



ELSEVIER

Journal of Financial Markets 8 (2005) 377–399

Journal of
FINANCIAL
MARKETS

www.elsevier.com/locate/finmar

Duration, volume and volatility impact of trades[☆]

Simone Manganelli*

European Central Bank, DG-Research, Kaiserstrasse 29, 60311 Frankfurt am Main, Germany

Available online 19 July 2005

Abstract

This paper presents a framework to model duration, volume and returns simultaneously, obtaining an econometric reduced form that incorporates causal and feedback effects among these variables. The methodology is applied to two groups of stocks, classified according to trade intensity. We find that: (1) all stocks exhibit trading volume clustering (which is significantly higher for frequently traded stocks); (2) times of greater activity coincide with a higher number of informed traders present in the market only for the frequently traded stocks; (3) the more frequently traded stocks converge more rapidly (in calendar time) to their long-run equilibrium, after an initial perturbation.

© 2005 Elsevier B.V. All rights reserved.

JEL classification: C32; G14

Keywords: Autoregressive conditional duration; GARCH; Ultra high frequency data; Empirical market microstructure

[☆]I would like to thank Rob Engle for the many helpful discussions and for his valuable comments. I also would like to thank an anonymous referee, Bruce Lehmann (the editor), Wouter Den Haan, Graham Elliott, Clive Granger, Philipp Hartmann, Lutz Kilian and the participants in the Conference on Market Microstructure and High-Frequency Data in Finance held in Sønderborg (Denmark) in August 2001. The data set used in the analysis was constructed together with Rob Engle and Andrew Patton.

*Tel.: +49 69 1344 7347; fax: +49 69 1344 6575.

E-mail address: simone.manganelli@ecb.int.

1. Introduction

How many trades and how many hours does it take for new information to be incorporated into prices? In which markets is information revealed faster? Do times of greater activity coincide with a higher number of informed traders present in the market? Do we observe a stronger clustering of trading volumes for firms with better public reporting? And, in particular, do the answers to these questions depend on the frequency with which stocks are traded? These are the main empirical questions addressed in this paper.

In theoretical models, trading in financial markets occurs either for information or for liquidity reasons. Informed agents are motivated by relevant private information about the real value of the asset not known to the rest of the population. They will act strategically in the attempt to make a profit out of this information. As a consequence, both the market maker and the uninformed agents face a learning problem: how to infer the true value of the traded asset by observing the behaviour of the other traders. They can learn by looking at market information, such as prices, volumes or times between trades. Several papers provide precise predictions about how these variables should be interrelated. In this paper we develop an econometric framework that can be used to test whether some of these predictions are supported by the data.

From a statistical point of view, each transaction can be viewed as the realisation of a stochastic process, whose variables of interest are the time between trades (duration), prices and volumes. A problem when dealing with ultra-high frequency data (such as transaction data) is that they are not regularly spaced. [Engle and Russell \(1998\)](#) suggest tackling this issue by modelling directly the duration between trades. [Zhang et al. \(2001\)](#) improve upon the original Engle and Russell model by allowing the expected duration to depend nonlinearly on past information variables. [Bauwens and Giot \(2000, 2003\)](#) suggest modeling the logarithm of durations. However, as pointed out by [Engle \(2002\)](#), this solution is not very satisfactory from a theoretical point of view, because it does not work if there are zeros in the dependent variable and it imposes implicit constraints on the distribution of the residuals. Moreover, modeling the logarithm makes it more difficult to study the link between the random variables of interest and their expectations (since the expectation of the logarithm is different from the logarithm of the expectation).

There are two main approaches in the empirical microstructure literature to study the impact of trades: (1) [Hasbrouck \(1991\)](#), who focuses on the dynamics of returns' first moments, and (2) [Engle \(2000\)](#), focusing on the dynamics of returns' second moments. [Hasbrouck \(1991\)](#) uses a vector autoregressive system to model the interactions among the economic variables of interest, namely prices and signed volume. [Dufour and Engle \(2000\)](#) incorporate durations in Hasbrouck's VAR. While these models provide useful insights about the dynamics of expected returns, it is not clear how they can be extended to study the interplay between time, volume and volatility. Since—as we will see in Section 2—the questions we are interested in concern exactly the relationship between these three variables, it is important to have an econometric framework where these relationships can be analysed directly.

In a different line of research, Engle (2000) incorporates the duration approach of Engle and Russell (1998) into a volatility framework. He analyses a time series of IBM stock transactions, suggesting a model of the timing of the trades and measuring the impact of this timing on price volatility. He also introduces economic variables—such as bid-ask spreads and large volume dummies—into the volatility specification, finding significant and positive effects. More recently, Grammig and Wellner (2002) suggest a model for the interdependence of intra-day volatility and trade duration processes, whose unknown parameters are estimated via GMM. A drawback of this approach is that it assumes that returns follow a weak GARCH model (see Drost and Nijman, 1993). Since in a weak GARCH setting only linear projections (and not conditional expectations) of the squared innovation process are considered, the results cannot be strictly interpreted as estimates of conditional variances.

The econometric contribution of this paper is to extend the approach of Engle (2000), by elaborating a system where returns and volatilities directly interact with duration and volume.¹ This is accomplished in two steps. First, the modeling idea behind the Autoregressive Conditional Duration by Engle and Russell (1998) is extended to model volumes. We treat volumes as a stochastic process and model it using an autoregressive specification that multiplies an i.i.d. error term. Since volume can take only positive values, the support of the distribution of the error terms is the positive real line.

Next, duration, volume and returns are modeled simultaneously, using a special type of vector autoregression. We allow expected duration, expected volume and variance of returns to depend on current and lagged values of the variables under study. In this way, we can build a system that incorporates causal and feedback effects among these variables. We also construct impulse–response functions that show how the system reacts to a perturbation of its long-run equilibrium. The impulse–response function analysis will be helpful to understand market “resiliency”, i.e., the speed with which prices tend to converge to the underlying liquidation value of the asset.

This modeling methodology can be easily extended to incorporate any number of non-negative variables (such as spreads, the number of transactions or the difference between high and low prices in any given period). The Multiplicative Error Model of Engle (2002) provides a framework for multivariate non-negative processes, generalizing the approach proposed in this paper.

The econometric methodology proposed in this paper is applied to study the characteristics and long-run equilibrium properties of different markets. The model is estimated with a sample of ten stocks from the NYSE, which covers the period from January 1998 to June 1999. The stocks are divided into two groups according to the intensity with which they are traded. In the empirical analysis, we find strong evidence that the dynamics of frequently traded stocks differ significantly from those of the infrequently traded ones.

¹Since we are dealing with ultra high frequency data, throughout the paper we use the term volume as a synonym of trade order size.

We test three hypotheses, in particular. First, we want to check whether we observe clustering in trading volumes, as predicted by some microstructure models (Admati and Pfleiderer, 1988; Foster and Viswanathan, 1990). We find that, consistent with theoretical predictions, volume is a highly persistent process. That is, the empirical regularities found for duration and volatility models are confirmed also for high frequency volumes. In addition, we find that volume autocorrelation (our measure of persistence) is significantly stronger for more frequently traded stocks. This finding is consistent with the prediction by Foster and Viswanathan (1990): informational advantages for firms with better and more frequent public reporting (proxied by the frequency of trading) are shorter lived, and therefore force informed traders to trade more aggressively.

Second, our results show that times of greater activity coincide with a higher number of informed traders present in the market only for the frequently traded stocks. Results for the infrequently traded stocks show a very weak relationship between duration, volume and price volatilities. This suggests that the predictions of Easley and O'Hara (1992) are valid only for frequently traded stocks.

A third interesting finding indicates that the more frequently traded the stock, the faster the market returns to its full information equilibrium after an initial perturbation. This result confirms that infrequently traded stocks might be characterised by a different mechanism of information transmission. In particular, this is consistent with the theoretical predictions of Holden and Subrahmanyam (1992), who show that an increase in the number of transactions causes more information to be revealed earlier. Our results are also consistent with those provided by Zhang et al. (2001), who find strong evidence that fast and slow transacting periods have quite different dynamics, although their empirical analysis is limited to the IBM stock.

