A Discrete-State Continuous-Time Model of Financial Transactions Prices and Times: The ACM-ACD Model[1].

Jeffrey R. Russell[**]
and
Robert F. Engle[***]

January 2004

Abstract

Financial transaction prices typically lie on a discrete grid of values and arrive at random times. This paper proposes an econometric model with this structure. The distribution of each price change is a multinomial, conditional on past information and the time interval between the transactions. The proposed autoregressive conditional multinomial model is not restricted to be markov or symmetric in response to shocks, however such restrictions can be imposed. The duration between trades is modeled as an ACD model following Engle and Russell (1998). Maximum likelihood estimation and testing procedures are developed. The model is estimated with 12 months of tick data on a moderately frequently traded NYSE stock, Airgas. The preferred model is estimated with three lags for the ACM and two lags for the ACD. Both price returns and squared returns influence future durations and present and past durations affect price movements. The model exhibits reversals in transaction prices in the short run due to bid-ask bounce and clustering of large moves of either sign in the longer run. Evidence of symmetry in the dynamics of prices is presented, but the response to durations is clearly non-symmetric. It is found that the volatility per second of trades is highest for short duration trades and that expected returns are lower for longer duration trades.

Keywords: Discrete valued time series, marked point process, high frequency data, ACD, transaction prices, bid-ask bounce, markov chain, multinomial.

---

[1] The authors are very grateful to the associate editor and referees who provided extensive suggestions. We would also like to thank David Brillinger, Xiaohong Chen, Clive Granger, Jim Hamilton, Alex Kane, Bruce Lehman, Peter McCullagh, Glenn Sueyoshi, George Tiao, and Hal White for valuable input. The first author acknowledges financial support from the Sloan Foundation, the University of California, San Diego Project in Econometric Analysis Fellowship and the University of Chicago Graduate School of Business. The second author acknowledges financial support from the National Science Foundation grant. SBR-9422575

[**] University of Chicago, Graduate School of Business
email: jeffrey.russell@gsb.uchicago.edu
[***] NYU Stern School of Business email: rengle@stern.nyu.edu

1. Introduction

The role of computers in modern business has generated a new type of economic data where every single transaction is recorded. Nowhere is this refinement in data collection more extensive than for financial data where transaction by transaction data sets contain detailed information about precisely when an asset is traded, as well as characteristics such as the transaction price and quantity. These new data sets provide us with an unprecedented microscopic view of the structure of financial markets that was previously impossible with time aggregated data.

Econometric modeling of transaction by transaction price dynamics is complicated by several features of the data. First, in a continuous auction market, such as the NYSE or the NASDAQ, transactions can occur at any point in time that the market is operational. As such, transactions do not occur at regularly spaced time intervals. The times between trades are random and are potentially informative about the underlying processes. Second, every financial market structure specifies a minimum unit of price measurement called ticks. That is, transaction prices must fall on a grid. When viewed over long time horizons, the variance of price changes greatly exceeds the effects of discreteness so that treating the data as continuous is unlikely to have a meaningful impact on the analysis. At the transaction by transaction frequency, however, price discreteness becomes a dominant feature of the data[2]. Often institutional rules in markets restrict the magnitude of transaction by transaction price movements. In the NYSE this is done by the specialist who is responsible for "price continuity". In other electronic markets, like the Taiwan stock exchange, price restrictions are directly imposed so consecutive prices can differ by no more than a fixed number of ticks. Hence the observed price changes often take just a handful of values. For the transactions data

[2] Institutional rules for a given market determine the granularity of the discreteness. The data that is considered in this paper is from the NYSE and price changes are restricted to \$1/16th's. The NYSE and NASDAQ began transitioning to a decimal system in August of 2000. The process was completed in January 2001.

studied in this paper, for example, we find that 99.3% of the price changes fall on one of just five different values.

An econometric model of such a process must be a continuous time process in that it should give at every instant of time the probability of observing a transaction at a particular discrete price, conditional on the past history of these processes. Most econometric models are not capable of this task and proceed by first converting to calendar time and second, ignoring the discreteness of prices. Any such procedure inevitably involves a loss of information and frequently leads to bias.

This paper proposes a new approach to modeling high-frequency transaction price dynamics that addresses both the spacing of the data and the discreteness. We propose treating the transactions data as a sequence of arrival times and characteristics associated with those arrival times. This is commonly referred to as a marked point process in the statistics literature. Following Engle (2000) we decompose the joint distribution into the product of the conditional distribution of price changes and the marginal distribution of the time interval. We propose using a variation of the ACD model of Engle and Russell (1998) for the marginal distribution of arrival times. A new model is then proposed for the discrete price movements that is flexible enough to capture the complex temporal dependence typically displayed by high-frequency transactions data. The model is termed the Autoregressive Conditional Multinomial (ACM) model.

While the proposed ACM model seems particularly well suited for analysis of financial transactions data, or movements in the midpoint of the discrete bid and ask quotes, it could be useful in many other applications involving time series of discrete random variables. For example, traders face a fixed number of possible order flow strategies involving market orders vs. limit orders. The model may be useful in the study of the dynamics of order flow without the need to impose a markovian structure. Alternatively, the ACM model may prove useful for credit risk ratings in financial markets or in marketing where the product brand purchased by consumers over time is of interest. The model can be applied to fixed interval data or, if the time series is viewed at irregular intervals, jointly modeled using the durations and the discrete random variable of interest. Hence, while the immediate application is to financial transactions data we believe the model could prove useful in a variety of other settings.

The model is applied to an NYSE traded stock. Given the joint distribution, the nature of the dependence between contemporaneous durations and transaction price changes can be examined. Two strands of theoretical market microstructure models provide predictions regarding the nature of dependence between transaction rates and price adjustments. First, Easley and O'Hara (1992) suggest that high trading rates may be associated with the presence of informed traders. In a rational expectations setting the specialist will make price adjustments more sensitive to order flow, thereby increasing volatility. Hence rapid trading should be associated with higher volatility. Second, Diamond and Verrecchia (1987) suggest that short sale constraints restrict the ability of privately informed agents possessing "bad news" to transact and capitalize on their better information. Since no such restraints exist in long positions high trading rates tend to be associated with good news and rising prices while the converse is true for slow trading rates.

The preferred model is estimated with three lags for the ACM and two lags for the ACD. Both price returns and squared returns influence future durations and present and past durations affect price movements. The model exhibits reversals in transaction prices in the short run due to bid-ask bounce and clustering of large moves of either sign in the longer run. Evidence of symmetry in the dynamics of prices is presented, but, consistent with the theory of short sale constraints and information in Diamond and Verrecchia (1987), we find that the response to durations is clearly non-symmetric with long durations predicting falling prices. Consistently with Easley and O'Hara (1992), we find that the volatility per second of trades is highest for short duration trades and that expected returns are lower for longer duration trades. Finally, both price changes and squared price changes are found to influence future durations.

This paper is organized as follows. Section 2 discusses the general modeling approach for irregularly spaced, discrete valued transactions data advocated in the paper. Section 3 introduces the ACM model for the discrete price changes. Some theoretical properties of the model are developed and estimation and model diagnostics are also presented in this section. Parameter restrictions for the ACM model derived from a symmetry condition are also considered. Section 4 presents model estimates and analysis for an NYSE traded stock and section 5 concludes.

2. An approach to joint modeling of arrival times and price changes.

The modeling strategy we adopt is conveniently motivated by first considering a short, representative sample of trades for the NYSE listed stock Airgas (ARG). In Figure 1a calendar time (or clock time) is on the horizontal axis and the price is on the vertical axis. Each diamond represents a point in time that a transaction occurred and the associated price. Two features of the data immediately become apparent. First, the intertrade durations vary significantly within the sample. This can be observed in the plot by noting the occasional long horizontal stretch between observations. Second, the transaction price changes take just 5 different values in this sample. Hence price discreteness is a dominant feature of the data.

In this paper we propose decomposing the time series plotted in Figure 1a into a bivariate system. This is shown graphically in Figures 1b and 1c. Figure 1b presents the discrete price changes from transaction to transaction. Figure 1c presents the counting function that denotes the number of transactions that have occurred by time t. These series comprise the bivariate system. More formally, let t_i denote the arrival time of the i^{th} transaction. Let $N(t)$ denote the counting function which describes the number of events that have occurred by time t. A sequence of arrival times that are strictly increasing is called a simple point process. At each transaction time t_i we denote the associated realization of the trade-to-trade change in the asset price by y_i. The bivariate process of arrival times and marks is called a marked point process. Since transaction price changes are typically discrete in nature we consider y to be a multinomial random variable.

A simple point process can be completely described by the sequence of arrival times t_i or the durations $\tau_i = t_i - t_{i-1}$. Our goal is to develop a model for the joint distribution of the discrete price changes and durations conditional on the bivariate filtration of arrival times and price changes. We denote this conditional bivariate density by $f\left(y_i, \tau_i \middle| y^{(i-1)}, \tau^{(i-1)}\right)$ where $y^{(i-1)} = \left(y_{i-1}, y_{i-2}, ..., y_1\right)$ and $\tau^{(i-1)} = \left(\tau_{i-1}, \tau_{i-2}, ..., \tau_1\right)$. Our discussion entails a substantial use of super/subscripts. The reader is referred to Appendix A for a summary of the superscript/subscript notation.

As discussed in Engle (2000), without any loss in generality we can decompose the joint conditional density of y_i and τ_i into the product of the conditional density of the mark and the marginal density of the arrival times, both conditioned on the past filtration of the joint information set. That is, if $f\left(y_i, \tau_i\right)$ denotes the joint density of y_i and τ_i then

$$(1) \qquad f\left(y_i, \tau_i \middle| y^{(i-1)}, \tau^{(i-1)}\right) = g\left(y_i \middle| y^{(i-1)}, \tau^{(i)}\right) q\left(\tau_i \middle| y^{(i-1)}, \tau^{(i-1)}\right)$$

Where $g(\cdot)$ denotes the probability density function associated with the price changes y_i conditional on $\tau^{(i)}$ and $y^{(i-1)}$. $q(\cdot)$ denotes the density function of the i^{th} duration conditional on $\tau^{(i-1)}$, $y^{(i-1)}$. Clearly the joint density in (1) allows us to study the relationship between price changes and contemporaneous durations as well as analyze their joint dynamics.

