, ask and bid price comparisons[a]

Discount	*L*1	*L*2	*L*3	*L*4	*L*5	*L*6	*L*7	*L*8	*L*9	*L*10	All
Panel A: *Spread discounts*											
$\bar{\theta}_i^{\mathrm{s}}(M2)$	45.31	42.23	36.58	38.12	40.04	31.99	36.26	33.93	31.81	24.85	34.39
$\bar{\theta}_i^{\mathrm{s}}(M3)$	−2.57	17.95	6.94	9.98	−31.32	7.89	4.51	−1.77	19.65	12.91	3.61
Panel B: *Price discounts*											
$\bar{\theta}_i^{\mathrm{B}}(M2)$	0.07	0.08	0.07	0.08	0.08	0.08	0.10	0.10	0.10	0.12	0.09
$\bar{\theta}_i^{\mathrm{A}}(M2)$	0.07	0.09	0.08	0.10	0.12	0.08	0.12	0.13	0.13	0.13	0.10
$\bar{\theta}_i^{\mathrm{B}}(M3)$	0.04	0.03	0.02	0.03	0.01	0.02	0.01	0.03	0.02	0.01	0.02
$\bar{\theta}_i^{\mathrm{A}}(M3)$	0.04	0.02	0.01	0.02	0.04	0.02	0.03	0.02	0.03	−0.03	0.03

[a] This table gives the mean spread discounts (Panel A) obtained using the M2 and M3-based schedules versus the naive (M1-based) schedule. These discounts are denoted $\bar{\theta}_i^{\mathrm{s}}(M2)$ and $\bar{\theta}_i^{\mathrm{s}}(M3)$, respectively. Also, the mean ask and bid price discounts (Panel B) obtained using the various schedules are denoted $\bar{\theta}_i^{\mathrm{A}}(M2)$, $\bar{\theta}_i^{\mathrm{A}}(M3)$, $\bar{\theta}_i^{\mathrm{B}}(M2)$ and $\bar{\theta}_i^{\mathrm{B}}(M3)$, respectively. Mean discounts are given in percentage terms. The liquidity deciles are denoted *L*1 for the most liquid stocks to *L*10 for the least liquid stocks.

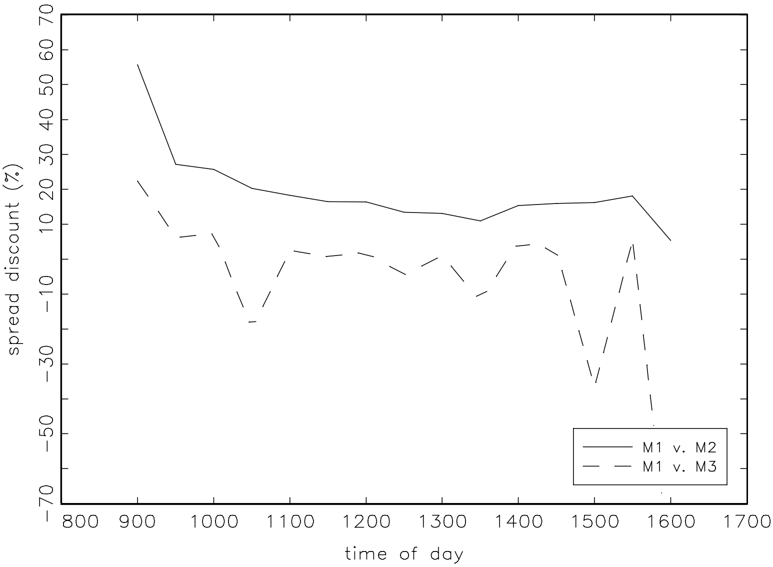

Fig. 2. Intraday spread discounts.

