

UC San Diego

Recent Work

Title

Econometric Analysis of Discrete-Valued Irregularly-Spaced Financial Transactions Data Using a New Autoregressive Conditional Multinomial Model

Permalink

<https://escholarship.org/uc/item/00m2c5hk>

Authors

Russell, Jeffrey
Engle, Robert F

Publication Date

1998-04-01

98-10

UNIVERSITY OF CALIFORNIA, SAN DIEGO

DEPARTMENT OF ECONOMICS

ECONOMETRIC ANALYSIS OF DISCRETE-VALUED
IRREGULARLY-SPACED FINANCIAL TRANSACTIONS DATA USING A
NEW AUTOREGRESSIVE CONDITIONAL MULTINOMIAL MODEL

BY

JEFFREY R. RUSSELL

AND

ROBERT F. ENGLE

**DISCUSSION PAPER 98-10
APRIL 1998**

Econometric analysis of discrete-valued irregularly-spaced financial transactions data using a new Autoregressive Conditional Multinomial model*

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

February 1998

This paper proposes a new approach to modeling financial transactions data. A new model for discrete valued time series is proposed in the context of generalized linear models. Since the model is specified conditional on both the previous state, as well as the historic distribution, we call the model the Autoregressive Conditional Multinomial (ACM) model. When the data are viewed as a marked point process, the ACD model proposed in Engle and Russell (1998) allows for joint modeling of the price transition probabilities and the arrival times of the transactions. In this marked point process context, the transition probabilities vary continuously through time and are therefore duration dependent. Finally, variations of the model allow for volume and spreads to impact the conditional distribution of price changes. Impulse response studies show the long run price impact of a transaction can be very sensitive to volume but is less sensitive to the spread and transaction rate.

Keywords: Discrete valued time series, marked point process, high frequency data.

The authors would like to thank David Brillinger, Xiaohong Chen, Clive Granger, Alex Kane, Bruce Lehman, Peter McCullagh, Glenn Sueyoshi, George Tiao, and Hal White for valuable input. The first author is grateful for 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 would like to acknowledge financial support from the National Science Foundation grant. SBR-9422575

** University of Chicago, Graduate School of Business
email: jeffrey.russell@gsb.uchicago.edu

*** University of California, San Diego
email: rengle@weber.ucsd.edu

1. Introduction

The recent development and distribution of high frequency transaction by transaction financial data has generated a large amount of research from both theoretical and empirical market microstructure perspectives. Frequently, empirical market microstructure issues cannot be addressed on intertemporally aggregated data since the questions at issue involve the potentially dynamic impact of characteristics of *individual* trades such as volume, whether the trade was buyer or seller initiated, or the impact of particular sequences or frequency of trades.

Our primary econometric interest is the dynamics of the price process and potentially its interaction with other features of the market. The price is only observed, however, at particular points in time when transactions occur. These transaction times are not equally spaced in time and Engle and Russell (1998) provide strong evidence that the arrival rate of traders is intertemporally correlated. That is, trades tend to cluster in time in both a deterministic and stochastic manner. In addition to observing the price at these points in time, each point has other associated characteristics such as the volume and the spread. Following Engle and Russell(1998) we treat the arrival times as a point process consider jointly modeling arrival times and price changes possibly as a function of predetermined or weakly exogenous variables.

We propose decomposing the joint distribution of price changes and arrival times into the product of the conditional distribution of price changes and the marginal distribution of the arrival times. Engle and Russell (1998) suggest the Autoregressive Conditional Duration (ACD) model for the marginal distribution of arrival times so we now turn our attention to the conditional distribution of price changes. Since transactions prices are required to fall on discrete quantities, usually $1/8^{\text{th}}$ of a dollar¹, we view the discrete price changes as multinomial time series data. Market microstructure issues such as bid ask bounce, inventory control behavior of the specialist, price smoothing requirements of the specialist, and dynamic strategic behavior all suggest a rich class of dynamics will be required to successfully capture the price dynamics. While bid ask bounce induces strong negative correlation in price changes at high frequencies the other characteristics mentioned above are likely characterized by longer range dependence. We therefore propose a new class of models for multinomial time series data that is able to account for these dynamic features. Because the model depends on both the historic distribution of the data as well as past realizations, the model is called the Autoregressive Conditional Multinomial (ACM) model.

We show that the model can be interpreted in the context of a competing risks model. The waiting time associated with the i^{th} transaction can exit into one of several states corresponding to discrete price movements. We also develop measures of the instantaneous expected price change and the instantaneous expected volatility. Expressing the transition probabilities in continuous time we examine the relationship between price distribution and trading rates. From these expression it can be seen that transaction rates have potentially two ways of affecting the volatility. First, the distribution of price changes from one transaction to the next may depend on the contemporaneous duration or on the expected duration. Second, this expression provides an explicit link between

¹ For the IBM data analyzed in this paper, 99.3% of the price changes fall on just 5 unique values.

transaction rates and the rate at which the process evolves. This idea has been referred to as time deformation as studied by Tauchen and Pitts, Andersen, and Ghysels to name a few. Hence simply speeding up or slowing down the transaction rates can affect volatility when measured in calendar time as is traditionally done.