Models that explain the probability of each possible outcome of a discrete random variable at time τ, are often called competing risks models[3]. The joint density in equation (1) clearly yields such a set of probabilities. Competing risks models are generally specified by the instantaneous probability of exit to state y given a duration τ. The hazard function denoting the instantaneous probability that the i^{th} trade exits to state y given the duration τ since the last event:

$$(2) \qquad \theta_i\left(y, \tau\right) = \lim_{\Delta t \to 0} \frac{\Pr\left(Y_i = y, \tau \le \tau_i < \tau + \Delta t \mid \tau_i \ge \tau, y^{(i-1)}, \tau^{(i-1)}\right)}{\Delta t}$$

For small values of Δt, $\theta_i\left(y, \tau\right)\Delta t$ is approximately equal to the probability of exit to state y over the time period $\left[t, t + \Delta t\right]$ given no transaction has occurred by duration τ. The hazard functions are easily obtained from (1) and are given by:

---

[3] See Kalbfleish and Prentice (1980) or, more recently, Lancaster (1990) for references on duration models.

(3) $\theta_i\left(y,\tau\right) = \kappa\left(\tau \middle| y^{(i-1)}\tau^{(i-1)}\right) g\left(y \middle| y^{(i-1)},\tau^{(i)}\right)$

where $\kappa\left(\tau \middle| y^{(i-1)},\tau^{(i-1)}\right) = \dfrac{q\left(\tau \middle| y^{(i-1)},\tau^{(i-1)}\right)}{1-\int_0^\tau q\left(s \middle| y^{(i-1)},\tau^{(i-1)}\right) ds}$ is the hazard function associated with the distribution of the waiting times between transactions.

Clearly, conditional moments of the price changes can be directly obtained from (1). Moments of price changes and trading rates over more than one transaction can sometimes be expressed analytically, but calendar time measures will generally require simulations.

Instantaneous moments of the price process are also easily obtained. We define the price at time t as the price associated with the most recent transaction and denote it by $p\left(t\right) = p\left(t_{N(t)}\right)$. The instantaneous mean and variance of the price change at time t can be conveniently expressed using the counting function $N(t)$ as follows:

$$(4) \qquad \mu\left(t\right) = \lim_{\Delta t\to 0} \frac{E\left(\left(p\left(t+\Delta t\right)-p\left(t\right)\right) \middle| N(t), y^{(N(t))}, t^{(N(t))}\right)}{\Delta t} = \sum_{\text{all y}} y\theta_{N(t)}\left(y, t-t_{N(t)}\right)$$

and

$$(5) \qquad \sigma^2\left(t\right) = \lim_{\Delta t\to 0} \frac{E\left\{\left[\left(p(t+\Delta t)-p(t)\right)-\mu\left(t\right)\right]^2 \middle| N(t), y^{(N(t))}, t^{(N(t))}\right\}}{\Delta t}$$

$$= \sum_{\text{all y}} y^2\theta_{N(t)}\left(y, t-t_{N(t)}\right) - \left(\mu\left(t\right)\right)^2$$

The instantaneous moments conveniently characterize the evolution of the price process. All else equal, the magnitude of the instantaneous mean and variance will be larger when the probability of a transaction is higher.

Engle and Russell (1998) propose the Autoregressive Conditional Duration (ACD) model for *q(·)* and find the model performs well for transactions data. A wide range of empirical studies have now compared various specifications of this general ACD form. See for example Bauwens, Giot, Grammig and Veredas (2003). Given an ACD formulation for $q(\cdot)$ the only remaining task is to specify a model for the prices $g(\cdot)$. We now focus our attention on modeling the conditional distribution of price changes $g(\cdot)$.

3. The ACM model.

Unlike their low frequency counterparts, high frequency price changes tend to exhibit strong and often complex temporal dependence. Any good model for the price changes must therefore be flexible and capable of generating strong dependence spanning many transactions. This section develops a time series model for discrete random variables consistent with this goal. The model is termed the Autoregressive Conditional Multinomial (ACM) model. We also establish some theoretical properties for an ACM model.

3.1 Model Specification.

We restrict our attention to the class of observation driven models in the sense of Cox (1981)[4]. Let k denote the number of states the multinomial random variable y_i can take. Let $\tilde{x}_i$ be a kx1 vector indicating the discrete price change y_i. $\tilde{x}_i$ takes the j^{th} column of the kxk identity matrix if the j^{th} state occurred. Let $\tilde{\pi}_i$ denote the kx1 vector of conditional (on information available at time t_{i-1}) probabilities associated with the states. That is, the j^{th} element of $\tilde{\pi}_i$ corresponds to the probability that the j^{th} element of $\tilde{x}_i$ takes the value 1. Clearly the conditional distribution of y_i is completely characterized by $\tilde{\pi}_i$. A natural starting point is a Markov chain. A first order Markov chain can be expressed as:

[4] Many models have been suggested in the context of parameter driven models and associated hidden markov models. While this literature is rich the models are often difficult to estimate and forecast. See MacDonald and Zucchini (1997) for a recent survey. Relatively little work has been done on discrete valued observation driven models. Jacobs and Lewis (1983) pursued a class of models for discrete valued time series data called DARMA models. These models often have unrealistic properties such non-negative autocorrelation restrictions. Furthermore, these models appear better suited for marginally Poisson, or Binomial data. The model proposed here is applicable to multinomial data.

(6) $$\tilde{\pi}_i = P\tilde{x}_{i-1}$$

where P is a kxk transition matrix that must satisfy:

a) all elements are nonnegative
b) all columns must sum to unity

In the more general setting P may be a conditional transition matrix and will vary with information available at period (i-1). In this context, we can include information on longer lags of $\tilde{x}$, perhaps past values of $\tilde{\pi}$ and the past arrival times of the transactions. An early discussion of time varying transition probabilities can be found in MacRae (1977), although their emphasis was on estimation when only aggregate data is available.

The restrictions on P are directly satisfied by simple estimators when the transition matrix is constant, but become quite difficult to impose in simple extensions. Here we propose using an inverse logistic transformation that imposes such conditions directly for any set of covariates[5]. Let π_{im} and x_{ij} denote the m^{th} and j^{th} elements of π_i and x_i. The log odds of the m^{th} ($m<k$) state relative to the k^{th} state is given by:

(7) $$\begin{aligned}\log(\tilde{\pi}_{im} / \tilde{\pi}_{ik}) &= \log\left(\sum_{j=1}^{k} P_{mj}\tilde{x}_{(i-1)j}\right) - \log\left(\sum_{j=1}^{k} P_{kj}\tilde{x}_{(i-1)j}\right) \\ &= \sum_{j=1}^{k} \log\left(P_{mj} / P_{kj}\right)\tilde{x}_{(i-1)j} \\ &= \sum_{j=1}^{k-1} P^{*}_{mj} x_{(i-1)j} + c_m\end{aligned}$$

[5] Hausman Lo and Mackinlay (1992) have proposed modeling discrete price changes using a Probit model with time varying mean and variance. Their approach, however, allows for very limited dependence due to its markov structure and is far less flexible regarding the impact of new information on the transition probabilities. Since the first draft of this paper numberous interesting alternative approaches have been proposed in the literature. Rydberg and Shephard (2000, 2003) proposed an alternative model for discrete price movements, Bauwens and Giot (2003) proposed an alternative competing risk model, and Prigent et al (2001) have applied the two-state ACM model suggested here in an option pricing setting. Models that assume continuous distributions for returns were considered in Ghysels and Jasiak (1998) and in Grammig and Wellner (2002).

where x is now a (k-1) dimensional vector and P^* is a (k-1)x(k-1) matrix constructed from the log odds and c_m is a scalar constant[6]. Rewriting the k-1 probabilities as the vector π, and defining the vector of logs of the probability ratios as $h(\pi_i) = \log(\pi_i / (1 - \iota' \pi_i))$, where ι is a conforming vector of ones, we get:

$$(8) \qquad h(\pi_i) = P^* x_i + c$$

where P^* is an unrestricted (k-1)x(k-1) matrix and c is a (k-1) vector with m^{th} element given by c_m. For any values of P^* and c the conditional probabilities are easily recovered from the logistic transformation:

$$(9) \qquad \begin{aligned} \frac{\pi_i}{1 - \iota' \pi_i} &= \exp\left[P^* x_{i-1} + c \right] \\ \pi_i &= \frac{\exp[P^* x_{i-1} + c]}{1 + \iota' \exp\left[P^* x_{i-1} + c \right]} \end{aligned}$$

where, again, ι denotes a conforming vector of ones and $\exp(P^*)$ is interpreted as a matrix with m,n element $\exp(P^*_{mn})$. Now, all probabilities will be positive including the probability of the k^{th} state are obtained from condition b) and will sum to unity. An expression for the transition probabilities is then obtained:

$$(10) \qquad P_{mn} = \frac{\exp[P^*_{mn} + c_m]}{1 + \sum_{j=1}^{k-1} \exp[P^*_{jn} + c_j]}$$

and again these are all positive and have columns that sum to unity.

We now consider generalizing (8) to allow for a more elaborate dynamic structure with dependence on a richer information set than just the most recent price movement. In

[6] Tildes on vectors denote a k dimensional vector and, unless otherwise specified, vectors without a tilde denote (k-1) dimensional vectors.

doing so, it is clear that we are generalizing the transition matrix in (2) from a time invariant transition matrix to one that varies over time.

Definition 1: An Autoregressive Conditional Multinomial model of order (*p*,*q*) is given by:

$$h(\pi_i) = \sum_{j=1}^{p} A_j \left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} B_j h\left(\pi_{i-j}\right) + \chi z_i \qquad (11)$$

where $h(\cdot)$ is the inverse logistic function. A_j and B_j denote the j^{th} (k-1)x(k-1) parameter matricies. z_i is an r+1 dimensional vector with 1 in the first element forming a constant and r other explanatory variables. χ denotes a (k-1)x(r+1) conforming matrix of parameters. These explanatory variables may contain predetermined variables such as characteristics of past trades including volume or spreads or, as of interest in our application, the vector z may include information about the timing of trades. The terms $(x_i - \pi_i)$ form a martingale difference sequence characterizing the new information associated with the i^{th} transaction[7].