manager. More formally, we test the null hypothesis that the M2 and M3-based investment managers' mean spread equals a naive investment managers' mean spread. A simple t-test cannot be used to test this hypothesis because of the non-normality of the underlying data. Indeed, normality tests applied to the spreads incurred by both types of investment manager universally reject the null. Therefore, we make use of the studentized bootstrap method. [16] In this case, the usual two-sample t-statistic

$$Z = \frac{\bar{s}_{i,t} - \bar{s}'_{i,t} - (\mu_{\mathrm{s}} - \mu_{\mathrm{s}'})}{\sqrt{\sigma^2_{\mathrm{s}}/T + \sigma^2_{\mathrm{s}'}/T}} \tag{17}$$

is approximately pivotal, that is, it has a distribution that is approximately independent of unknown parameters. Using the selected spreads under the M2 and M3-based and naive trading schedules, the following test statistic is calculated:

$$z_0 = \frac{\bar{s}_{i,t} - \bar{s}'_{i,t}}{\sqrt{\hat{\sigma}^2_{\mathrm{s}}/T + \hat{\sigma}^2_{\mathrm{s}'}/T}}, \tag{18}$$

where $\bar{s}_{i,t}$ is the mean spread incurred by the M2 or M3-based investment managers, $\bar{s}'_{i,t}$ is the mean spread incurred by the M1-based investment man-

[16] For more details on this technique, see Davison and Hinkley (1997).

ager, $\hat{\sigma}_s^2$ and $\hat{\sigma}_{s'}^2$ denote the sample variances of the spreads incurred by investment managers using M2 or M3-based and naive trading schedules, respectively. The distribution of this test statistic is calculated using bootstrap re-sampling. In particular, R values of

$$z^* = \frac{\overline{s}^*_{i,t} - \overline{s}'^*_{i,t} - (\overline{s}_{i,t} - \overline{s}'_{i,t})}{\sqrt{\hat{\sigma}_s^{*2}/T + \hat{\sigma}_{s'}^{*2}/T}} \tag{19}$$

are generated, where statistics denoted by * indicate that they are based on the shuffled (using replacement) dataset containing T observations. Finally, the z_0 statistic is compared with the R separate z^* statistics and P-values are calculated. In this particular application, we set $R = 999$.

The results of carrying out the above test are given in Table 4. The entries in this table represent the mean differences in spreads between the M2-based and naive (M1-based) schedules and are subdivided into fifteen (non-overlapping) intervals over the trading day. In addition, the mean differences over the entire sample period are given. The entries in this table also give an indication of the null hypothesis rejection decision. For all liquidity deciles, the results indicate that mean spreads are significantly different from each other. The mean spread differences are greatest between 9.00 a.m. and 9.30 a.m. and are generally, significantly different throughout the trading day. When the entire sample is used mean spreads are significantly different.

It is possible that the above result is dependent on the large differences observed during the first hour of trading. To examine this conjecture we remove these data from the sample and repeat the analysis. The results indicate that it is still possible to reject the null of no spread differences at significance levels below 1%. The test is repeated by pooling the spreads for all stocks. This test is carried out for the differences between the M2-based and naive (M1-based) schedules and between the M3-based and naive (M1-based) schedules. The results indicate that the M2-based schedule always delivers preferable spread levels while the M3-based schedule is only marginally preferable. Indeed, when considering the differences over the entire trading day, the null hypothesis cannot be rejected even at the 10% level when the M3-based schedule is used. The overall conclusion is, therefore, that the spreads incurred by an investment manager who makes use of a model-based trading schedule are significantly less than those spreads incurred by an investment manager who trades at fixed times throughout the day.

4.7. Spread differences and market depth

It could be argued that the above touch half-spread differences can only be enjoyed by an investment manager wishing to buy or sell certain amounts of a

Table 4
Intraday spread differences[a]