The paper is structured as follows. In the next section we discuss the testable empirical hypotheses derived from existing market microstructure theories. In Section 3 we introduce our econometric model. Section 4 presents the empirical results. Section 5 concludes and suggests new directions for future research.

2. The testable hypotheses

There are at least three possible ways through which economic agents can learn from observing market information: prices, volumes and trade duration. Market microstructure research has devoted considerable efforts to explaining why and how these variables interact and can convey information. The basic framework was set up by Bagehot (1971) and assumes that in the market there exist traders with superior information who try to transform their informational advantage into net profits. A common feature of this type of model is that new information is incorporated into prices through the interaction of informed and uninformed traders. The underlying assumption is that prices are only semi-strong form efficient, in the sense that they reflect all public information, but not the private information. If this private information is valuable, then informed traders can make positive profits. Many

possible extensions of this basic idea have been considered. See, for example, Copeland and Galai (1983), Kyle (1985), Glosten and Milgrom (1985), Easley and O'Hara (1987, 1992), Diamond and Verrecchia (1987), Admati and Pfleiderer (1988). See O'Hara (1995) for a comprehensive survey.

Although the early microstructure models did not recognise any explicit role for time and volume, several contributions have now emphasized the impact of the random arrival of new pieces of information in the information disclosure process. The general intuition is that market activity and volatility change over time because new information becomes available to traders at a varying rate. When not much information is available, trading is slow and there are few price changes. On the contrary, when new unexpected information hits the markets, trades become more frequent, volumes increase and prices move much faster. The implication is that volumes and time between trades influence prices because they are correlated with private information about the stock's true value. Since informed traders wish to trade more frequently and in larger quantities when they have valuable information, a rational market maker will interpret large orders and short durations as evidence of trading by informed agents and will adjust beliefs and prices accordingly (Easley and O'Hara, 1987, 1992). At the same time the lack of trade can be associated with the fact that no new information exists and that prices efficiently incorporate all the available information. More precisely, after a period in which no trade occurs, the market maker updates her beliefs, raising the probability she attaches to the no-information event and to the prior expected value of the asset (Easley and O'Hara, 1992).

In this paper I will address three specific empirical questions. For each of these questions, I first briefly discuss the theoretical background and then derive some testable implications.

Question No. 1: Do we observe clustering in trading volumes?

Admati and Pfleiderer (1988) and Foster and Viswanathan (1990) provide models to explain time dependent patterns in securities trading. The modeling structure is the one introduced by Kyle (1985), with a market maker, informed traders and liquidity traders. Both studies allow the possibility that liquidity traders may have some discretion regarding the timing of their trading. Admati and Pfleiderer (1988) show that, in equilibrium, all discretionary liquidity traders choose to trade together. Therefore, according to their model, we should observe a clustering of trading volume at some time of the day. Foster and Viswanathan (1990) consider a variation of this model. They allow new information about the security to arrive in each trading period. However, this private information advantage is eroded both by the disclosure of public information and by the market maker's learning process. These elements induce informed agents to trade more intensely, thus causing private information to be revealed more quickly and inducing a clustering in trading volumes.

Translated into statistical terms, both models predict a strong autocorrelation in the volume process: high-volume trades tend to be followed by high-volume trades, and, vice versa, low-volume trades tend to be followed by low-volume trades. In addition, Foster and Viswanathan (1990) predict that these clustering effects should

be stronger for firms with better public reporting and more discretionary liquidity traders. The reason is that for these firms informational advantages are shorter lived, because they will be eventually eliminated by public announcements. Moreover, the decision of discretionary liquidity traders to postpone their trading until the announcement has been released will make it easier for the market maker to infer the informed agent's reasons for trading. There are two testable implications of these models:

- (a) *volumes should exhibit a certain degree of autocorrelation (trading volume clustering);*
- (b) *volume autocorrelation should be stronger for more frequently traded firms (assuming, as in Foster and Viswanathan (1993), that the frequency of trading is a good proxy for the quality of public information).*

Question No. 2: Do times of greater activity coincide with a higher number of informed traders present in the market?

The key intuition of Easley and O'Hara (1992) is that, if time can be correlated with information about the true value of an asset, then both trade and lack of trade can convey information to market participants. In particular, agents can learn not only about the direction of any new information by looking at trades, but also whether any new information exists by looking at the absence of trades. The main implication of this analysis is that time is no longer exogenous to the price process. In addition, volume plays a role. Since trades are positively serially correlated, the probability of a trade (and therefore the probability of observing a transaction price) will be positively correlated with volume. The intuition behind this result is that since informed agents always trade, expected duration will be lower and expected volumes will be higher after an information event occurs. In other words, the greater the volume, the less likely the absence of trade and therefore the more likely that new information exists.² A consequence of this is that (i) the duration between trades can affect prices, (ii) duration should be correlated with volume, and (iii) both duration and volume should be correlated with transaction price variances. The testable implications of this model are:

- (c) *short durations and high volumes should decrease the expected duration and increase the expected volume of the next trade;*
- (d) *short durations and high volumes should increase the price variance of the next trade.*

Question No. 3: How long does it take for new information to be impounded into prices?

²In the model of Easley and O'Hara (1992) agents are allowed to trade only one unit of the asset at each period. Therefore, they define volume as the cumulative sum of buys and sells up to time t . The use of this definition of volume would present serious challenges in an empirical analysis, because it would result in a non stationary process interacting with stationary processes, such as duration and volatility. To circumvent this difficulty, in our empirical analysis we use order size instead.

Holden and Subrahmanyam (1992) develop a multi-period auction model in which privately informed agents act strategically to exploit their informational advantage. The basic set up is the one developed by Kyle (1985), who shows that a single informed trader tends to trade gradually, in order to maximise the benefits deriving from his monopoly power. Holden and Subrahmanyam (1992) extend this model to incorporate multiple informed agents and show that they tend to trade more aggressively than their monopolist counterparts. The driving element of this result is the competition between informed traders. At the limit, as the number of traders goes to infinity, all their private information is revealed immediately. The same result holds with just two informed traders, as the number of auctions (i.e., trades) increases and tends to infinity. They illustrate the difference between the imperfectly competitive case and the monopolist case with a series of numerical simulations. In Fig. 2, they show how price variance declines very rapidly to zero through time. As the number of trades increases, the speed with which it drops increases very sharply. The testable implication of this model is the following:

- (e) *price variance of frequently traded stocks should converge more rapidly than that of infrequently traded stocks to their long run equilibrium, after an initial perturbation;*

Answers to this last question might have considerable practical relevance. Theoretical models show that prices do eventually converge to full information values. However, this convergence takes place only at the limit and not much is said about how long this adjustment process takes. Knowing the speed at which new information is incorporated into prices might shed some light on the nature of market efficiency. In addition, research along these lines can yield new insights on how the institutional features of different markets contribute to the attainment of this type of efficiency.

3. Econometric methodology

When studying market microstructure, a major problem faced by the econometrician is that transaction data arrive at irregular times. The practice has been to ignore the time occurring between two trades and analyze the data with fixed time interval econometric techniques, in many cases by taking averages of the values of interest over a given, arbitrary interval.³ However, as the previous discussion points out, there might be a significant loss of information by adopting this modeling strategy, since the time elapsing between two trades may be informative about the behaviour of the traders.

The Autoregressive Conditional Duration (ACD) model of Engle and Russell (1998) provides a solution to this problem. Engle and Russell model the arrival times

³See, for example, Hartmann (1999), Fleming and Remolona (1999), Ederington and Lee (1993), Jones et al. (1994).

of trades as random variables that follow a point process. Associated with each arrival time are random variables called marks, such as volume, spread and price. It is often the case that the variables of interest are those in the marks, for example the price and the volatility of the asset under study. The objective of this paper is to model explicitly the components of the mark and to construct a general framework to study the interaction among them.