Clearly, (11) can be interpreted as specifying dynamics for the conditional log odds for all states with respect to a base state and therefore specifies the dynamics of the conditional log odds for all pairs of states. It follows that the specification in (11) completely describes the transition probabilities and hence the dynamics of the multinomial random variable y_i. The linear structure of (11) implies that the choice of the base state is arbitrary. This is easily verified since parameters for the choice of any base state can be expressed as an exact function of the parameters of any other choice of base state.

From (11) it is immediately apparent how the history impacts the transition probabilities. The structure of this equation is recursive. At the time of the i-1 transaction, knowing all past x and π gives from (9) a calculated value of the next π. Consequently, subject to some starting values, the full sequence of transition probabilities π can be constructed from observations on x. This allows evaluation of the likelihood

[7] Shephard (1995) considers generalized linear autoregressive time series models in the same spirit as (7). This paper can be referenced at http://www.nuff.ox.ac.uk/economics/papers/index1995&4.htm

function and its numerical derivatives. It can now be seen that the first $(k$-1$)$ conditional probabilities are easily recovered from

(12) $$\pi_i = \frac{\exp[\sum_{j=1}^{p} A_j \left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} B_j h\left(\pi_{i-j}\right) + \chi z_i]}{1 + \iota' \exp[\sum_{j=1}^{p} A_j \left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} B_j h\left(\pi_{i-j}\right) + \chi z_i]}$$

The k^{th} probability is determined by condition b). Hence the transition probabilities are given by π_i and the conditional covariance matrix of x can be defined as

(13) $$V_i \equiv V\left(x_i \mid x^{(i-1)}, z_i\right) = diag\{\pi_i\} - \pi_i \pi_i'$$

We now turn to some theoretical properties of the model. For illustrative purposes consider the ACM(1,1) model when p=q=1 and r=0 so that z_i is simply a constant denoting the intercept. When the eigenvalues of B are distinct and lie inside the unit circle, we can rewrite the ACM(1,1) model as:

(14) $$h\left(\pi_i\right) = \sum_{j=1}^{\infty} P \Lambda^{j-1} P^{-1} A\left(x_{i-j} - \pi_{i-j}\right) + \left(I - B\right)^{-1} \chi$$

where $B= P\Lambda P^{-1}$ and Λ is the diagonal matrix with the eigenvalues of B along the diagonal. Since the $\left(x_i - \pi_i\right)$ form a martingale difference sequence the dynamics of the ACM model are easily understood from (14). The impact of past information is determined by A while the decay of past information is determined by the eigenvalues of B. Generalizing (14) to an ACM(p,q) we can construct bounds for the transition probabilities determined by the parameters of the ACM model. These results are summarized in the following theorem.

Theorem 1: Consider the ACM(p,q) model given by (11) with r=0 (denoting a constant term only) and let the A_i (i=1,...,p) and B_j (j=1,...,q) be of full rank. Let I_{k-1} denote the k-1 dimensional identity matrix. If all the values of z satisfying $\left|I_{k-1} - B_1 z - B_2 z^2 - ... - B_q z^q\right| = 0$ lie outside the unit circle, then the elements of π_i are strictly positive.

Proof: See appendix B.

Corollary 1: Under the conditions of theorem 1, y_i is irreducible meaning that regardless of the initial condition, every state will be visited infinitely often as $i \to \infty$. Furthermore, y is aperiodic in the sense that the minimum recurrence time is one period.

Proof: This follows trivially from Theorem 1 and the fact that k is finite.

Corrollary 1 insures that in the long run all states will be visited infinitely often and that the transition matrix will always be fully saturated, that is, any state is attainable regardless of the sequence of preceding price moves. It would seem that any good model for the transition probabilities of transaction prices should have these properties[8].

3.2 Estimation and Diagnostics for the ACM model.

Given initial conditions, the entire path of π_i can be constructed. Hence the likelihood can be constructed as the product of the conditional densities. Letting π_{ij} denote the j^{th} element of π_i the log likelihood is then expressed as:

(15) $$L = \sum_{i=1}^{N} \sum_{j=1}^{K} \left(\tilde{x}_{ij} \log\left(\tilde{\pi}_{ij}\right)\right) = \sum_{i=1}^{N} \tilde{x}_i' \log\left(\tilde{\pi}_i\right)$$

For a general ACM(p,q) model the derivatives of the log likelihood take a recursive form analogous to those of GARCH models. We therefore propose estimating the model by

[8] We note that Theorem 1 does not necessarily provide sufficient conditions for stationarity of the transaction price process.

maximum likelihood using a numerical optimization algorithm such as Berndt, Hall, Hall and Hausman (1974) (BHHH). Under the usual regularity conditions, we will obtain consistent asymptotically normal parameter estimates.

Model diagnostic tests are suggested by considering the sequence of errors

$$(16) \qquad v_i^* = x_i - \pi_i .$$

This sequence should form a heteroskedastic martingale difference sequence where the conditional variance covariance matrix is given by V_i in (13). Standardized errors are constructed by pre-multiplying v_i^* by the Cholesky factorization of the conditional variance covariance matrix. The standardized errors are given by:

$$(17) \qquad v_i = U_i v_i^* \text{ where } U_i V_i U_i = I$$

Now, v_i should be uncorrelated with the past and have a variance covariance matrix equal to the (k-1) identity matrix. Moreover, v_i should be uncorrelated with the filtration of price moves and any information in z_i. Given parameter estimates we can construct the series of standardized residuals $\hat{v}_i$, the sample counterpart to (17). Tests can then be performed to check if $\hat{v}_i$ is uncorrelated. The s^{th} sample cross correlations associated with the standardized residuals are calculated by

$$(18) \qquad P_s = \frac{1}{N-(s+1)} \sum_{i=s+1}^{N} \hat{v}_i \hat{v}'_{i-s}$$

A formal test of the null hypothesis that the elements of the standardized vector are white noise can be done with a multivariate version of the Portmanteau statistic. Li and McLeod (1981) propose a test based on the statistic

$$(19) \qquad Q = N \sum_{s=1}^{M} Trace\left(P_s P'_s\right)$$

This test statistic will be distributed as a chi-squared with $(k\text{-}1)^2 * M$ degrees of freedom.

3.3 Symmetry in Price Dynamics

In this section we propose some parameter restrictions for the general ACM(p,q) model. Harris (1990) provides a detailed discussion of the effects of price discreteness on estimates of autocorrelations (and the variance) calculated from the observed return series. The cornerstone of Harris' work is the idea that the observed transaction price is the "true" price of the asset plus an upward (downward) departure for buyer (seller) initiated trades rounded to the nearest tick. Harris assumes that the arrivals of buyers and sellers can be described by an iid Bernoulli with constant probability .5. If order flow is correlated, as suggested in Hasbrouck (1991), analysis of the effects of discreteness on the price dynamics becomes much more complicated. In the presence of (unobserved) time varying risk the magnitude of departures of the observed price from the efficient price may also be time varying, further complicating any analysis of the price dynamics. In this more realistic setting we have little theory to guide us in determining the dynamics of discrete bid and ask prices.

Nevertheless, there is a particular type of symmetry that we might expect in the dynamics of the price movements. This hypothesis is most easily understood by examining a simple special case. Consider the simple 2-state time invariant markov model given in (6). Let state 1 denote a downward price movement and state 2 an upward price movement and let p_{ij} denote the i,j element of P. Our symmetry hypothesis restricts $p_{12}=p_{21}$ and $p_{11}=p_{22}$. These restrictions impose that the probability of a price continuation is the same regardless of whether the price is moving up or down. Similarly the probability of a price reversal is the same regardless of whether the price moved up or down. Following state 1 the conditional distribution is $\left[p_{11}\ p_{21}\right]'$ and, imposing the restrictions above, the conditional distribution following a downward price movement is given by $\left[p_{21}\ p_{11}\right]'$. The restrictions imply that the conditional distribution following an upward (downward) price move is the mirror image of the conditional distribution following a downward (upward) price move, and vice-versa.

We now generalize the symmetry hypothesis beyond the simple first order markov model and to more than just two states. Define the matrix Q to be a rotated identity matrix:

$$\text{(20)} \qquad \mathbf{Q} = \begin{bmatrix} 0 & & 1 \\ & \iddots & \\ 1 & & 0 \end{bmatrix}$$

The elements of a conforming vector are reversed when pre-multiplied by Q. Arrange x_i in the natural ordering with the first element corresponding to the extreme down price move and the last element corresponding to the extreme upward price move. The zero price move is taken to be the base state given by the zero vector. We define the mirror image history by Qx_{i-1}, $Qx_{i-2}, \ldots$. That is, upward price movements become downward price movements of equal magnitude in the mirror image. For a general k-state model with $(k\text{-}1)/2$ upward price movement states, and corresponding $(k\text{-}1)/2$ downward price movement states we make the following definitions:

Definition 2: An nxn matrix W is response symmetric if for the nxn matrix Q defined in (16) $QW = WQ$. That is, Q and W commute.

Definition 3: A vector w is symmetric if $Qw = w$.

Definition 4: For an ACM(p,q) model we say the transaction price process is dynamic symmetric for prices if all A_i and B_i are all response symmetric matricies.

Definition 5: For an ACM(p,q) model we say the transaction price process is dynamic symmetric for the j^{th} element of z_i if the j^{th} column of χ is a symmetric vector.

In this case, the marginal impact of a dynamic symmetric element of z on the log odds is the same for price moves of equal but opposite direction.

Proposition 1: If an ACM(p,q) model is dynamic symmetric for prices and elements of Z then $Q\pi_i\left(x_{i-1}, x_{i-2}, \ldots z_i\right) = \pi_i\left(Qx_{i-1}, .Qx_{i-2}, \ldots z_i\right)$.

Proof is in Appendix B.

When x_i arranged in its natural ordering then the conditions of Proposition 1 imply that the mirror image history of price changes will produce the mirror image transition probabilities. When the ACM model is dynamic symmetric for prices but not for all elements of z there is a remaining symmetry in the marginal impacts the past price changes on the log odds of the transition probabilities.

Proposition 2: Let H_s denote a matrix with m,n element given by $\dfrac{dh_{im}}{dx_{i-s,n}}$ where m and n denote the m^{th} and n^{th} element of h and x respectively. If an ACM(p,q) model is dynamic symmetric for prices then H_s is a response symmetric matrix for all s>0.