Interval	*L*1	*L*2	*L*3	*L*4	*L*5	*L*6	*L*7	*L*8	*L*9	*L*10	All
Panel A: *M1 vs. M2*											
1	29.82***	15.44***	11.80***	27.65***	34.88***	22.23***	34.62***	17.21***	19.27***	32.58***	24.55***
2	23.33***	6.81***	16.71***	17.08***	21.50***	28.56***	22.10***	31.27***	10.63***	9.77***	18.78***
3	4.16***	4.80***	5.25***	15.83***	6.68***	12.88***	16.10***	20.85***	28.42***	14.67***	12.96***
4	4.54***	14.99***	9.51***	5.37***	8.43***	6.49***	20.05***	7.28***	5.16***	13.01***	9.48***
5	3.80***	4.34***	4.11***	7.50***	4.83***	4.50***	11.58***	7.23***	6.48***	12.14***	6.65***
6	4.10***	3.60***	3.77***	3.84***	4.95***	4.99***	7.39***	6.40***	2.58*	12.24***	5.39***
7	2.27***	3.83***	3.12***	6.37***	3.27***	4.90***	8.15***	6.41***	3.92***	6.68***	4.89***
8	3.94***	2.29***	3.34***	2.95***	5.39***	4.49***	6.73***	3.26***	6.44***	3.84***	4.27***
9	3.62***	5.01***	3.77***	2.66***	3.58***	4.70***	5.52***	4.51***	9.87***	6.57***	4.98***
10	2.77***	5.50***	2.17***	3.57***	3.04***	3.51***	3.82***	3.10***	4.59***	3.86***	3.59***
11	3.40***	15.64***	5.69***	6.54***	2.99***	4.51***	5.29***	5.43***	4.12**	8.99***	6.26***
12	3.26***	3.76***	5.20***	3.94***	4.70***	4.73***	3.45***	5.33***	7.72***	18.39***	6.04***
13	4.61***	14.64***	11.87***	15.60***	5.64***	5.96***	4.51***	13.09***	29.15***	9.83***	11.49***
14	4.34***	5.24***	15.77***	6.70***	19.91***	7.66***	5.82***	6.08***	12.25***	12.20***	9.60***
15	4.96***	20.43***	17.97***	6.62***	16.41***	8.66***	7.04***	30.82***	21.32***	21.96***	15.62***
*S*1	6.86***	8.42***	8.00***	8.82***	9.75***	8.59***	10.81***	11.22***	11.46***	12.45***	9.64***
*S*2	3.83***	8.01***	7.04***	6.73***	6.91***	6.00***	8.11***	9.21***	10.93***	11.11***	7.79***
Panel B: *M1 vs. M3*											
*S*1	−1.71	5.22	1.78***	2.58***	−13.60	2.70	1.52**	−0.87	7.94	9.69	1.53
*S*2	−2.33	5.47	1.56***	1.56***	−4.57	1.84	1.01	−2.02	8.21	10.83	2.16

[a] This table gives the mean differences between the naive (M1-based) schedule spreads and the M2 or M3-based schedule spreads for various stocks and during different intraday periods. The intervals considered range from 9.00 a.m. to 9.30 a.m. (interval 1) to 4.00 p.m. to 4.30 p.m. (interval 15). The *S*1 row gives the mean differences during the whole period and the *S*2 row gives the mean differences during the whole period but with observations obtained before 10 a.m. during each day being excluded from the analysis. Statistical inference is carried out using the studentized bootstrap technique. The liquidity deciles are denoted *L*1 for the most liquid stocks to *L*10 for the least liquid stocks.

[*] Denotes rejection of the null hypothesis that there is no difference in mean spreads at the 10% level.

[**] Denotes rejection of the null hypothesis that there is no difference in mean spreads at the 5% level.

[***] Denotes rejection of the null hypothesis that there is no difference in mean spreads at the 1% level.

particular stock. This is because these spreads are only good for trade sizes up to that specified in the best order. Therefore, it is quite possible that spreads may be larger for alternative trade sizes. This issue is examined by assigning the touch half-spreads obtained by the M2-based schedule to various depth deciles. [17] For example, if the M2-based trading schedule indicates that a trade should occur at 10.15 a.m. then the depth of the market at that time is recorded and the trade is assigned to that particular depth decile. The above test is then carried out on the spread differences within each depth decile. Space limitations prevent us from presenting these results. However, the results indicate that the spread differences remain highly significant irrespective of the depth of the market. Thus, the observed spread discounts can be enjoyed regardless of the amount of shares one wishes to trade.