Let $d_t = z_t - z_{t-1}$ be the duration between two consecutive trades, where z_t denotes the time of the t th transaction. At the t th transaction, the t th mark x_t is realised. Typically, x_t will be a vector containing the price and the volume of the transaction, plus the bid-ask spread posted by the market maker. These data can be viewed as the sample realization from a hypothetical stochastic process. We assume that the true DGP generating each pair (d_t, x_t) can be written as follows:

$$(d_t, x_t) \sim f(d_t, x_t | \Omega_t; \theta), \quad (1)$$

where Ω_t denotes the information available at time t and θ is a (finite) vector of parameters of interest.

Engle (2000) points out that this joint density can be rewritten as the product of the marginal density of d times the conditional density of x given d . Assuming that the marks of interest are just returns (y) and volume (v), we can write the conditional density of x given d as the marginal conditional density of v given d times the density of y conditional on d and v :

$$\begin{aligned} (d_t, v_t, y_t) &\sim f(d_t, v_t, y_t | \Omega_t; \theta) \\ &= g(d_t | \Omega_t; \theta_1) \cdot h(v_t | d_t, \Omega_t; \theta_2) \cdot k(y_t | d_t, v_t, \Omega_t; \theta_3). \end{aligned} \quad (2)$$

Although other parameterizations are possible, the one proposed here seems natural, given the widespread use of strategic models in the market microstructure literature. In Kyle's (1985) model, for example, insider traders act strategically, i.e., they take into account the effect their orders have on price, by conditioning on the behavior of both the market maker and uninformed traders. At the moment in which new private information becomes available, informed traders face a given supply/demand schedule from the market maker. They will thus choose the amount of shares to trade in order to maximise their informational advantage, discounting the market maker reaction to their trade. Since the timing of the trade as well can convey information to the market maker (see Easley and O'Hara, 1992), informed traders will take this aspect into account also when determining the timing of their trades. This implies that both duration and volume will determine the price at time t . A similar argument can be used regarding the relationship between duration and volume. When new private information arrives, the informed trader would like to exploit it by trading a larger amount of shares. However, by following this strategy she would immediately reveal herself to be an informed trader and the market maker would adjust prices accordingly. In order to hide her type, the informed trader will split her buy and sell orders in many trades. It is then this trading intensity that determines the amount of shares to be sold or bought at each trade. These arguments support the parameterization used in (2), suggesting a plausible causality relation running from duration to volume and from duration and volume to prices.

By modeling each marginal and conditional distribution in (2), it is possible to develop a complete framework for the triple (d_t, v_t, y_t) . As discussed in Section 2 economic theory suggests that clusters of activity should be observed in the market every time some unexpected piece of information arrives. In terms of duration, this translates in periods in which we observe very frequent transactions, with very short durations. The ACD models pick up this aspect.

The typical ACD model can be stated as follows:

$$\begin{aligned} d_t &= \psi_t \varepsilon_t, & \varepsilon_t &\sim \text{i.i.d.}(1, \sigma_\varepsilon^2) \\ \psi_t &\equiv E(d_t | \Omega_t; \theta_d), \end{aligned} \quad (3)$$

where Ω_t is, as usual, the information available at time t , and θ_d is a vector of parameters that need to be estimated.

Since ψ_t is the expectation of d_t , ε_t must be a random variable with mean 1. Moreover, both ψ_t and the support of the distribution of ε_t must be positive, as duration cannot be negative. Engle and Russell (1998) suggest modeling ψ_t as an autoregressive process. This choice is reasonable, as the autoregression can be interpreted as the statistical way to model the clustering of activity discussed above. Moreover, autoregressive processes are consistent with the fact that the whole sequence of past trades might become informative. Pushing along these lines, it seems plausible to model volume in a similar fashion. A possible model for volume could be the following:⁴

$$\begin{aligned} v_t &= \phi_t \eta_t, & \eta_t &\sim \text{i.i.d.}(1, \sigma_\eta^2) \\ \phi_t &\equiv E(v_t | \Omega_t; \theta_v). \end{aligned} \quad (4)$$

As in the ACD model, η_t must be an i.i.d. error term with mean 1. Since volume is a variable that assumes only positive values, the distribution of the error term will be defined only for positive values. By imposing different distributional assumptions on this error term and specifying the process followed by the expected volume ϕ_t , we can get a variety of models. Moreover, the quasi-maximum likelihood results of GARCH and ACD models can be directly extended to this model: imposing an exponential distribution for η_t , it is possible to obtain consistent estimates of the parameters θ_v , provided that the expected volume ϕ_t is correctly specified.⁵

Given the predictions of economic theories and the known empirical regularities that characterize durations and volatilities, we model expected volume as an autoregressive process. To stress the analogy with duration and volatility models, we call these models autoregressive conditional volume (ACV). The simplest possible

⁴Here we assume that volume is a stationary process. In practice, however, volume might have a tendency to increase over time. For example, while the average daily trading volume in the broad-based S&P 500 index was around 190 million shares back in 1993, it has since increased steadily to around 1.5 billion shares in 2000. Given the relatively short sample period used in our empirical analysis, this is unlikely to create any problem. Nevertheless, when longer time series are studied, one might want to take this aspect into consideration when data are cleaned for the analysis.

⁵See Bollerslev and Wooldridge (1992) and Engle and Russell (1998) for details.

specification is an ACV(1,1):

$$\phi_t = \omega + \alpha v_{t-1} + \beta \phi_{t-1}. \tag{5}$$

Other more general ACV specifications can be introduced. However, instead of pursuing this strategy, we propose a more general approach. Modeling the returns with a GARCH process whose mean and variance are denoted by μ_t and σ_t^2 , and specifying a model for each of the marginal and conditional density functions in (2), we can get a complete and general framework for the random variables of interest. More precisely, we assume the following:

$$\begin{aligned} d_t &= \psi_t(\theta_d; \Omega_t) \varepsilon_t, & \varepsilon_t &\sim \text{i.i.d.}(1, \sigma_\varepsilon^2) \\ v_t &= \phi_t(\theta_v; d_t, \Omega_t) \eta_t, & \eta_t &\sim \text{i.i.d.}(1, \sigma_\eta^2) \\ y_t &= \mu_t + \sigma_t(\theta_y; d_t, v_t, \Omega_t) \zeta_t, & \zeta_t &\sim \text{i.i.d.}(0, 1). \end{aligned} \tag{6}$$

Note that in this framework the error terms are uncorrelated with each other by construction. Taken separately, these are, respectively, the ACD, ACV and GARCH models. However, under this framework it is possible to allow interactions among the three variables under study.⁶ According to how we specify the functional forms for expected duration, expected volume and variance, we can get many different models and construct an econometric reduced form to evaluate the theoretical predictions from market microstructure. One possible specification is the following:

$$\begin{aligned} \psi_t &= a_0 + \sum_{i=1}^q (a_{1i} \psi_{t-i} + a_{2i} \phi_{t-i} + a_{3i} \sigma_{t-i}^2) + \sum_{j=1}^p (a_{4j} d_{t-j} + a_{5j} v_{t-j} + a_{6j} y_{t-j}^2), \\ \phi_t &= b_0 + \sum_{i=1}^q (b_{1i} \psi_{t-i} + b_{2i} \phi_{t-i} + b_{3i} \sigma_{t-i}^2) + b_4 d_t \\ &\quad + \sum_{j=1}^p (b_{5j} d_{t-j} + b_{6j} v_{t-j} + b_{7j} y_{t-j}^2), \\ \sigma_t^2 &= c_0 + \sum_{i=1}^q (c_{1i} \psi_{t-i} + c_{2i} \phi_{t-i} + c_{3i} \sigma_{t-i}^2) + c_4 d_t + c_5 v_t \\ &\quad + \sum_{j=1}^p (c_{6j} d_{t-j} + c_{7j} v_{t-j} + c_{8j} y_{t-j}^2). \end{aligned} \tag{7}$$

The nice feature of this specification is that it has a vector autoregressive moving average representation. Indeed, rewriting (7) in matrix form, we get, in more compact notation:

$$\chi_t = \gamma + A_1 \chi_{t-1} + \dots + A_q \chi_{t-q} + B_0 \tau_t + B_1 \tau_{t-1} + \dots + B_p \tau_{t-p}, \tag{8}$$

⁶Model (6) provides a link between [Hasbrouck \(1991\)](#) and [Engle \(2000\)](#), the two approaches discussed in the introduction. By incorporating the trade sign into the specification of μ_t and adding an extra equation for the trade sign, it is possible to merge into a single empirical specification the approaches of [Hasbrouck \(1991\)](#) and [Engle \(2000\)](#). I am grateful to an anonymous referee for pointing this out.

where $\chi'_t = (\psi_t, \phi_t, \sigma_t^2)$, $\tau'_t = (d_t, v_t, y_t^2)$, γ is a (3,1) vector of coefficients and $A_1, \dots, A_q, B_0, B_1, \dots, B_p$ are (3,3) matrices of coefficients.