Proof is in Appendix B.

The implications of proposition 2 are most easily understood by returning to the simple 2 state model described above. Proposition 2 says that the marginal impact of a down tick on the log odds of a subsequent uptick is identical to the marginal impact of an uptick on the log odds of a subsequent downtick. Similarly the marginal impact of an uptick on a subsequent uptick is identical to the marginal impact of a downtick on a subsequent downtick. If an ACM(p,q) is dynamic symmetric for prices and the constant is a symmetric vector, then the number of parameters is reduced from $\left(k-1\right)\left(1+\left(k-1\right)\left(p+q\right)\right)$ to $\dfrac{\left(k-1\right)}{2}\left(1+k\left(p+q\right)\right)$ or almost by half. Clearly this restriction can be tested in practice.

A final model restriction that we consider is to set off-diagonal elements of all B_j to zero. Under this assumption the partial derivative of the log odds with respect to a unit shock will decay at a geometric rate determined by the diagonal elements of the B_j. Thus the impact of new information is generously specified while the long run decay is more parsimoniously formulated. We refer to this restriction as the diagonal specification.

4. Data and Estimation

In this section we provide a description of the financial transaction data and consider estimation and diagnostic tests of an ACM-ACD model. Tests for the symmetry hypothesis discussed in section 3.3 are also presented. Upon finding a good representation for the data, the nature of dependence between transaction price changes and durations is analyzed. Our findings are related to predictions obtained from existing market microstructure theory. We begin with a more detailed discussion of the data.

4.1 Data.

The number of transactions per day varies greatly from stock to stock. For example, IBM may experience 10,000 transactions in a single day (or about a trade every 2 or 3 seconds) while other stocks trade very infrequently often going an entire day without any transactions occurring. We try to strike a balance in this application by selecting a stock that trades about once every 3 minutes on average. This trading frequency provides a large number of transactions per day, but remains tractable enough to analyze one complete year of data. The stock analyzed is Airgas (ticker symbol ARG). This is the first stock alphabetically in decile 8 of the stocks examined by Engle and Patton (2003). Their selection was based on trade frequency during the previous year, with decile 10 as the most frequently traded stocks. The data used in this paper were abstracted from the TAQ (Trades and Quotes) data set distributed by the NYSE.

Over the one year period January 1999 to December 1999 there were 21,837 transactions. The average transaction price for the sample is \$10.44. The minimum price change for ARG during this period is 1/16th of a dollar. Following Engle and Russell (1998) we omit the first half hour of trades since some of these will contain trades recorded during the opening batch auction. We also delete overnight price changes leaving a sample of 18,573 transactions of which 61.72% of the transaction prices are unchanged from their previous value. The distribution of transaction price changes is roughly symmetric with 16.20% and 16.36% down one tick and up one tick respectively.

Down and up two ticks occurred with 2.60% and 2.39% frequency respectively. Downward moves greater than two ticks occurred with frequency 0.30% and upward moves greater than two ticks occurred with frequency 0.41%. A histogram of the transaction price changes is presented in Figure 2 and the raw frequencies are given in appendix C.

Given the sparseness of the data beyond two tick moves we use a five-state model. The two extreme states therefore include upward price movements of 2 or more ticks and downward price movements of 2 or more ticks. We use the natural ordering for the state vector given by:

$$x_i = \begin{cases} [1,0,0,0]' & \text{if } \Delta p_i \leq \text{-2 ticks} \\ [0,1,0,0]' & \text{if } \Delta p_i = -1 \text{ tick} \\ [0,0,0,0]' & \text{if } \Delta p_i = 0 \\ [0,0,1,0]' & \text{if } \Delta p_i = +1 \text{ tick} \\ [0,0,0,1]' & \text{if } \Delta p_i \geq +2 \text{ ticks} \end{cases}$$

The state vector provides an interesting perspective from which to view the dynamics of the transaction price changes. Using the multivariate summary methodology initially proposed by Tiao and Box (1981) the intertemporal cross correlations of the state vector are presented in matrix form with the correlations replaced by the symbols "+", "-", and "·". If the price changes were i.i.d., an asymptotic 95% confidence interval would be given by $1.96 * N^{-1/2}$. A dot indicates that a correlation does not exceed this 5% significance level. Plus and minus signs indicate positive and negative exceedences respectively.

Denoting the sample mean of x_i by the *k*-dimensional vector $\bar{x}$, the s^{th} sample cross correlation matrix is calculated by

$$\text{(21)} \qquad \begin{aligned} &\mathrm{P}_s = R_0^{-1} R_s \\ &\text{where } R_s = \frac{1}{N-(s+1)} \sum_{i=s+1}^{N} (x_i - \bar{x})(x_{i-s} - \bar{x})' \end{aligned}$$

The Tiao Box plot for the state vector temporal correlations up through lag 15 is given in Figure 3. The *i,j* element of the s^{th} matrix gives the correlation of state *i* with state *j* lagged *s* periods. The upper right and lower left quadrants correspond to price reversals. The upper left and lower right quadrants correspond to price continuations. For s=1 the cross correlations are generally positive in the upper right and lower left quadrants and negative in the upper left and lower right quadrants indicating that price reversals are more likely to occur than price continuations immediately after a price move. The transaction price "bounces" back and forth between buy and sell prices generating negative autocorrelation in transaction price changes and the observed positive signs associated with the price reversals. This is often referred to as bid-ask bounce.

Beyond the first lag, the significant correlations are generally positive. Significant elements tend to occur most often in the corners of the matrix and sometimes in the center. This pattern implies that large price changes (of 2 ticks or more) tend to follow large price changes of either direction and small price changes tend to follow small price changes of either direction. This pattern is an expression of volatility clustering in the discrete price moves.

Finally, we notice a particular symmetry in the correlations. For many of the correlations, the signs of the correlation reflected through the origin are the same. If the symmetry conditions discussed in section 4 are satisfied, this is the exact pattern the correlations should display. The symmetry condition will be formally tested later in this section.

The time intervals between trades are known to contain a periodic U shaped pattern throughout the trading day. Durations tend to be shortest in the morning just after the open and in the afternoon just prior to the close. Intraday volatility also exhibits a similar periodic pattern although this pattern is typically examined using price data observed over fixed time intervals[9]. We examine the state vector x_i to check for any evidence of these diurnal effects in the distribution of transaction by transaction price movements. In doing so patterns in the variance or any other moments of the price changes may be detected if present. We treat each of the four elements of x_i as a

---

[9] See, for example Engle and Russell (1998) for diurnal patterns in durations and McInish and Wood (1992) for an early analysis of volatility periodicity.

univariate time series and fit by least squares a linear spline in the time of day. Nodes are placed at each hour of the trading day and the spline is restricted to be continuous. The result for the j^{th} series is an estimate of the probability that state j occurs at any point during the day. This parallels the two-step procedure used to estimate deterministic patterns in Engle and Russell (1995). If there are no diurnal effects, the coefficients should be zero with only the intercept non-zero. An F-statistic is provided to assess the null hypothesis that there are no diurnal effects. The estimates for the time of day effects for each price change state as well as for the durations are given in Table 1. For each regression, d_j denotes the j^{th} spline coefficient estimate.

From Table 1 we see that the durations tend to be shortest near the open and close of the market and longest in the middle of the day - this is the inverted U-shape typically observed. The p-value for the null hypothesis of no time of day effects is near zero and the hypothesis is easily rejected for the durations. Alternatively, the price changes do not exhibit any indication of periodicity. The coefficients appear random with no real pattern and all are insignificant. The p-value for the null of no time of day effects is large for all of the price states. The smallest p-value is for state 1 (down 2 or more ticks) which has a p-value of 6.43%.

It is interesting not to find evidence of a deterministic pattern in the transaction by transaction price dynamics and, in particular, no pattern in the magnitude of the price changes since periodicity in volatility patterns is well documented for intraday prices measured in fixed time intervals. This suggests that the time of day patterns discovered for volatility using within-day fixed-interval analysis are driven by time of day patterns in the transaction rates rather than the magnitude of transaction by transaction price changes. This result is similar in spirit to findings reported in Ane and Geman (1999) and Jones Kaul and Lipson (1994). These studies, however, focus on the role played by the random number of transactions in directing the stochastic component of volatility.

Since the durations exhibited strong intraday deterministic patterns we follow the two step procedure discussed in Engle and Russell (1995) by first partialing out the deterministic pattern by taking the durations and dividing by their expectation based on time of day alone. The expectation is obtained from the splines in Table 1. The resulting series will have an unconditional mean near 1 and should be free of any deterministic patterns. We will simply refer to this series in what follows as the durations. We now turn to the specification and estimation of the ACM-ACD model.

4.2 Specification and Estimation

This section provides a specific parameterization for the ACM and ACD models that will be estimated using the ARG data.

We consider the following specification for the ACM model:

$$h(\pi_i) = c + \sum_{j=1}^{p} A_j \left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} B_j h\left(\pi_{i-j}\right) + \sum_{j=1}^{r} \chi_j \ln\left(\tau_{i-j+1}\right) \tag{22}$$

where χ_j is a (k-1) parameter vector. Clearly we have taken $z_i = \left[1, \ln\left(t_i\right), \ln\left(t_{i-1}\right), \ldots, \ln\left(t_{i-r+1}\right)\right]'$ so that the ACM model now depends on the log of the contemporaneous duration as well as the first (r-1) lags of the log duration. Since the log of the probability appears on the left hand side it seems natural to take logs of the durations that appear on the right hand side.