4.8. Portfolio cost savings

Having found that lower spreads are incurred by a model-based investment managers, we now consider the bid and ask prices incurred by such a manager. The prices incurred by an M2-based investment manager correspond to the prices that prevail during the period in which the expected touch half-spread is minimised. By contrast, an M3-based investment manager will face prices which occur when the expected quote is minimised (ask prices) or maximised (bid prices). The mean percentage discounts on bid and ask prices are given in Table 4. Although the discounts on these prices are small they can reach levels of 0.13% when considering the ask prices of the $L8$, $L9$ and $L10$ liquidity deciles if M2 is used. Following the earlier results, an M2-based investment manager enjoys more heavily discounted prices than an M3-based investment manager. This undoubtedly reflects the superior quality of the M2 forecasts.

We now attempt to motivate use of a model-based trading schedule by considering an M2 and M3-based investment manager who wishes to build (unwind) a position in all 50 stocks. We assume that over the trading day, this investment manager buys (sells) £250 worth of each stock during each 30-minute trade horizon. The cumulative cash savings of this investment manager are calculated in comparison to an investment manager who uses a naive trading schedule. The results indicate that substantial savings are possible, particularly when building positions. For instance, an M2-based investment manager achieves savings of £49,564 and £48,007 when building and unwinding positions in all 50 stocks, respectively. This compares to respective

[17] These depth deciles represent the sum of the depths specified in all the best ask and bid orders placed on SETS for a particular stock.

savings of £11,799 and £9,018 when an M3-based investment manager trades the portfolio of stocks.

The cumulative cash savings obtained when an investment manager does not trade during the first hour are also calculated. [18] The latter piece of analysis measures the sensitivity of results to the huge spread discounts observed at the beginning of the trading day. Both the M2 and M3-based investment managers experience a fall in savings regardless of whether they are buying or selling stocks. For instance, an M2-based manager who wishes to build a position in the portfolio of stocks can now only achieve savings of £23,572. Similarly, an M3-based manager pursuing a similar strategy can only achieve savings of £4,541.

5. Conclusion

The economic and statistical significance of forecasts generated by two econometric models of bid–ask spreads are examined in this paper. We show that it is possible to construct forecasts of bid–ask spreads that can be beneficial to investment managers who wish to save money on the prices they pay and receive for individual stocks. Spread savings of up to 45% can be enjoyed by investment managers willing to make use of an unrestricted VAR model. Similar savings have been brought about by the introduction of entirely new trading systems. For example, Gemmill (1998) reports a 39 basis point spread for large companies and a 79 basis point spread for small companies before the introduction of SETS. By contrast, the respective spreads after the introduction of SETS were 32 basis points and 53 basis points. A comparison of these spread reductions with the ones observed in this paper indicates clear evidence of the economic significance of forecasts generated by econometric models of bid–ask spreads.

The paper also shows that spread discounts are not evenly distributed throughout the day. During the beginning of the trading day spreads are higher than at any other time of the day. However, when this effect is controlled for, the results indicate that substantial spread discounts are still available throughout the rest of the trading day and appear to be evenly distributed after 10 a.m. Indeed, a simple example shows that substantial cost savings can be achieved even if an investment manager cannot trade when the market opens.

[18] As an hour of trading is lost, the investment manager is assumed to trade additional stock throughout the day such that the daily trading value remains unchanged. This amounts to assuming that the investment manager now trades approximately £288 ($=250\times(15/13)$) worth of each share during each 30-minute trade horizon.

Table 5
The subsample of FTSE100 stocks[a]