Estimation is performed on the joint likelihood function. The relationship between price changes and the arrival rate of traders is examined. We show that both the expected duration and the realized duration affect the distribution of price changes at the transaction level. Price dynamics are further examined via impulse response studies. We find that while spreads and expected duration between transactions can affect the long run price impact of a transaction or sequence of transactions the affects of volume can be much more pronounced.

The paper is organized as follows. Section 2 introduces the ACM model. Section 3 suggests some parameter restrictions motivated by economic intuition. Section 4 examines the model from a continuous time perspective with duration dependence. Section 5 introduces the data. Section 6 presents results for various models and section 7 examines volume and impulse response functions. Finally, section 8 concludes.

Section 2. The Autoregressive Conditional Multinomial Model

In this paper we view the transaction price process as a marked point process. In this context the arrival times of the transactions are denoted by t_i . At each transaction time t_i there is an associated realization of the price of the asset denoted by y_i . It is convenient to measure these as changes from the previous transaction price. Since transaction prices fall on discrete values we assume that y_i can take on K values ($k=1,2,\dots,K$). We are interested in modeling the conditional joint distribution of price changes and arrival times denoted by:

$$(1) \quad f(y_i, t_i | y^{(i-1)}, t^{(i-1)})$$

where $y^{(i-1)} = (y_{i-1}, y_{i-2}, \dots)$ and $t^{(i-1)} = (t_{i-1}, t_{i-2}, \dots)$

In the spirit of Engle (1996) we decompose the joint distribution of the mark and the arrival time into the product of the conditional distribution of the mark and the marginal distribution of the arrival times.

$$(2) \quad f(y_i, t_i | y^{(i-1)}, t^{(i-1)}) = g(y_i | y^{(i-1)}, t^{(i)}) q(t_i | y^{(i-1)}, t^{(i-1)})$$

where $g(\cdot)$ denotes the density function associated with the discrete valued random variable y_i conditional on the current arrival time and the filtration of y and t . $q(\cdot)$ denotes the density function of the waiting time between the t_{i-1} and t_i arrival times conditional on just the filtration of y and t . Engle and Russell (1998) propose the Autoregressive Conditional Duration (ACD) model specification for $q(\cdot)$ and find the model is able to

explain transaction arrival rates for IBM transactions data. Hence, we now focus our attention on the conditional distribution $g(x)$.

Toward this end, we restrict our attention to the class of observation driven models in the sense of Cox (1981)². We propose a new class of multinomial time series models. The probability of each state is modeled via a multivariate ARMA structure allowing for complex dynamic structure in the conditional distribution. A similar structure is proposed by Shephard(1995) as a GLAR, a generalized linear autoregression.

Let x_i denote a $(K-1)$ dimensional random vector where the k^{th} element of x_i is one if $y_i=k$ [$k=1, \dots, K-1$] occurred and zero otherwise. In this example, the state will correspond to a particular transaction price change and uniquely determines y_i . Denote the conditional expectation of x_i by

$$(3) \quad \mathbf{p}_i \equiv E\left(x_i | Z_i, x^{(i-1)}\right)$$

where Z_i might consist of t_i or other weakly exogenous variables in the sense of Engle Hendry and Richard (1987) or perhaps deterministic functions of time. We arbitrarily omit state K since the probability of state K is given by $(1-\mathbf{1}'\mathbf{p}_i)$ where $\mathbf{1}$ denotes the $(K-1)$ unit vector. Hence, \mathbf{p}_i uniquely describes the distribution of y_i conditional on the filtration of x and Z . The conditional covariance matrix of x can similarly be defined as

$$(4) \quad V_i \equiv V(x_i | Z_i, x^{(i-1)}) = \text{diag}\{\mathbf{p}_i\} - \mathbf{p}_i \mathbf{p}_i'$$

We now consider parameterizations for (3). Of course, \mathbf{p}_i must satisfy all the usual conditions associated with a distribution function for a discrete valued random variable and must have no error term since it is defined as a conditional expectation. In particular, $\mathbf{1}'\mathbf{p}_i \leq 1$ and the j^{th} element of \mathbf{p}_i denoted by \mathbf{p}_i^j must be positive for all i and j . A natural modeling strategy would be to assume that an appropriate transformation of the conditional expectation \mathbf{p}_i is some function of the conditioning variables. That is, for some appropriate link function $h(\cdot): (K-1) \rightarrow (K-1)$ such as the logistic or probit, and a measurable function \mathbf{h} ,

$$(5) \quad \mathbf{h}(\mathbf{p}_i) = \mathbf{h}(E(x_i | x^{(i-1)})) = \mathbf{h}(x^{(i-1)})$$

Equation (5) is a type of Generalized Linear Model in a time series context. Clearly the success of (5) in characterizing the dynamics of x_i lies in the choice of \mathbf{h} and $h(\cdot)$.