The dynamics for the durations are assumed follow an ACD model. An ACD model is characterized by 1) $\psi_i = E\left(\tau_i \mid I_{i-1}\right)$ where I_i is an information set available at time t_i and 2) $\frac{\tau_i}{\psi_i} = \varepsilon_i$ is iid. In our analysis we assume an exponential distribution for ε and the following form for ψ:

$$\ln\left(\psi_i\right) = \omega + \sum_{j=1}^{u} \alpha_j \varepsilon_{i-j} + \sum_{j=1}^{v} \beta_j \ln\left(\psi_{i-j}\right) + \sum_{j=1}^{w} \left(\rho_j y_{i-j} + \zeta_j y_{i-j}^2\right) \tag{23}$$

This specification differs from the original application of the ACD model in that the log of the expectation appears on the left-hand side. Additionally, it is the "innovation" ε_i and past values of the logged expectation that appears on the right hand side. Those familiar with the exponential GARCH model of Nelson (1991) will recognize the connection and we therefore refer to this model as the *Nelson-form ACD* model[10]. The Nelson-form ACD model is useful because it automatically insures that the conditional expectation of the duration is non-negative even in the presence of additional explanatory variables such as the past price changes. The ACD dynamics also depend on the first w lags of both the past price change and its square. Clearly durations can now depend on both the direction and the magnitude of past price changes. The conditional density function associated with the i^{th} duration is then given by:

$$(24) \qquad q\left(\tau_i \left| y^{(i-1)}, \tau^{(i-1)}\right.\right) = \frac{1}{\psi_i} \exp\left(-\frac{\tau_i}{\psi_i}\right)$$

Estimation of the ACD and ACM parameters can be performed by separately maximizing the two log-likelihoods or by joint estimation of (1) although there may be a loss of efficiency if estimation is performed separately. Clearly, maximizing the log of the joint likelihood in (1) is obtained by maximizing the sum of the ACM log likelihood given in (15) and $\sum_{i=1}^{N} \log\left(q\left(\tau_i \left| y^{(i-1)}, \tau^{(i-1)}\right.\right)\right)$ where $q\left(\tau_i \left| y^{(i-1)}, \tau^{(i-1)}\right.\right)$ is given by (23) and (24). In our work we perform joint estimation using the BHHH algorithm.

The dynamic structure of price changes associated with the closing transaction one evening and the opening transaction the next morning is unlikely to have the same dynamic structure as two consecutive trades within the same day. We therefore reinitialize lagged variables at the beginning of each day to their unconditional mean. The martingale terms in the ACM model is set to zero and the innovations in the ACD model are set to one at the start of each day. The lagged values of h and $\ln(\psi)$ at the beginning of each day are treated as parameters to be estimated although in interest of parsimony we restrict them to be the same across days.

We impose the diagonal structure for the B matricies but initially do not impose any of the symmetry conditions discussed in section 3.3. In the interest of conserving space we do not present results for all estimated models. An ACD(2,2) model often

[10] Models of this form are advocated in Bauwens and Giot (2000) who refer to it as a log-ACD$_2$ model, and applied in Engle and Lunde (2003). See Bauwens, Galli, and Giot (2002) for a discussion of the theoretical properties of this model.

provides a very good starting point for modeling durations and we begin by jointly estimating an ACM(2,2)-ACD(2,2) model given by (21) and (22) with r=w=2. We then test if the ACM(2,2) model is sufficient or if a higher order model is needed. The likelihood ratio test strongly rejects the null of an ACM(2,2)-ACD(2,2) in favor of an ACM(3,3)-ACD(2,2) model with a p-value value is $6.7\text{x}10^{-5}$. The ACM(3,3)-ACD(2,2) model, however, is not rejected for the ACM(4,4)-ACD(2,2) model with a p-value of .29. Further diagnostics suggest that no additional lags of the duration are needed in the ACM specification. Similarly, we also find that no additional lags of prices are needed in the ACD specification. The parameter estimates for the ACM(3,3)-ACD(2,2) with r=m=2 model are given in Table 2. Standard errors of the estimates are given in parenthesis.

Before discussing the parameter estimates we examine the model diagnostics. As discussed in section 3.2 the standardized residuals should be temporally uncorrelated. Figure 4 presents the correlations constructed for the standardized residuals given in (17). We again denote significant (at the 5% level) correlations with a "+" or "-" indicating the sign. The strong correlations at lag 1 have vanished. Furthermore, the long sets of positive correlations in the raw series have also disappeared. A formal test for the null hypothesis that the standardized series is white noise can be obtained from the test statistic in (19). The p-value is .31 providing no evidence of remaining correlation. The one-step ahead prediction errors associated with each state are not correlated with past errors indicating that the model is well specified.

Engle and Russell (1998) suggest using the standardized durations, given by $e_i = \frac{\tau_i}{\hat{\psi}_i}$, to asses the fit of the ACD model. Under correct specification this series should be distributed as an iid unit exponential. The Ljung Box test for the null hypothesis that the series is uncorrelated through the first 15 lags has a p-value of .11. The variance of the series indicates some remaining excess dispersion, however.

4.2 Interpretation of Results and Hypothesis Tests

We now turn to the interpretation of the results. We begin by summarizing the parameter estimates and testing the symmetry conditions suggested in section 3.3.

Finally we provide a detailed discussion of the nature of dependence between price changes and the durations implied by the model estimates. This relationship is related to market microstructure theories proposed in the literature.

The sum of the j^{th} diagonal elements of the B_i matricies of an ACM model characterize the persistence associated with the j^{th} state. This persistence measure is similar across all states with states 1&5 summing to 0.874 and 0.876 respectively and states 2&4 suming to .803 and .88 respectively. The impact of the contemporaneous and lagged durations on the j^{th} price transition probabilities are given by χ_{j1} and χ_{j2} respectively. The coefficients on the contemporaneous duration are positive and significant for all 4 states. Since the log odds is taken with respect to the base state of no price change this indicates that the probability of a price move increases with the duration. These estimates imply that the timing of transactions and the distribution of transaction by transaction price changes are related. We will investigate this relationship later in this section.

Further examination of terms in the *A* and *B* matricies reveals structure. In particular, for each *A* and *B* matrix, row 5 is roughly the reverse, or mirror image of row 1. Similarly, row 4 is roughly the reverse, or mirror image of row 2. This is exactly what we should expect to find if the price symmetry hypothesis discussed in section 3.3 holds. We now proceed to test several hypothesis of symmetry.

We first consider a test for dynamic price symmetry by jointly testing that the *A* and *B* matricies are all response symmetric. The likelihood ratio test associated with the null hypothesis for the *A* and *B* matricies, is marginally not rejected at the 5% level with a p-value of .058. Hence we are unable to reject symmetry in the marginal impact of the past price changes on the future log odds. We are also unable to reject that the constant vector *c* is symmetric. The p-value for this likelihood ratio test 0.39. However, we strongly reject the null hypothesis that the coefficient vectors on the durations χ_1 and χ_2 are symmetric vectors; the p-value is 1.05×10^{-3}. The parameter estimates for the price symmetric model with symmetric intercept but asymmetric duration impact are presented in Table 3.

The coefficients on the contemporaneous duration associated with the downward price moves are larger than those associated with upward price moves. Hence the probability of a downward price move increases more than the probability of an upward price move as the elapsed time since the last trade increases. Short elapsed time since the last transaction is associated with rising prices and long elapsed time since the last transaction is associated with falling prices.

We now examine the relationship between the transaction price change and the elapsed time. Given the filtration of past price changes and durations the variance of transaction price changes can be expressed conditionally as a function of the contemporaneous duration given a duration τ. We plot the transaction price variance as a function of τ fixing the lagged values at their unconditional means and evaluating the conditional variance of price changes obtained by varying the contemporaneous duration. This plot is given in Figure 5. The variance is an upward sloping function of the contemporaneous duration. For reference, each standardized unit of time is, on average, a little over 3 minutes. It is interesting to compare the conditional variance to the variance obtained when the log stock price follows a random walk. We calculate the variance of the open to close returns. The variance per unit of standardized time implied by Brownian motion is then calculated. The upward sloping line in Figure 5 is the Brownian motion variance as a function of the standardized time interval which is linear in elapsed time. We convert from the variance of returns to variance of price changes by multiplying by the square of the sample average price.

We might expect that the conditional transaction price variance should be larger than the Brownian motion variance since the transaction price process includes both the variance of Brownian motion and market microstructure effects such as price discreteness and other transitory effects. Even very short durations may result in a price change equal to or greater than the minimum tick size. Indeed, for very short durations the conditional variance from the ACM model is much higher than that implied by Brownian motion. What is interesting, however, is that the volatility is a concave function of the duration τ with slope becoming smaller than that of the Brownian motion, so that the variance per unit time declines with duration. For short duration trades the variance is above that of the Brownian motion volatility and for long duration trades it is below.

This analysis shows that for the Airgas stock, the timing of trades and not merely the passage of time affects volatility. This relationship is predicted by theoretical models where time varying transaction rates are driven by discretionary informed traders that only trade when they posses superior information For an early reference see Easley and O'Hara (1992). In these models, the price adjustments, and hence volatility, will be larger in periods of frequent trading and smaller when trades are infrequent.

Figure 5 also presents a plot of the expected transaction price change given a duration τ and the filtration of past price changes obtained using the same method as for the variance. The conditional mean is a downward sloping function of the time since the last trade. This result is in the same direction and much more significant than the mean effect found in Engle (2000)[11]. The fact that long durations are associated with falling prices is consistent with the theoretical model of Diamond and Verrecchia (1987) who suggest that in the presence of short selling constraints periods of infrequent trading are indicative of bad news. Agents that posses bad news about the asset and would like to short the asset may be unable to do so given short selling constraints.

The coefficient on the lagged duration is negative for all states. For the 2-tick price moves (states 1 and 5) the coefficient on the lagged duration is larger in magnitude than the coefficient on the contemporaneous duration. All else equal, the net effect of a long duration beyond one period is to decrease the probability of a large price change. The coefficient on the lagged duration is slightly smaller in magnitude for the one-tick price moves. The net effect of a long duration beyond one period is to slightly lower the probability of a one tick move. Long durations increase the probability of contemporaneous price moves but have a slight tendency to decrease the probability of price moves expected multiple periods ahead.

Examining the coefficients associated with the ACD model we find both lag one coefficients of the price change and its square are negative and significant indicating that durations are expected to be shorter following upward price moves and/or larger price changes. This effect is partially offset by the lag two coefficients that are positive and significant for both the price changes and squared price changes.

[11] In a previous version of the paper we found similar evidence of falling prices associated with longer durations for 3 months of IBM data from the TORQ data set.