To keep things simple, in this paper we concentrate our attention on a first order model of (8):

$$\chi_t = \gamma + A\chi_{t-1} + B\tau_t + C\tau_{t-1}. \tag{9}$$

The parameterization in (2) imposes that B is lower diagonal with all the elements on the diagonal equal to zero. Further restrictions can be imposed. For example, by imposing that A and C are diagonal and B is equal to zero, we get back the three independent models ACD, ACV and GARCH. Note that assuming that A is diagonal is equivalent to assuming that the parameters of this system are variation free as in Engle et al. (1983). Under this assumption, the likelihood in (2) can be rewritten as

$$f(d_t, v_t, y_t | \Omega_t; \theta) = g(d_t | \Omega_t; \theta_d) \cdot h(v_t | d_t, \Omega_t; \theta_v) \cdot k(y_t | d_t, v_t, \Omega_t; \theta_y), \tag{10}$$

where $\theta' = (\theta'_d, \theta'_v, \theta'_y)$. This makes it possible to estimate the three models separately.

The VAR system in (8) will be covariance stationary if and only if $|\lambda| < 1$ for all values of λ satisfying $|I_n \lambda^p - H_1 \lambda^{p-1} - H_2 \lambda^{p-2} \dots - H_p| = 0$, where $H_i = (B_0 - I)^{-1} \times (B_i + A_i)$, $i = 1, \dots, p$.

The model described by (6) and (9) allows one to compute the effect that an unexpected trade today has on future expected duration, volume and volatility. Suppose the system is in steady state at $\chi_0 = (I - A - B - C)^{-1} \gamma$. Then the impulse–response function of the model (6) and (9) for $t > 0$ is $\partial E(\chi_t | \Omega_0) / \partial \xi'_0 \equiv \Phi_t = D^{t-1} (I - B)^{-1} (AB + C) \text{diag}(\chi_0)$, where $D = (I - B)^{-1} (A + C)$, and $\xi'_0 \equiv (\varepsilon_0, \eta_0, \zeta_0)$.⁷

The standard errors for the impulse–response can be computed as follows. Let $\theta = [\theta'_d, \theta'_v, \theta'_y]'$, $\kappa_t \equiv \text{vec}(\Phi_t(\theta))$ and $\sqrt{T}(\hat{\theta} - \theta) \xrightarrow{a} N(0, Q)$. Then, under standard regularity conditions, $\sqrt{T}(\hat{\kappa}_t - \kappa_t) \xrightarrow{a} N(0, G_t Q G'_t)$, where $G_t = \frac{\partial \kappa_t}{\partial \theta'}$. The i th column of the matrix of derivatives G_t can be easily estimated numerically by $G^i_t = [\kappa_t(\hat{\theta} + e_i \Delta) - \kappa_t(\hat{\theta})] / \Delta$, where e_i denotes the i th column of the $(p \times p)$ identity matrix and Δ is a small number.

4. Empirical analysis

4.1. Data

The econometric models discussed in the previous section were tested on a sample of ten stocks. The data are taken from the Trades and Quotes (TAQ) database. All stocks are quoted in the NYSE and cover the period from January 1, 1998 to June 30, 1999. They were randomly chosen using the following procedure. We constructed

⁷See Proposition 2 of Manganelli (2002) for a derivation of this result.

Table 1
Stocks used in the analysis

A. Frequently traded		B. Infrequently traded	
CP	CDN PACIFIC	DTC	DOMTAR INC
GAP	GREAT A & P	FTD	FORT DEARBN INCM
COX	COX COMMUNICATION	GBX	GREENBRIER CO
DLP	DELTA & PINELAND	GSE	GUNDLE/SLT ENV
AVT	AVNET INC	JAX	J ALEXANDERS

ten deciles on the basis of the 1997 total number of trades of all the stocks quoted on the NYSE. Then we randomly picked five stocks from the eighth decile and five from the second decile. All the stocks that didn't survive the whole period of study were discarded and substituted by other randomly chosen stocks from the same decile. The tickers and names of the ten stocks are reported in Table 1.

To prepare the data for the analysis, first we dropped any transaction that occurred before the opening time, 9:30 am, and after the closing time, 16:00. Then, we computed the durations between trades treating the overnight period as if it didn't exist, so that, for example, the time elapsing between 15:59:30 and 9:30:10 of the following day is only 40 seconds. We adopted this strategy because in our sample we have stocks that are very rarely traded. Eliminating the durations for the overnight period would have caused the loss of important data for these stocks. Third, we computed the price of each transaction as the average of the bid and ask quotes that appear at least five seconds before the transaction itself. This procedure is standard in market microstructure studies. Taking the average of the bid and ask quotes limits the bid-ask bounce problems, while considering only 5 seconds old quotes was originally proposed by Lee and Ready (1991) and its purpose is to correct reporting errors in the sequence of trades and quotes.⁸ Fourth, we eliminated all the transaction data with zero duration. We treated these transactions as one single transaction, summing up all the volumes. Fifth, we adjusted the data for stock splits, by multiplying prices and volumes by the stock split ratio. Sixth, the returns were computed as the difference of the log of the prices. Seventh, we adjusted the data for dividend payments and trading halts, by simply deleting the first observation whose return incorporated the dividend payment or the trading halt.

It is a well-known fact that both durations and volatilities exhibit a typical daily pattern over the course of the trading day, with very high trading activity at the beginning and at the end of the day. In order to remove this feature from the data, the time series of durations, prices and volumes were diurnally adjusted as in Engle

⁸In the following empirical analysis, all the variables of interest are measured whenever a transaction occurs. However, in the real world, there is a difference between a transaction and a print. Consider, for example, a 1000-share market order for which the specialist takes 600 shares and the rest is passed onto the limit order book. This will show up in the TAQ data set as two prints, but in fact it was just one transaction with three counterparties. This is a limitation of the data set that is hard to overcome. The only way to get around this problem is use order data, which are rarely available.

Table 2

Summary statistics for the ten stocks used in the analysis. The sharp difference between the number of trades for the two groups of stocks reflects the selection criteria used. Notice that this difference is reflected in the average and median duration, but not in volume.

Ticker	# Obs	Duration		Volume	
		Average	Median	Average	Median
DLP	65304	134.15	45	1483.84	500
GAP	46826	187.3	80	827.03	300
CP	71672	122.43	65	2978.55	500
COX	88917	98.6	35	2692.57	800
AVT	58389	150.02	64	1075.55	400
JAX	2761	3165.7	1032	1002.32	500
GSE	1968	4438.42	1523	1523.53	600
GBX	5154	1695.49	620	1483.95	500
FTD	3264	2416.44	1390	736.95	500
DTC	4161	2094.22	526	2143.91	800

(2000). We regressed durations, volumes and absolute values of returns on a piecewise linear spline with knots at 9:30, 10:00, 11:00, 12:00, 13:00, 14:00, 15:00, 15:30 and 16:00 and constructed the adjusted series by dividing each original series by the spline forecast. The two extra knots in the first and last half-hour reflect the typical different trading intensity during the day. Less frequently traded stocks, in general, do not exhibit any regular pattern during the day. More frequently traded stocks, instead, are characterised by the typical inverted U pattern for duration, an L pattern for absolute returns (i.e., high absolute returns at the opening and flat shape during the rest of the day), and no regular pattern for volume.