Stock	Sector	Liquidity decile
3I Group	Investment trusts	$L10$
Abbey National	Banks	$L7$
Alliance and Leicester	Banks	$L6$
Allied Domecq	Alcoholic beverages	$L8$
Asda Group	Food retailers	$L7$
Associated British Foods	Food manufacturers	$L10$
BAA	Transport	$L1$
Bank of Scotland	Banks	$L9$
Barclays	Banks	$L1$
British Gas	Oil and gas	$L1$
BOC Group	Chemicals	$L6$
Boots	General retailers	$L3$
British Airways	Transport	$L4$
British Energy	Electricity	$L5$
British Land	Properties	$L8$
British Telecommunications	Telecommunications	$L1$
Cable and Wireless	Telecommunications	$L7$
Cadbury Schweppes	Food manufacturers	$L3$
Carlton Communications	Media	$L10$
Centrica	Oil and gas	$L8$
EMI Group	Media	$L7$
General Electric	Electricity	$L9$
Glaxo Welcome	Pharmaceuticals	$L4$
Great Universal Stores	General retailers	$L3$
Halifax	Banks	$L3$
Hays	Support services	$L2$
HSBC Holdings	Banks	$L10$
Imperial Chemical Industries	Chemicals	$L10$
Ladbroke Group	General retailers	$L6$
Land Securities	Properties	$L7$
Legal and General	Insurance	$L7$
Lloyds TSB Group	Banks	$L8$
Lucasvarity	Engineering	$L8$
Marks and Spencer	General retailers	$L4$
National Power	Electricity	$L3$
National Westminster Bank	Banks	$L1$
Norwich Union	Insurance	$L7$
Orange	Telecommunications	$L9$
Pearson	Media	$L5$
Peninsular and Orient Steam	Transport	$L6$
Powergen	Electricity	$L1$
Prudential Corporation	Insurance	$L4$
Railtrack Group	Transport	$L5$
Reckitt and Coleman	Household goods	$L9$
Reed International	Media	$L2$
Rentokil Initial	Support services	$L4$
Reuters Group	Media	$L2$
Rio Tinto	Mining	$L2$
Rolls-Royce	Engineering	$L6$

[a] This table gives the stocks examined and their respective sectors and liquidity deciles ($L1$–$L10$).

Acknowledgements

I thank two anonymous referees for helpful comments on an earlier draft of this paper and the Institute of Chartered Accountants in England and Wales for providing financial support.

Appendix A

See Table 5.

References

Biais, B., 1993. Price formation and equilibrium: Liquidity in fragmented and decentralised markets. Journal of Finance 48, 1765–1789.

Davison, A., Hinkley, D., 1997. Bootstrap Methods and their Application. Cambridge University Press, Cambridge.

Diebold, F., Mariano, R., 1995. Comparing predictive accuracy. Journal of Business and Economic Statistics 13, 253–265.

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69–90.

Gemmill, G., 1998. Touched by success. Financial News Briefing Notes.

Glosten, L., 1994. Is the electronic open limit order book inevitable? Journal of Finance 49, 1127–1161.

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

Hansen, L., 1982. Large sample properties of generalized method of moments estimators. Econometrica 50, 1029–1054.

Ho, T., Stoll, H., 1983. The dynamics of dealer markets under competition. Journal of Finance 38, 1053–1075.

Huang, R., Stoll, H., 1994. Market microstructure and stock return predictions. Review of Financial Studies 7, 179–215.

Huang, R., Masulis, R., 1999. FX spread and dealer competition across the 24-hour trading day. Review of Financial Studies 12, 61–93.

Menyah, K., Paudyal, K., 1996. The determinants and dynamics of bid–ask spreads on the London Stock Exchange. Journal of Financial Research 19, 377–394.

Miller, M., 1990. Index arbitrage and volatility. Financial Analysts Journal 46, 6–7.

Mizzi, J., 1999. Transaction costs: A practitioner's experience in the UK. Working Paper. Henderson Investors.

Naik, N., Yadav, P., 1999. The effects of market reform on trading costs of public investors: Evidence from the London Stock Exchange. Working Paper. London Business School.

Shah, N., 1999. The European equity trading day: When should we trade? Working Paper. Equity Derivatives Research, Goldman Sachs.

Snell, A., Tonks, I., 1999. Measuring microstructure effects in less-liquid stocks on the London Stock Exchange. Working Paper. University of Bristol.

Stoll, H., 1978a. The supply of dealer services in securities markets. Journal of Finance 33, 1133–1151.

Stoll, H., 1978b. The pricing of security dealer services: An empirical study of NASDAQ stocks. Journal of Finance 33, 1153–1172.