5. Conclusion

This paper proposes modeling financial transactions data as a marked point process where the points are the transaction times and the associated marks are information about the transaction such as the price. We propose decomposing the joint density of arrival times and price changes into the product of a conditional distribution for the price changes and a marginal distribution for the arrival times. Institutional features restrict prices to fall on discrete values. For our sample the overwhelming majority of the price changes take one of just 5 different values. We therefore treat the price changes as a multinomial random variable and propose an autoregressive model for the price transition probabilities. Some theoretical properties of the Autoregressive Conditional Multinomial (ACM) model are examined. The ACD model of Engle and Russell (1998) is used for the marginal distribution of the arrival times.

We find little evidence of time of day effects in the distribution of transaction by transaction price for the stock analyzed. This has the interesting implication that the time of day patterns typically observed in within day volatility measured over fixed time intervals is an artifact of the diurnal patterns in the transaction rate.

The joint ACM-ACD model is estimated by maximum likelihood. A simple to general model selection approach suggests that moderately simple models appear adequate. This modeling approach provides us with a microscopic view of the intraday dynamics of asset prices. For the NYSE stock analyzed, we test for and find evidence of a type of symmetry in the marginal impacts of the history of price changes on the transition probabilities.

We find that the transaction price variance increases with the duration but at a slower rate than would be implied by simple geometric Brownian motion. In fact, the variance is virtually constant after a length of time equal to the mean duration has passed since the last transaction. This is consistent with predictions from Easley and O'Hara (1992) where the absence of transactions is indicative of no private information in the market and slow adjustment of prices. Conversely, rapid transactions are associated with the presence of informed trading so prices adjust quickly.

We also test for and find that long durations are associated with falling prices. This result is consistent with the theoretical predictions of Diamond and Verrecchia (1987) where short selling constraints suggest that long durations are indicative of "bad news".

We believe that the ACM-ACD model may provide a useful tool for analyzing other discrete, potentially irregularly spaced data such as credit risk dynamics or marketing data where consumers face a discrete product choice set such as different brands.

Recently the NYSE and NASDQ completed its move to decimalization. It is also worth noting that for many stocks the histogram of stock price changes measured in ticks looks remarkably similar in post decimalization transactions data. Hence there is no reason to think the model will not perform equally well on the current decimalized data. The ACD model may well provide a good approach to analyzing how these changes affect the transaction costs (effective cost) or price dynamics.

Appendix A

Summary of superscript/subscript notation

For the random variable y

y_i denotes the random variable associated with transaction i.

$y^{(i-1)} = (y_{i-1}, y_{i-2}, \ldots)$

For a random vector w

w_i denotes a vector associated with transaction i

w_{ij} denotes the j^{th} element of a vector associated with transaction i

Unless otherwise specified, vectors with tildes have dimension k and vectors without tildes have dimension k-1.

For a matrix W

W_j denotes the j^{th} matrix

W_{mn} denotes the m,n element

For a constant vector c

c_m denotes the m^{th} element.

Appendix B

Proof of Theorem 1

The ACM(p,q) model with a vector of constants χ is given by:

(1') $$\left(I-\sum_{j=1}^{q}B_jL^j\right)h(\pi_i)=\left(\sum_{j=1}^{p}A_jL^j\right)(x_i-\pi_i)+\chi$$

Since the roots of $\left|I_{k-1}-B_1z-B_2z^2-\ldots-B_qz^q\right|=0$ lie outside the unit circle then we can write:

(2') (2') $$h(\pi_i)=\sum_{j=1}^{\infty}\Psi_j\left(x_{i-j}-\pi_{i-j}\right)+\chi^*$$

where $\chi^*=\left(I-\sum_{m=1}^{q}B_m\right)^{-1}\chi<\infty$ and $\sum_{m=1}^{\infty}\left|\Psi_m\right|<\infty$.

$|\Psi|$ is element by element absolute value. Next, note that the elements of $|x_i-\pi_i|$ must lie between zero and one inclusive so it follows:

(3') $$\left|h(\pi_i)\right|\le\sum_{j=1}^{\infty}\left|\Psi_j\left(x_{i-j}-\pi_{i-j}\right)\right|+\left|\chi^*\right|\le\sum_{j=1}^{\infty}\left|\Psi_j\right|\left|\left(x_{i-j}-\pi_{i-j}\right)\right|+\left|\chi^*\right|\le\sum_{j=1}^{\infty}\left|\Psi_j\right|\iota+\left|\chi^*\right|<\infty$$

where ι is a k-1 vector of ones. So $h(\pi_i)$ is bounded. Bounds on the probabilities are then obtained by setting

(4') $$\begin{aligned}M^u&=\exp\left(\sum_{j=1}^{\infty}\left|\Psi_j\right|\iota+\left|\chi^*\right|\right)\\ M^l&=\exp\left(-\sum_{j=1}^{\infty}\left|\Psi_j\right|\iota-\left|\chi^*\right|\right)\end{aligned}$$

where the exponential function is understood to be element by element. Since the probabilities are given by the logistic transformation it follows that that the elements of π_i are bounded strictly away from zero.

(5') $$\pi_i\ge\frac{1}{\left(\iota' M^u+1\right)}M^l>0$$

Proof of Proposition 1

(6') $$Qh(\pi_i.) = \sum_{j=1}^{p} QA_j\left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} QB_j h\left(\pi_{i-j}\right) + Q\chi Z_i$$

From the symmetry assumptions it follows that

(7') $$Qh(\pi_i.) = \sum_{j=1}^{p} A_j Q\left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} B_j Qh\left(\pi_{i-j}\right) + \chi Z_i$$

So, given fixed initial conditions $\pi_1,\ldots,\pi_p$ the mirror image history generates the mirror image log odds. Finally, recall that the i^{th} element of h is simply the $\log(\pi_i / \pi_0)$ so it follows that

$Q\pi_i\left(x_{i-1}, x_{i-2},\ldots Z_i\right) = \pi_i\left(Qx_{i-1}, Qx_{i-2},\ldots Z_i\right)$

Proof of Proposition 2

We must show that Ψ_j of (2') are response symmetric.

(8') $$\Psi(L) = \left(I - \sum_{j=1}^{q} B_j L^j\right)^{-1}\left(\sum_{j=1}^{p} A_j L^j\right)$$

First notice that $Q^{-1}=Q$ so that

(9') $$\begin{aligned} Q\left(I - \sum_{j=1}^{q} B_j L^j\right)^{-1} &= \left(Q^{-1}\left(I - \sum_{j=1}^{q} B_j L^j\right)\right)^{-1} = \left(\left(I - \sum_{j=1}^{q} B_j L^j\right)Q^{-1}\right)^{-1} \\ &= \left(I - \sum_{j=1}^{q} B_j L^j\right)^{-1} Q \end{aligned}$$

The last equality follows from the fact that the prices are dynamic symmetric.

Response symmetric Ψ_j then follows immediately from the response symmetry of

$\left(I - \sum_{j=1}^{q} B_j L^j\right)^{-1}$ and $\left(\sum_{j=1}^{p} A_j L^j\right)$.

Appendix C

Bin	Freq.	Percent
-0.8125	1	0.0054
-0.75	2	0.0108
-0.6875	0	0.0000
-0.625	0	0.0000
-0.5625	0	0.0000
-0.5	1	0.0054
-0.4375	2	0.0108
-0.375	1	0.0054
-0.3125	0	0.0000
-0.25	14	0.0754
-0.1875	34	0.1831
-0.125	482	2.5952
-0.0625	3009	16.2009
0	11463	61.7186
0.0625	3039	16.3625
0.125	444	2.3906
0.1875	49	0.2638
0.25	21	0.1131
0.3125	4	0.0215
0.375	1	0.0054
0.4375	1	0.0054
0.5	1	0.0054
0.5625	0	0.0000
0.625	2	0.0108
0.6875	1	0.0054
0.75	0	0.0000
0.8125	1	0.0054

References

1. Ane, T., and H. Geman 1999, Stochastic Volatility and Transaction Time: an Activity Based Volatility Estimator, *The Journal of Risk* 2, 18

2. Berndt, E., B. Hall, R. Hall, and J. Hausman, 1974, Estimation and Inference in Nonlinear Structural Models, *Annals of Economic and Social Measurement*, 3, 653-665

3. Bauwens, L., F. Galli, and P. Giot, 2002, The Moments of Log-ACD Models, CORE Discussion Paper

4. Bausens, L., and P. Giot, 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

5. Bauwens, L., and P. Giot 2003, Asymmetric ACD models: introducing price information in ACD models with a two-state transition model *Empirical Economics* 28(4), 709-731

6. Bauwens, L., P. Giot, J. Grammig, and D. Veredas, 2003, A Comparison of Financial Duration Models Via Density Forecasts, forthcoming *International Journal of Forecasting*

7. Cox, D. R., 1981, Statistical Analysis of Time Series: Some Recent Developments. *Scandinavian Journal of Statistics*, 8, 93-115

8. Diamond, D.W., and R.E. Verrecchia, 1987, Constraints on Short-selling and Asset Price Adjustments to Private Information, *Journal of Financial Economics* 18, pp 277-311

9. Easley, D., and M. O'Hara, 1992, Time and the Process of Security Price Adjustment. *The Journal of Finance* 19, 69-90

10. Engle, R., 2000, The Econometrics of Ultra-High Frequency Data, *Econometrica,* 2001, 1-22

11. Engle, R., and A. Lunde, 2003, Trades and Quotes: A Bivariate Point Process, *Journal of Financial Econometrics* 1, 159-188

12. Engle, R., and A. Patton, 2003, Impacts of Trades in an Error-Correction Model of Quote Prices, forthcoming *Journal of Financial Market*

13. Engle, R., and J. Russell, 1995, Autoregressive Conditional Duration: A New Model for Irregularly Spaced Data, University of California, San Diego Working Paper Series

14. Engle, R., and J. Russell, 1998, Autoregressive Conditional Duration: A New Model for Irregularly Spaced Data, *Econometrica* 66, 5, 1127-1162

15. Ghysels, E., and J. Jasiak, 1998, GARCH for Irregularly Spaced Data: The ACD-GARCH model, *Studies in Nonlinear Economics and Econometrics* 2, 133-149

16. Grammig J., and M. Wellner, 2002, Modeling the Interdependence of Volatility and Intertransaction Duration Processes, *Journal of Econometrics* 106, 369-400

17. Harris, Lawrence E., 1990, Estimation of Stock Price Variance and Serial Covariances from Discrete Observations, *Journal of Financial and Quantitative Analysis* 25, 291-306

18. Hasbrouck, J., 1991, Measuring the Information Content of Stock Trades, *The Journal of Finance* 66,1, 179-207

19. Hausman, J., A. Lo, and C. MacKinlay, 1992, An Ordered Probit Analysis of Transaction Stock Prices, *Journal of Financial Economic* 31, 319-379

20. Jacobs, P., and P. Lewis, 1983, Stationary Discrete Autoregressive-Moving Average Time Series Generated by Mixtures. *Journal of Time Series Analsysis.* 4, 19-36.