Throughout the empirical section, we concentrate our attention on volume and variance per transaction. Even if volatility per unit of time is the natural and commonly used measure of volatility, in market microstructure we observe only transaction data. Transforming these data into per unit data would bias the analysis. Indeed, some of the theories predict that the absence of trade must be interpreted as if bad news has arrived, and hence trades with longer durations should have a higher price impact, measured in transaction (not calendar) time. In this case, if we divide returns by the square root of the durations (in order to get volatility per second), the price impact of the transaction will be considerably reduced, since long durations will almost always be associated with low returns per second.

In Table 2 we report some summary statistics for the ten stocks used in the analysis. For each stock we report the number of observations, and average and median of duration and volume. The sharp difference between the number of observations for the two groups of stocks reflects the criteria used to select them. It is interesting to note that this difference is reflected in the average and median duration of the trades, but not in the volumes, whose average magnitude seems comparable across the two groups of stocks.

Table 3

Autoregressive coefficients (β) of the three models for duration, volume and trade variance

- $d_t = \psi_t \varepsilon_t, \quad \varepsilon_t \sim \text{i.i.d.}(1, \sigma_\varepsilon^2)$
 $\psi_t = \omega + \alpha d_{t-1} + \beta \psi_{t-1}$.
- $v_t = \phi_t \eta_t, \quad \eta_t \sim \text{i.i.d.}(1, \sigma_\eta^2)$
 $\phi_t = \omega + \alpha v_{t-1} + \beta \phi_{t-1}$.
- $y_t = \sigma_t \zeta_t, \quad \zeta_t \sim \text{i.i.d.}(0, 1)$
 $\sigma_t^2 = \omega + \alpha y_{t-1}^2 + \beta \sigma_{t-1}^2$

These models were estimated imposing the strong exogeneity assumption. In terms of the notation developed in the paper, we set the matrix B equal to zero and imposed the matrices A and C to be diagonal. Standard errors are omitted, because all the coefficients were statistically significant at the 1% confidence level.

Ticker	ACD	ACV	GARCH
DLP	0.9324	0.9525	0.9618
GAP	0.9521	0.9726	0.9658
CP	0.9651	0.9707	0.9583
COX	0.9373	0.9718	0.9505
AVT	0.8950	0.9154	0.9508
JAX	0.9533	0.7163	0.9072
GSE	0.8303	0.8496	0.9493
GBX	0.9295	0.7154	0.9149
FTD	0.8764	0.2925	0.9487
DTC	0.8857	0.8004	0.9519

4.2. Results

Question No. 1: Do we observe clustering in trading volumes?

First we estimated the three independent ACD, ACV and GARCH models. In terms of the general model described in the previous section, we set the matrix B equal to zero and impose the matrices A and C to be diagonal. In Table 3 we report the estimated autoregressive coefficients (the diagonal elements of A) of the three models. All these coefficients are highly significant. We omit to report the t -statistics in order to present a cleaner picture of the results.

This exercise is novel in two aspects. First the ACV model is estimated and second these models are estimated with two groups of stocks, frequently and infrequently traded stocks. Volume appears to be a very persistent process for the frequently traded stocks, as indicated by the autoregressive coefficients, which are always above 0.9, and above 0.95 for four out of five stocks. This finding confirms that the empirical regularities found for duration and volatility models hold for volume as well. It is also consistent with the predictions of market microstructure theories,

according to which markets should be very active every time an unexpected piece of information arrives or when there is a clustering of liquidity traders.⁹ The persistence, however, is significantly lower for the five less frequently traded stocks. Here the autoregressive coefficient drops to 0.7–0.8, and down to 0.3 in one case. This is a first piece of evidence that the dynamics of frequently traded stocks differ significantly from those of infrequently traded stocks. Looking also at the ACD and GARCH models, we notice an analogous drop in the level of the autoregressive coefficients for the infrequently traded stocks, but much less dramatic. This provides empirical evidence in favour of the hypotheses (a) and (b) described in Section 2. Volume does exhibit a certain degree of autocorrelation, as predicted by the models of Admati and Pfleiderer (1988) and Foster and Viswanathan (1993). Moreover, given that frequently traded stocks are also more likely to be characterized by better disclosure of public information, the higher clustering of trading volumes for frequently traded stocks is consistent with the insights of Foster and Viswanathan (1993): since private information advantage is shorter lived for frequently traded stocks, informed agents will try to trade more intensively, thus inducing a greater clustering in trading volume.

Question No. 2: Do times of greater activity coincide with a higher number of informed traders present in the market?

With only this framework at hand, it is not possible to address question 2. We need models that allow interactions among the three variables under study. In the estimation, we imposed the assumption of weak exogeneity (that is we assume that the matrix A in (9) is diagonal), so that the optimization could be done separately for each variable, and we estimated the following model:

$$\begin{aligned} \psi_t &= a_0 + a_1\psi_{t-1} + a_2d_{t-1} + a_3v_{t-1} + a_4\mu_{t-1}^2, \\ \phi_t &= b_0 + b_1\phi_{t-1} + b_2v_{t-1} + b_3d_t + b_4d_{t-1} + b_5\mu_{t-1}^2, \\ \left| \begin{aligned} y_t &= X_t'\beta + u_t, & u_t &= \rho u_{t-1} + \mu_t + \theta\mu_{t-1}, \\ \sigma_t^2 &= c_0 + c_1\sigma_{t-1}^2 + c_2\mu_{t-1}^2 + c_3d_t + c_4d_{t-1} + c_5v_t + c_6v_{t-1}. \end{aligned} \right. \end{aligned} \quad (11)$$

In (11) we model the mean of y_t as an ARMA(1,1). The vector X_t contains current and lagged values of d_t and v_t . We don't report the results for the mean because our interest lies principally in the second moment dynamics.

In Table 4 we report the Ljung–Box statistics with 15 lags for raw and fitted data for duration, volume and returns. The statistics for the raw data are overwhelmingly significant, indicating the existence of strong autocorrelation in the variables. The analysis of the LB statistics for the residuals of the fitted models reveals that these models succeed in capturing most of the autocorrelation, as indicated by the sharp drop in value of the statistic. In some cases, the Ljung–Box statistic is still significant, but this is common with such large time series (see, for example, Engle, 2000).

⁹Although the persistence of volume at the ultra high frequency level is a new finding in the empirical literature of market microstructure, similar evidence was found by Hartmann (1999) who shows how the log of daily foreign exchange market trading volumes displays conditional heteroscedasticity.

Table 4

Ljung–Box (LB) statistics for raw data and fitted residuals. The fitted residuals are obtained from estimation of the full model of Eq. (12).

- $\psi_t = a_0 + a_1\psi_{t-1} + a_2d_{t-1} + a_3v_{t-1} + a_4\mu_{t-1}^2,$
- $\phi_t = b_0 + b_1\phi_{t-1} + b_2v_{t-1} + b_3d_t + b_4d_{t-1} + b_5\mu_{t-1}^2,$
- $\begin{cases} y_t = X_t'\beta + u_t, & u_t = \rho u_{t-1} + \mu_t + \theta\mu_{t-1}, \\ \sigma_t^2 = c_0 + c_1\sigma_{t-1}^2 + c_2\mu_{t-1}^2 + c_3d_t + c_4d_{t-1} + c_5v_t + c_6v_{t-1}. \end{cases}$

The sharp drop in value of the LB statistics reveals that this model successfully captures most of the autocorrelation present in the raw data.

	Duration		Volume		Return (mean)		Return (variance)	
	LB Raw	LB Fitted	LB Raw	LB Fitted	LB Raw	LB Fitted	LB Raw	LB Fitted
CP	3704.96	118.07	2513.63	31.94	376.17	12.18	3610.92	15.09
GAP	5052.96	56.16	357.02	58.16	45.78	19.62	3539.06	42.14
COX	30163.28	87.65	361.96	25.19	864.21	15.10	694.16	48.29
DLP	20027.87	114.03	2010.17	84.81	604.70	16.16	7922.59	65.76
AVT	6781.75	80.10	484.38	36.14	484.07	12.69	659.27	39.54
DTC	161.73	10.46	401.08	10.80	78.68	15.66	182.09	11.38
FTD	17.68	10.63	82.88	22.63	46.89	33.67	326.16	9.82
GBX	1023.15	15.05	389.99	12.25	31.09	13.84	77.60	11.79
GSE	85.56	6.81	82.69	25.89	37.77	18.77	71.15	15.57
JAX	104.33	7.50	95.01	14.84	31.66	17.46	330.35	11.84

To evaluate the effect that duration and volume have on prices, we need to look at their lagged values, because we are interested in the reaction of the market maker to the last trade. For this purpose, we need to take into account the effect that lagged duration and volume have on lagged expected volume and variance per trade. Rewriting (11) in matrix form as in (9), we get:

$$\begin{aligned} \chi_t &= \gamma + A\chi_{t-1} + B\tau_t + C\tau_{t-1} \\ &= \sum_{k=0}^{t-1} A^k\gamma + B\tau_t + \sum_{k=1}^{t-1} A^{k-1}(AB + C)\tau_{t-k} + A^{t-1}C\tau_0 + A^t\chi_0. \end{aligned} \tag{12}$$

The matrix of coefficients of τ_{t-1} is $(AB + C)$ and the null hypothesis that these coefficients are equal to 0 can be easily tested as shown at the end of Section 3. The results are reported in Table 5.

Let's start by looking at the variance equations. The coefficient of lagged duration in the variance equation is negative for 8 out of 10 stocks, but significant only in four cases. The sign of the duration coefficient is consistent with the results typically found in the empirical microstructure literature (see, for example, Engle, 2000; Dufour and Engle, 2000) and suggest that times of greater activity coincide with a

Table 5

Coefficients of the matrix (AB + C). These coefficients indicate the effect that lagged values of duration, volume and squared trade returns have on expected variables. Bollerslev–Wooldridge robust t -statistics in parenthesis. Variables significant at the 1% confidence level formatted in bold. The asterisk denotes significance at 5%.

	CP			GAP			COX			DLP			AVT		
	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2
ψ_t	0.0274 (23.90)	-0.0004 (-2.79)	-0.0001 (-3.19)	0.0450 (27.50)	-0.0013 (-3.37)	-0.0001 (-1.12)	0.0608 (44.70)	-0.0006 (-6.74)	-0.0002 (-3.48)	0.0679 (39.30)	-0.0026 (-5.36)	-0.0006 (-2.90)	0.0876 (36.08)	-0.0007 (-1.61)	-0.0007 (-7.38)
ϕ_t	-0.0011 (-1.55)	0.0227 (12.00)	0.0001 (0.36)	-0.0037 (-4.10)	0.0190 (10.40)	-0.0001 (-3.14)	0.0009 (1.34)	0.0241 (13.60)	0.0000 (-0.17)	0.0002 (0.33)	0.0408 (17.30)	0.0001 (1.60)	0.0028 (1.52)	0.0519 (11.29)	-0.0004 (-1.58)
σ_t^2	-0.0432 (-15.1)	0.0195 (3.68)	0.0421 (17.30)	-0.0043 (-0.88)	0.0198 (4.27)	0.0390 (14.30)	-0.0161 (-5.54)	0.0114 (4.91)	0.0407 (21.40)	-0.0033 (-1.18)	0.0013 (0.44)	0.0399 (9.31)	-0.0174 (-3.60)	0.0017 (0.54)	0.0373 (14.55)
	DTC			FTP			GBX			GSE			JAX		
	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2	d_{t-1}	v_{t-1}	μ_{t-1}^2
ψ_t	0.0850 (8.05)	0.0001 (0.03)	-0.0009 (-1.26)	0.032* (2.29)	-0.0430 (-46.9)	-0.0020 (-1.51)	0.0679 (10.03)	0.0000 (0.03)	0.0003 (0.8)	0.0964 (4.71)	-0.0089 (-1.59)	-0.0021 (-8.30)	0.0388 (5.92)	-0.0012 (-0.62)	-0.0002 (-0.67)
ϕ_t	-0.0110 (-1.58)	0.1210 (2.33)	0.0003 (0.19)	-0.0203 (-1.65)	0.1190 (5.02)	-0.0018 (-0.98)	-0.0087 (-1.33)	0.1429 (5.74)	0.0004 (0.21)	0.0012 (0.18)	0.0988 (4.90)	0.0016 (0.99)	-0.0078 (-0.91)	0.1530 (5.06)	0.0064 (1.09)
σ_t^2	0.045* (2.30)	0.0381 (1.90)	0.0650 (4.24)	-0.04* (-1.96)	0.068* (2.25)	0.0424 (5.49)	-0.0010 (-0.13)	0.0004 (0.11)	0.0588 (5.2)	0.0038 (0.17)	0.0282 (1.53)	0.0486 (3.94)	-0.0521 (-7.61)	0.029* (2.06)	0.0632 (7.25)

higher fraction of informed traders present in the market. The existing results, however, are limited to very liquid blue chip stocks.

The volume coefficient is always positive, but strongly significant only for the frequently traded stocks. This supports some of the predictions of Easley and O'Hara (1987, 1992), for example suggesting that trades with larger sizes are more likely to be executed by informed traders and thus have a greater impact on the price of the stock.

A very different picture emerges from the results for the infrequently traded stocks. Our findings provide evidence against the robustness of the above results, when the analysis is extended to less frequently traded stocks. In this case neither duration nor volume affect the variance. Thus the relationships among trading intensity, trade sizes and volatilities predicted by hypothesis (d) do not find confirmation for the less frequently traded stocks.

A second striking difference between the two groups of stocks is found in the duration equation. For the frequently traded stocks, both lagged volumes and squared returns affect the expected duration of the next trade. The signs of the two variables are always negative, implying that more frequent trading follows big price movements and high volumes. These results confirm those of Dufour and Engle (2000), who found that "short durations and thus high trading follow large returns and large trades," extending them to incorporate volume as well. However, the coefficients for the infrequently traded stocks are almost never significant. This suggests that the exogeneity of the trade arrival process imposed by Dufour and Engle (2000) and by some theoretical models may be a reasonable assumption only when dealing with infrequently traded stocks.

In the volume equation, only lagged volume is strongly significant, with lagged durations and squared returns playing practically no role for all the ten stocks. Therefore, we can conclude that hypothesis (c) finds partial confirmation only for frequently traded stocks.

The weak exogeneity assumption imposed in (11) could be restrictive, although this strategy has been adopted in a similar context by most empirical microstructure papers. As a robustness check, in Table 6, we test this restriction as suggested by Dufour and Engle (2000). We regress the standardized residuals of the three models against lagged expected variables. More precisely, we run the following three regressions:

$$\begin{aligned}\hat{\varepsilon}_t &= \alpha_0 + \alpha_1 \hat{\phi}_{t-1} + \alpha_2 \hat{\sigma}_{t-1}^2, \\ \hat{\eta}_t &= \beta_0 + \beta_1 \hat{\psi}_{t-1} + \beta_2 \hat{\sigma}_{t-1}^2, \\ \frac{\hat{\mu}_t}{\hat{\sigma}_t^2} &= \gamma_0 + \gamma_1 \hat{\psi}_{t-1} + \gamma_2 \hat{\phi}_{t-1}.\end{aligned}$$

If the model was correctly specified, then the standardised residuals should be i.i.d, and the estimated coefficients of the above regressions should not be significantly different from zero. On the other hand, if the expected variables belong to the model, then we should find significant correlation between the estimated residuals and the omitted variables. White (1980) type standard errors are reported in parentheses.

Table 6

Tests for omitted variables. We regressed the standardised residuals from model (12) on lagged expected variables:

$$\hat{\varepsilon}_t = \alpha_0 + \alpha_1 \hat{\phi}_{t-1} + \alpha_2 \hat{\sigma}_{t-1}^2,$$

$$\hat{\eta}_t = \beta_0 + \beta_1 \hat{\psi}_{t-1} + \beta_2 \hat{\sigma}_{t-1}^2,$$

$$\hat{\mu}_t / \hat{\sigma}_t^2 = \gamma_0 + \gamma_1 \hat{\psi}_{t-1} + \gamma_2 \hat{\phi}_{t-1}.$$

White robust *t*-statistics in parenthesis. Variables significant at the 1% confidence level are formatted in bold. The asterisk denotes significance at 5%.