21. Jones, C., G. Kaul, and M. Lipson, 1994, Transactions, Volume, and Volatility. *Review of Financial Studies* 7, 631-651

22. Kalbfleisch, J., and R. Prentice, 1980, *The Statistical Analysis of Failure Time Data*, John Wiley & Sons.

23. Lancaster, T., 1990, *The Econometric Analysis of Transition Data* Cambridge University Press

24. Li, W. K., and A.I. Mcleod, 1981, Distribution of the Residual Autocorrelations in Multivariate ARMA Time Series. *Journal of the Royal Statistical Society*, B, 43, 231-239.

25. MacDonald, I., and W. Zucchini, 1997, *Hidden Markov and Other Models for Discrete-valued Time Series*. Chapman & Hall

26. MacRae, E., 1977, Estimation of Time-Varying Markov Processes with Aggregate Data. *Econometrica* 45, 183-198

27. McInish, T., and R. Wood, 1992, An Analysis of Intradaily Patterns in Bid/Ask Spreads for NYSE Stocks, *Journal of Finance*, 47, 753-764.

28. Nelson, D., 1991, "Conditional Heteroskedasticity in asset returns: A new Approach, *Econometrica* 59, 347-370

29. Prigent, J., E. Renault, and O. Scaillet, 2001, An Autoregressive Conditional Binomial Option Pricing Model, Selected Papers form the First World Congress of the Bachelier Finance Society

30. Rydberg, T., and N. Shephard, 2000, "A Modelling Framework for the Prices and Times Made on the New York Stock Exchange", in W.J. Fitzgerald, R.L. Smith, A.T. Walden and P. C. Young (eds.), *Non-Stationary and Non-Linear Signal Extraction* (Cambridge: Issac Newton Institute Series, Cambridge University Press), 217-246

31. Rydberg, T., and N. Shephard, 2003, "The Dynamics of Trade-by-Trade Price Movements", *Journal of Financial Econometrics* 1, 2-25.

32. Shephard, N.,1995, "Generalized Linear Autoregressions", unpublished manuscript, Nuffield College, Oxford

33. Tiao, G., and G. Box, 1981, "Modeling Multiple Time Series with Applications", *Journal of the American Statistical Association* 76, 802-816

Table 1: Estimates of the deterministic pattern for durations and states.

	Const.	d_1	d_2	d_3	d_4	d_5	d_6	*F stat p-value*
Durations	239.30 (11.83)	-16.84 (17.02)	65.20 (14.24)	43.08 (15.03)	-41.88 (15.10)	-63.26 (14.39)	-118.25 (17.69)	0.00%
Down 2	.030 (.0053)	-0.0034 (.0076)	0.0083 (.0064)	-0.0000 (.0067)	-.0153 (.0067)	0.00625 (.0064)	0.00818 (.00795)	6.432%
Down 1	0.165 (.0117)	-.0077 (.0168)	.0060 (.0140)	-.0046 (.0148)	.00659 (.0149)	.00773 (.0142)	-.0292 (.0175)	61.3%
Up 1	.156 (.0117)	.0147 (.0169)	-.0106 (.0141)	.0196 (.0149)	-.029 (.0149)	.0030 (.0142)	.0227 (.0175)	25.6%
Up 2	.0395 (.0052)	-.0172 (.0075)	.0106 (.0063)	-.0042 (.0066)	-.0059 (.0067)	.0053 (.0064)	.0013 (.0078)	23.3%

Table 2: Parameter estimates for unrestricted ACM(3,3)-ACD(2,2) model

	State 1		State 2		State 4		State 5	
	c_1	-.4541 (.0794)	c_2	-.2327 (.0537)	c_4	-.1474 (.0394)	c_5	-.4424 (0867)
A_1	a_{11}	-.2955 (.3471)	a_{21}	-.7524 (.1910)	a_{41}	.8684 (.1211)	a_{51}	2.712 (.1439)
	a_{12}	-.4420 (.1763)	a_{22}	-.4429 (.0731)	a_{42}	.9362 (.0524)	a_{52}	.7013 (.1276)
	a_{14}	.9612 (.1153)	a_{24}	.9690 (.0528)	a_{44}	-.4162 (.0727)	a_{54}	.0023 (.1579)
	a_{15}	2.307 (.1443)	a_{25}	.7599 (.1239)	a_{45}	-.7185 (.2012)	a_{55}	.6706 (.2333)
A_2	a_{11}	1.178 (.3302)	a_{21}	.5554 (.1958)	a_{41}	-.7684 (.1852)	a_{51}	-1.037 (.3594)
	a_{12}	.5397 (.1998)	a_{22}	.3341 (.0858)	a_{42}	-.6031 (.0959)	a_{52}	-.3378 (1943)
	a_{14}	-.3928 (.1616)	a_{24}	-.3867 (.0867)	a_{44}	.3536 (.0958)	a_{54}	.2414 (.1912)
	a_{15}	-.6338 (.2643)	a_{25}	-.1502 (.1648)	a_{45}	.6007 (.2252)	a_{55}	.0663 (.2872)
A_3	a_{11}	-.0323 (.2179)	a_{21}	.3863 (.1441)	a_{41}	.1415 (.1454)	a_{51}	-.4421 (.3243)
	a_{12}	-.1220 (.1404)	a_{22}	-.1502 (.1648)	a_{42}	-.1482 (.0788)	a_{52}	-.0995 (.1521)
	a_{14}	-.3557 (.1393)	a_{24}	-.2381 (.0746)	a_{44}	.1550 (.0731)	a_{54}	.1022 (.1365)
	a_{15}	-.3678 (.2392)	a_{25}	-.0922 (.1429)	a_{45}	.2684 (.1476)	a_{55}	.3905 (.2311)
B_1	b_{11}	.7765 (.0734)	b_{22}	.6879 (.0581)	b_{44}	.8329 (.0696)	b_{55}	.7866 (.1155)
B_2	b_{11}	.0225 (.0838)	b_{22}	.1655 (.0540)	b_{44}	.0698 (.0733)	b_{55}	.0601 (.1383)
B_3	b_{11}	.0754 (.0567)	b_{22}	-.0370 (.0337)	b_{44}	-.0227 (.0366)	b_{55}	.0292 (.0735)
χ	χ_{11}	.3155 (.0290)	χ_{21}	.2612 (.0136)	χ_{41}	.2047 (.0133)	χ_{51}	.2081 (.0305)
	χ_{12}	-.3412 (.0285)	χ_{22}	-.2277 (.0156)	χ_{42}	-.1825 (.0139)	χ_{52}	-.2274 (.0304)

ACD	
ω	-.0588 (.0076)
α_1	.0688 (.0040)
α_2	-.0059 (.0085)
β_1	.8692 (.1384)
β_2	.1086 (.1356)
ρ_1	-.0330 (.0064)
ρ_2	.0117 (.0067)
ζ_1	-.0291 (.0044)
ζ_2	.02221 (.0043)

ACM: $$h(\pi_i) = c + \sum_{j=1}^{p} A_j \left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} B_j h(\pi_{i-j}) + \chi \begin{pmatrix} \ln(\tau_i) \\ \ln(\tau_{i-1}) \end{pmatrix}$$

ACD: $$\ln(\psi)_i = \omega + \sum_{j=1}^{u} \alpha_j \varepsilon_{i-j} + \sum_{j=1}^{v} \beta_j \ln(\psi_{i-j}) + \rho_1 \Delta p_{i-1} + \rho_2 \Delta p_{i-2} + \zeta_1 \Delta p_{i-1}^2 + \zeta_2 \Delta p_{i-2}^2$$

$$A = \begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,4} & a_{1,5} \\ a_{2,1} & a_{2,2} & a_{2,4} & a_{2,5} \\ a_{4,1} & a_{4,2} & a_{4,4} & a_{4,5} \\ a_{5,1} & a_{5,2} & a_{5,4} & a_{5,5} \end{bmatrix} B = \begin{bmatrix} b_{11} & 0 & 0 & 0 \\ 0 & b_{22} & 0 & 0 \\ 0 & 0 & b_{44} & 0 \\ 0 & 0 & 0 & b_{55} \end{bmatrix} \chi = \begin{bmatrix} \chi_{11} & \chi_{12} \\ \chi_{21} & \chi_{22} \\ \chi_{41} & \chi_{42} \\ \chi_{51} & \chi_{52} \end{bmatrix}$$

Table 3: Parameter estimates for restricted ACM(3,3)-ACD(2,2) model

	State 1		State 2		State 4		State 5	
	c_1	-.4567 .(0692)	c_2	-.2008 (.0435)	c_4	-.2008 -	c_5	-.4567
A_1	a_{11}	.2797 (.1899)	a_{21}	-.7440 (.1385)	a_{41}	.8161 -	a_{51}	2.513 -
	a_{12}	-.2046 (.1167)	a_{22}	-.4290 (.0522)	a_{42}	.9508 -	a_{52}	.8348 -
	a_{14}	.8348 (.0857)	a_{24}	.9508 (.0380)	a_{44}	-.4290 -	a_{54}	-.2046 -
	a_{15}	2.513 (.1013)	a_{25}	.8161 (.0871)	a_{45}	-.7440 -	a_{55}	.2797 -
A_2	a_{11}	.5736 (.2073)	a_{21}	.5871 (.1466)	a_{41}	-.4382 -	a_{51}	-.8774 -
	a_{12}	.3973 (.1382)	a_{22}	.3368 (.0639)	a_{42}	-.4731 -	a_{52}	-.3861 -
	a_{14}	-.3861 (.1258)	a_{24}	-.4731 (.0678)	a_{44}	.3368 -	a_{54}	.3973 -
	a_{15}	-.8774 (.2100)	a_{25}	-.4382 (.1241)	a_{45}	.5871 -	a_{55}	.5736 -
A_3	a_{11}	.1691 (.1588)	a_{21}	.3352 (.1041)	a_{41}	.0115 -	a_{51}	-.3360 -
	a_{12}	-.0273 (.0976)	a_{22}	.1683 (.0502)	a_{42}	-.2114 -	a_{52}	-.2104 -
	a_{14}	-.2104 (.1013)	a_{24}	-.2114 (.0558)	a_{44}	.1683 -	a_{54}	-.0273 -
	a_{15}	-.3360 (.1863)	a_{25}	.0115 (.1009)	a_{45}	.3352 -	a_{55}	.1691 -
B_1	b_{11}	.8032 (.0632)	b_{22}	.7401 (.0459)	b_{44}	.7401 -	b_{55}	.8032 -
B_2	b_{11}	-.0028 (.0723)	b_{22}	.1315 (.0448)	b_{44}	.1315 -	b_{55}	-.0028 -
B_3	b_{11}	.0721 (.0449)	b_{22}	-.0327 (.0243)	b_{44}	-.0327 -	b_{55}	.0721 -
χ	χ_{11}	.3133 (.0284)	χ_{21}	.2584 (.0135)	χ_{41}	.2073 (.0132)	χ_{51}	.2071 (.0301)
	χ_{12}	-.3378 (.0275)	χ_{22}	-.2274 (.0146)	χ_{42}	-.1789 (.0143)	χ_{52}	-.2298 (.0297)