	Exp Dur			Exp Vol			Volatility		
	<i>c</i>	Exp Vol	Volatility	<i>c</i>	Exp Dur	Volatility	<i>c</i>	Exp Dur	Exp Vol
CP	1.0068 (104.49)	0.0153* (2.13)	-0.0050 (-3.33)	0.9473 (17.70)	0.0228 (0.49)	0.0069 (1.70)	1.2054 (19.99)	-0.0846 (-1.46)	-0.1136 (-6.23)
GAP	0.9737 (44.58)	0.0004 (0.16)	0.0253 (1.23)	1.0011 (29.53)	0.0028 (0.10)	-0.0010 (-0.42)	1.3936 (15.18)	-0.1096 (-2.40)	-0.2840 (-5.88)
COX	0.9985 (120.59)	0.0123* (2.05)	-0.0027* (-2.16)	0.9676 (29.13)	0.0077 (0.30)	0.0059 (1.49)	1.0783 (45.57)	0.0082 (0.48)	-0.0800 (-6.46)
DLP	1.0079 (81.09)	-0.0019* (-2.31)	-0.0005 (-0.04)	1.0165 (67.34)	-0.0140 (-1.03)	-0.0005 (-0.52)	1.0682 (12.27)	-0.0264 (-1.33)	-0.0284 (-0.55)
AVT	0.9944 (67.13)	0.0054 (0.53)	0.0000 (0.01)	1.0259 (28.64)	-0.0075 (-0.29)	-0.0051 (-0.91)	1.0851 (24.91)	-0.0212 (-0.70)	-0.0546* (-2.01)
DTC	1.0129 (21.68)	0.0034 (0.10)	-0.0052 (-0.46)	0.9948 (9.72)	-0.0018 (-0.02)	0.0023 (0.12)	1.0260 (9.36)	0.0166 (0.17)	-0.0426 (-2.96)
FTD	0.9831 (8.73)	0.0427 (0.39)	-0.0054 (-0.61)	0.5641 (2.52)	0.3732 (1.76)	0.0181* (2.21)	-0.5320 (-0.47)	0.4833 (0.65)	1.0702 (2.39)
GBX	1.0209 (27.26)	-0.0246 (-1.15)	0.0010 (0.15)	1.0076 (15.20)	-0.0291 (-0.46)	0.0061 (0.60)	1.0299 (11.80)	0.0329 (0.50)	-0.0636* (-2.14)
GSE	0.9708 (13.77)	0.0185 (0.26)	0.0023 (0.16)	1.2034 (7.29)	-0.1797 (-1.37)	-0.0059 (-0.37)	1.1413 (4.32)	-0.0570 (-0.28)	-0.0783 (-0.77)
JAX	1.0095 (17.29)	0.0298 (0.66)	-0.0129 (-1.08)	1.0518 (8.70)	-0.0622 (-0.68)	0.0037 (0.16)	0.9273 (6.27)	0.1236 (0.97)	-0.0496 (-1.15)

The results show that the coefficients of the lagged expected variables are almost never significantly different from zero, especially for the infrequently traded stocks. Only the test for the variance specification shows evidence of misspecification for the frequently traded stocks, suggesting that lagged expected volume should also be taken into consideration.

Question No. 3: How long does it take for new information to be impounded into prices?

We can use our framework to compute the impulse-responses in calendar time, by using the formula derived at the end of Section 3. Since we are modeling

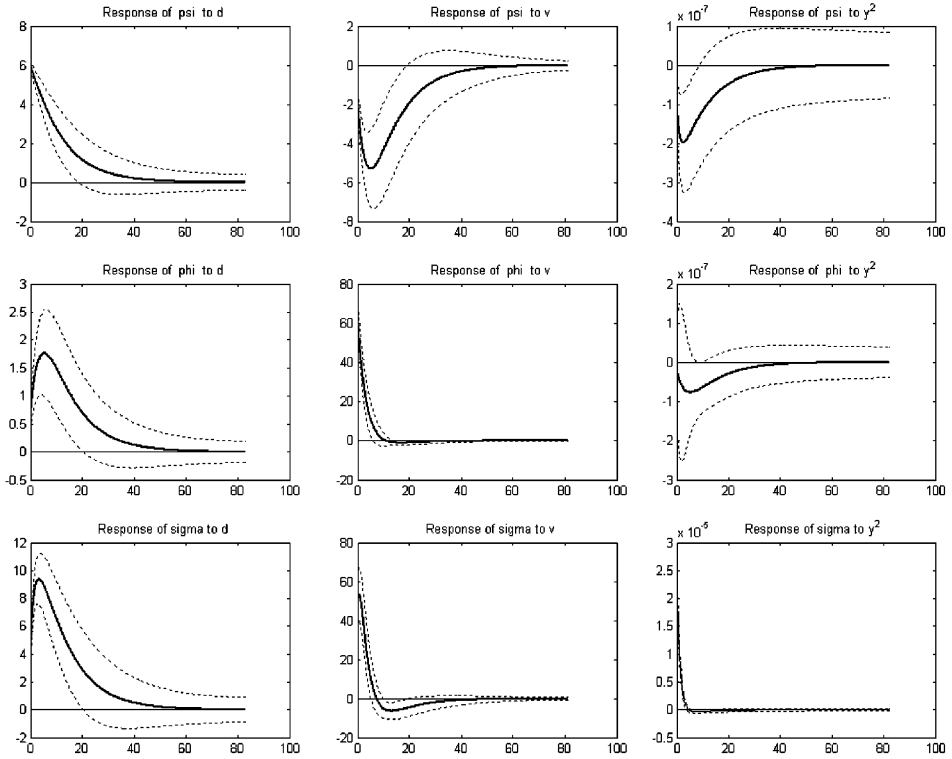


Fig. 1. Impulse–response function for COX in calendar time (hours).

simultaneously duration, volumes and volatility, it is possible to keep track of the expected duration at each transaction t .¹⁰ By progressively cumulating these durations, one can compute the expected timing of all future transactions. Obviously, the path of expected duration depends on the variable we decide to shock initially. In Figs. 1 and 2 we report the impulse–response functions for COX and JAX expressed in calendar time. Each expectation at time t (expected duration, expected volume and price variance) is plotted against the cumulative sum of the corresponding expected duration. It is interesting to note the striking similarity between these empirical impulse–responses and the simulated ones reported in Fig. 2 of Holden and Subrahmanyam (1992). Consistently with the theoretical predictions of their model, the speed with which price variance declines to zero after an initial shock increases dramatically with the trading frequency. Table 7 summarizes the results for the other stocks, confirming that the price variance of

¹⁰Since when cleaning the data we had diurnally adjusted the series (see Section 4.1), we put back in the system the unconditional means of the original series before computing the impulse–responses.

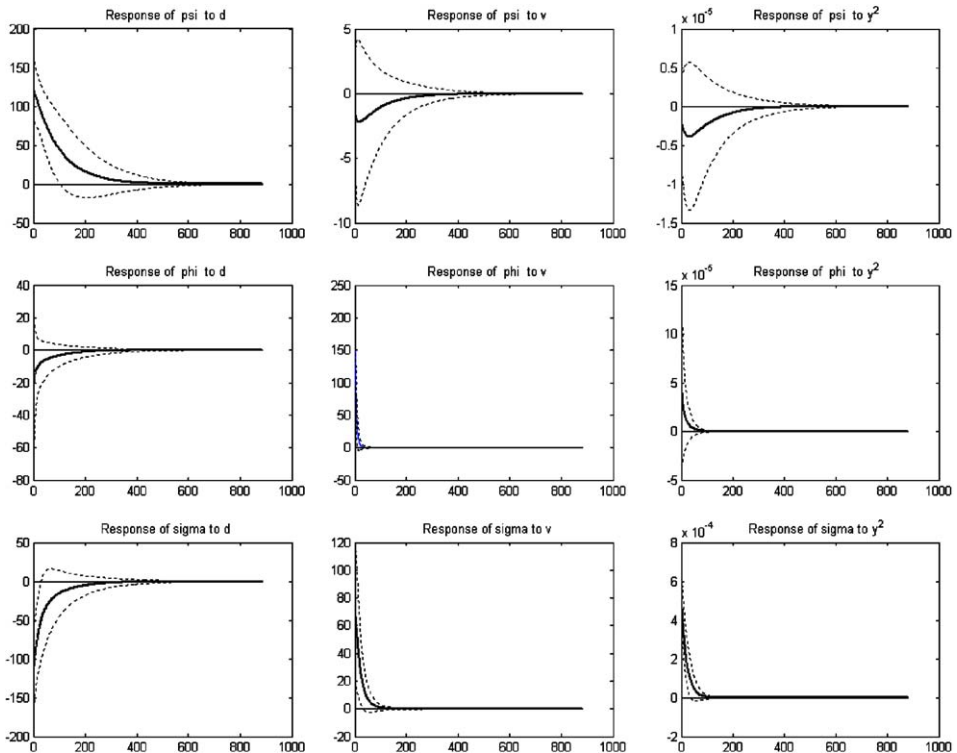


Fig. 2. Impulse–response function for JAX in calendar time (hours).

Table 7

Time (in hours) it takes to absorb shocks to the long term equilibrium variances

	AVT	COX	CP	DLP	GAP
Shock to duration	20	101	32	143	117
Shock to volume	21	108	53	156	105
Shock to volatility	15	23	13	76	21
	DTC	FTD	GBX	GSE	JAX
Shock to duration	164	264	2782	946	1006
Shock to volume	137	245	1611	1039	653
Shock to volatility	106	156	1398	880	506

frequently traded stocks converges much faster to its long-run equilibrium after an initial perturbation.¹¹

¹¹The threshold at which the shock producing the impulse–response is assumed to be absorbed is at $1e-7$ for shocks to ζ_0 and at 0.001 for shocks to ε_0 and η_0 . That is, Table 7 reports the time it takes for the impulse–response of the variance to fall below $1e-7$.

5. Conclusion

In this paper, we modeled the interaction between expected duration, expected volume and variance of returns. The methodology was applied to two groups of stocks, differing from each other according to their trade intensity. We found significant evidence that the frequency with which the stock is traded is important in determining its dynamic behavior. We found empirical evidence in favour of three working hypotheses:

- (1) All the stocks in the analysis exhibit a certain degree of trading clustering, the degree being significantly higher for frequently traded stocks. This prediction is in line with that of Foster and Viswanathan (1990), who claim that these clustering effects should be stronger for firms with better public reporting and more discretionary liquidity traders.
- (2) The key intuition of Easley and O'Hara (1992) that times of greater activity coincide with a higher fraction of informed traders present in the market is confirmed only for the frequently traded stocks. Infrequently traded stocks seem to be characterised by different information transmission mechanisms.
- (3) In line with the predictions of Holden and Subrahmanyam (1992), we find that more frequently traded stocks converge more rapidly (in calendar time) than the infrequently traded stocks to their long-run equilibrium, after an initial perturbation.

Regarding directions for future research, an interesting extension of the econometric framework constructed in this paper is to incorporate depth and spread, by modeling them in a similar way to volume and duration. By adding these two explanatory variables to the model, one could get a clearer picture about the trading environment of a specific market.

A second avenue for research, as discussed in footnote 6, would be to incorporate the trade sign of Hasbrouck (1991) into the approach proposed in this paper.

A third interesting possibility is a systematic study that explores the relationship between the characteristics of different markets and their dynamic properties. Many market microstructure models show that prices eventually converge to full information value, but provide very little insight into how long the adjustment process might take. We believe that the impulse–response framework suggested in this paper might prove a very valuable tool for this kind of analysis. For example, one could compare the estimates of the model across market structures (electronic versus specialist markets) or across different periods (crashes versus calm markets).

References

- Admati, A., Pfleiderer, P., 1988. A theory of intraday patterns: volume and price variability. *Review of Financial Studies* 1 (Spring), 3–40.
- Bagehot, W., 1971. The only game in town. *Financial Analysts Journal* 27, 12–14.

- Bauwens, L., Giot, P., 2000. The logarithmic ACD model: an application to the bid-ask quote process of three NYSE stocks. *Annales d'Economie et de Statistique* 60, 117–149.
- Bauwens, L., Giot, P., 2003. Asymmetric ACD models: introducing price information in ACD models. *Empirical Economics* 28, 709–731.
- Bollerslev, T., Wooldridge, J.M., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* 11 (2), 143–172.
- Copeland, T.E., Galai, D., 1983. Information effects and the bid-ask spread. *Journal of Finance* 38, 1457–1469.
- Diamond, D.W., Verrecchia, R.E., 1987. Constraints on short-selling and asset price adjustments to private information. *Journal of Financial Economics* 82, 33–53.
- Drost, F.C., Nijman, T.E., 1993. Temporal aggregation of GARCH processes. *Econometrica* 61, 907–927.
- Dufour, A., Engle, R.F., 2000. Time and price impact of a trade. *Journal of Finance* 55, 2467–2498.
- Easley, D., O'Hara, M., 1987. Price, trade size, and Information in securities markets. *Journal of Financial Economics* 19, 69–90.
- Easley, D., O'Hara, M., 1992. Time and the process of security price adjustment. *Journal of Finance* 47, 577–606.
- Ederington, L.H., Lee, J.H., 1993. How markets process information: news releases and volatility. *Journal of Finance* 48, 1161–1191.
- Engle, R.F., 2000. The econometrics of ultra high frequency data. *Econometrica* 68, 1–22.
- Engle, R.F., 2002. New frontiers for ARCH models. *Journal of Applied Econometrics* 17, 425–446.
- Engle, R.F., Russell, J., 1998. Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica* 66, 1127–1162.
- Engle, R.F., Hendry, D.F., Richard, J.-F., 1983. Exogeneity. *Econometrica* 51, 277–304.
- Fleming, M.J., Remolona, E.M., 1999. Price formation and liquidity in the U.S. Treasury market: the response to public information. *Journal of Finance* 54, 1901–1915.
- Foster, F.D., Viswanathan, S., 1990. A theory of the interday variations in volume, variance, and trading costs in securities markets. *Review of Financial Studies* 3 (4), 593–624.
- Foster, F.D., Viswanathan, S., 1993. Variations in trading volume, return volatility, and trading costs: evidence on recent price formation models. *Journal of Finance* 48 (1), 187–211.
- Glosten, L., Milgrom, P., 1985. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 13, 71–100.
- Grammig, J., Wellner, M., 2002. Modeling the interdependence of volatility and inter-transaction duration processes. *Journal of Econometrics* 106, 369–400.
- Hartmann, P., 1999. Trading volumes and transaction costs in the foreign exchange market. Evidence from daily dollar–yen spot data. *Journal of Banking and Finance* 23, 801–824.
- Hasbrouck, J., 1991. Measuring the information content of stock trades. *Journal of Finance* 46, 179–207.
- Holden, C.W., Subrahmanyam, A., 1992. Long-lived private information and imperfect competition. *Journal of Finance* 47, 247–270.
- Jones, C.M., Kaul, G., Lipson, M.L., 1994. Transactions, volume and volatility. *The Review of Financial Studies* 7, 631–651.
- Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica* 53, 1315–1336.
- Lee, C.M., Ready, M.J., 1991. Inferring trade direction from intraday data. *Journal of Finance* 46, 733–747.
- Manganelli, S., 2002. Duration, volume and volatility impact of trades. ECB Working Paper No. 125, European Central Bank, Frankfurt am Main.
- O'Hara, M., 1995. *Market Microstructure Theories*. Blackwell Publisher Inc, MA.
- White, H., 1980. A heteroscedasticity-consistent covariance matrix estimator and a direct test for heteroscedasticity. *Econometrica* 48, 351–372.
- Zhang, M.Y., Russell, J., Tsay, R.S., 2001. A nonlinear autoregressive conditional duration model with applications to financial transaction data. *Journal of Econometrics* 104, 179–207.