ACD	
ω	-.0582 (.0076)
α_1	.0689 (.0040)
α_2	-.0067 (.0084)
β_1	.8779 (.1374)
β_2	-.1001 (.1346)
ρ_1	-.0331 (.0064)
ρ_2	.0118 (.0067)
ζ_1	-.0290 (0044)
ζ_2	.0220 (.0043)

ACM: $$h(\pi_i) = c + \sum_{j=1}^{p} A_j\left(x_{i-j} - \pi_{i-j}\right) + \sum_{j=1}^{q} B_j h(\pi_{i-j}) + \chi\begin{pmatrix}\ln(\tau_i)\\ \ln(\tau_{i-1})\end{pmatrix}$$

ACD: $$\ln(\psi)_i = \omega + \sum_{j=1}^{u}\alpha_j\varepsilon_{i-j} + \sum_{j=1}^{v}\beta_j \ln(\psi_{i-j}) + \rho_1\Delta p_{i-1} + \rho_2\Delta p_{i-2} + \zeta_1\Delta p_{i-1}^2 + \zeta_2\Delta p_{i-2}^2$$

$$A=\begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,4} & a_{1,5} \\ a_{2,1} & a_{2,2} & a_{2,4} & a_{2,5} \\ a_{4,1} & a_{4,2} & a_{4,4} & a_{4,5} \\ a_{5,1} & a_{5,2} & a_{5,4} & a_{5,5} \end{bmatrix} B=\begin{bmatrix} b_{11} & 0 & 0 & 0 \\ 0 & b_{22} & 0 & 0 \\ 0 & 0 & b_{44} & 0 \\ 0 & 0 & 0 & b_{55} \end{bmatrix} \chi=\begin{bmatrix} \chi_{11} & \chi_{12} \\ \chi_{21} & \chi_{22} \\ \chi_{41} & \chi_{42} \\ \chi_{51} & \chi_{52} \end{bmatrix}$$

Figure 1: One day of transactions data.

a) Transaction Price

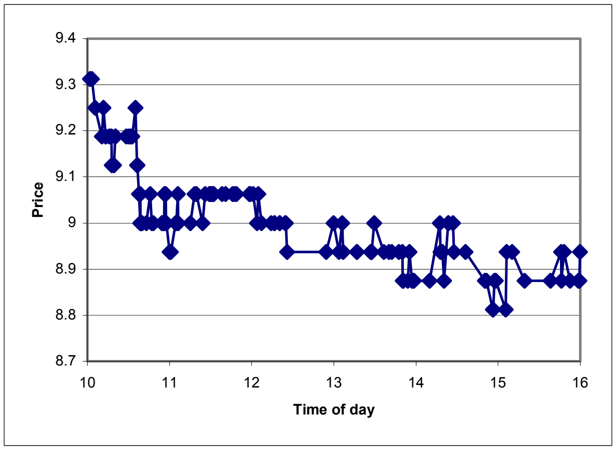

b) Transaction price changes

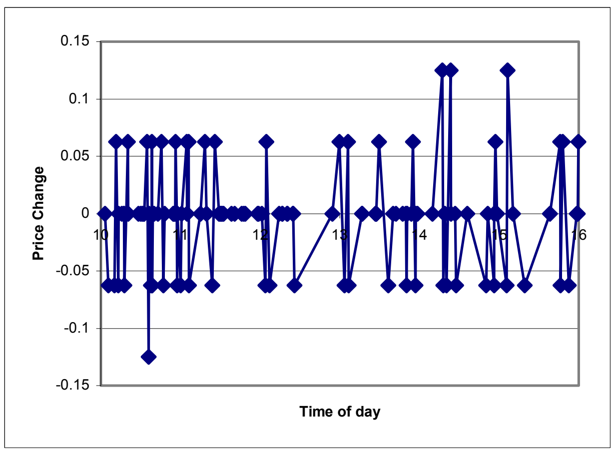

c) Number of transactions that have occurred.

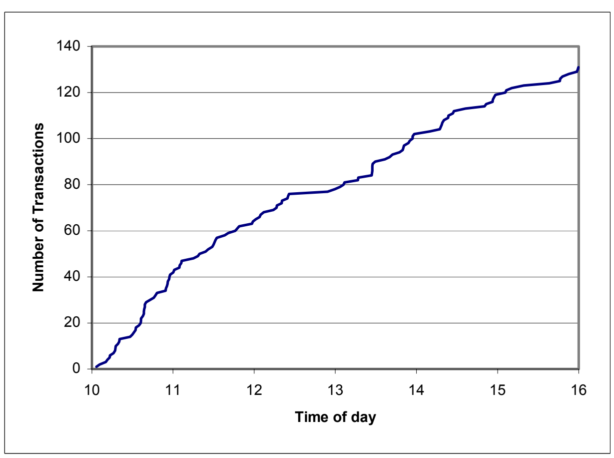

Figure 2: Histogram of Transaction Prices

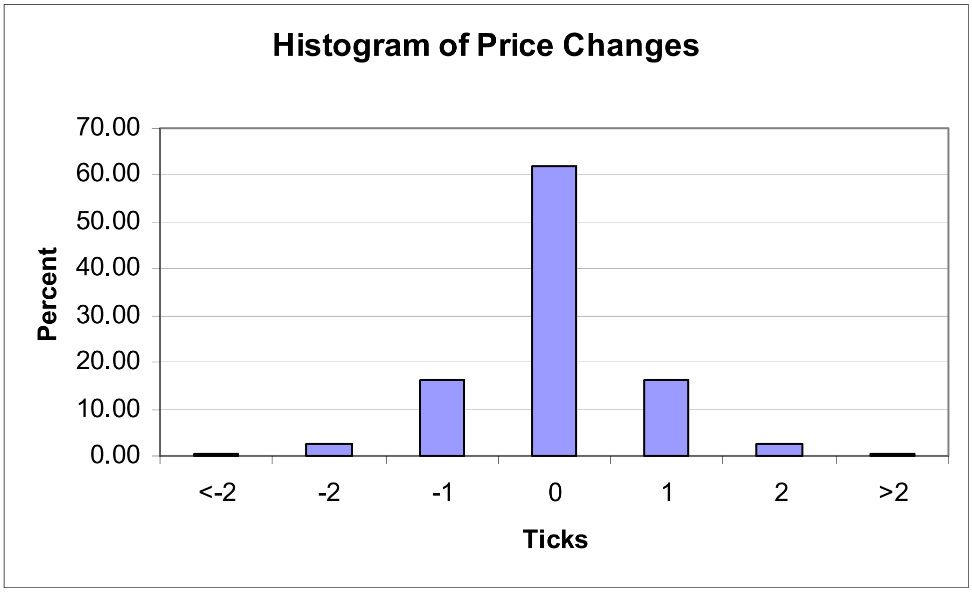

Figure 3: Box Tiao Representation of Sample Cross Correlations of x

$$R_s = \frac{1}{N-(s+1)} \sum_{i=s+1}^{N} x_i x_{i-s}' \qquad \mathrm{P}_s = R_0^{-1} R_s$$

s= 1 2 3 4 5

$$\begin{bmatrix} - & - & + & + \\ - & - & + & + \\ \cdot & + & - & - \\ + & + & \cdot & - \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & \cdot & + & + \\ \cdot & + & \cdot & \cdot \\ + & \cdot & \cdot & + \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & + & \cdot & \cdot \\ \cdot & \cdot & + & \cdot \\ + & \cdot & \cdot & + \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & + & \cdot & \cdot \\ \cdot & \cdot & + & \cdot \\ + & \cdot & \cdot & + \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & + & \cdot \\ + & \cdot & \cdot & + \end{bmatrix}$$

6 7 8 9 10

$$\begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & \cdot & + & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ + & \cdot & \cdot & + \end{bmatrix} \quad \begin{bmatrix} + & - & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ + & \cdot & \cdot & + \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & + \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & \cdot & + & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ + & + & \cdot & + \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & + & \cdot \\ \cdot & \cdot & \cdot & + \end{bmatrix}$$

11 12 13 14 15

$$\begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ + & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ + & \cdot & + & + \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & + \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ + & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & + \\ \cdot & \cdot & - & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & + \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & + & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & + & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix}$$

Figure 4: Box Tiao Representation of Sample Cross Correlations of Standardized Residual Vector

s= 1 2 3 4 5

$$\begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ + & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & + \end{bmatrix}$$

6 7 8 9 10

$$\begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & + & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & + & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & - & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} + & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & - & \cdot \\ \cdot & \cdot & + & + \\ \cdot & \cdot & \cdot & + \\ \cdot & + & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix}$$

11 12 13 14 15

$$\begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} - & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & - & - & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix} \quad \begin{bmatrix} \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \\ \cdot & + & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot \end{bmatrix}$$

Figure 5. Conditional Mean and Variance of Price Changes as a Function of Time Elapsed Since Previous Transaction