We define the ACM model specification as a linear function of its own past and the innovations in x , potentially interacted with Z . That is,

² 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 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. Given the success of ARMA models for continuous valued time series we are optimistic in our approach which will provide an ARMA structure for discrete valued time series.

$$(6) \quad h(\mathbf{p}_i) = \sum_{j=1}^p A_j M_{i-j} (x_{i-j} - \mathbf{p}_{i-j}) + \sum_{j=1}^q B_j x_{i-j} + \sum_{j=1}^r C_j h(\mathbf{p}_{i-j}) + GZ_i$$

Since the probability structure at time i depends both on the historic distribution as well as the past realizations we call this model the Autoregressive Conditional Multinomial model. The most simple version of the model which only depends on the history of the price process might be referred to as an ACM(p,q,r) model. The matrix M can be taken as $V^{-1/2}$, with V the conditional covariance matrix of x , or as the diagonal elements of this covariance matrix, or simply as the identity matrix. In some applications it could even be taken to depend upon predetermined variables.

The structure of this equation is recursive. At the time of the $i-1$ transaction, knowing all past x and π gives from (6) a calculated value of the next π . Consequently, subject to some starting values, the full history of π can be constructed from observations on x and z . This allows evaluation of the likelihood function and its numerical derivatives.

Several particular cases of this specification are familiar. Static models of probabilities have this form with $A=B=C=0$. When $K=2$, and the link function is simply the identity function, this is the linear probability model. In the same setting, if

$$h(\mathbf{p}) = \log(\mathbf{p} / (1 - \mathbf{p})),$$

the log odds ratio, then the model is the logistic. For the probit, $h(\cdot) = F^{-1}(\cdot)$ where F is the cumulative standard normal distribution function. For more than two states, the natural models are multinomial logit and probit. Hausman Lo and MacKinlay (1992) for example used an ordered probit to analyze financial transaction prices. In the logit case,

$$(7) \quad h(\mathbf{p}_j) = \log(\mathbf{p}_j / \mathbf{p}_K), \text{ for } j = 1, \dots, K-1$$

Dynamic models of course must include lagged information. A Markov chain requires only one past state to initiate all future probabilities. In this case, $h(\cdot)$ is the identity function and $A=C=0$, $Z=1$ and $q=1$. Higher order Markov chains set $q>1$. For full generality, additional terms in $x_{i-k} \otimes x_{i-j}$ for $j,k>0$ may then be needed. In this notation the first order Markov chain can be expressed as³

$$(8) \quad \mathbf{p}_i = Bx_{i-1} + \mathbf{m}$$

This model has a steady state set of probabilities $\bar{\mathbf{p}} = (I - B)^{-1} \mathbf{m}$ as long as all eigenvalues of B lie inside the unit circle. The parameterization of such a Markov chain in terms of transition probabilities insures that all probabilities will lie between zero and one.

Substituting for μ gives

$$(9) \quad \mathbf{p}_i = \bar{\mathbf{p}} + B(x_{i-1} - \bar{\mathbf{p}})$$

and multistep forecasts:

$$(10) \quad E_i \mathbf{p}_{i+k} = \bar{\mathbf{p}} + B^{k-1} (x_{i+1} - \bar{\mathbf{p}})$$

The introduction of additional information from the past, relaxes the Markov structure and may improve the performance of the model. Consider the simple linear ACM model with $B=0$, $Z=1$ and $p=r=1$, in the following parameterization:

$$(11) \quad \mathbf{p}_i = A(x_{i-1} - \mathbf{p}_{i-1}) + C\mathbf{p}_{i-1} + \mathbf{m} = Ax_{i-1} + (C - A)\mathbf{p}_{i-1} + \mathbf{m}$$

When C has all its eigenvalues within the unit circle, this model also has multistep forecasts and steady state probabilities given by

³ The intercept appears because of the elimination of the equation for the K^{th} state.

$$(12) \quad E_i(\mathbf{p}_{i+k}) = \bar{\mathbf{p}} + C^{k-1}(\mathbf{p}_{i+1} - \bar{\mathbf{p}}), \quad \bar{\mathbf{p}} = (I - C)^{-1} \mathbf{m}$$

Defining $\mathbf{e} = x - \mathbf{p}$, an innovation, equation (11) is seen to be a vector ARMA:

$$(13) \quad x_i = Cx_{i-1} + \mathbf{m} + (A - C)\mathbf{e}_{i-1} + \mathbf{e}_i$$

where clearly the eigenvalues of C control the long run properties. The parameters in A as well as those in C determine whether all probabilities lie in the unit interval and the dynamic response to particular states.

LEMMA 1 The probabilities in (11) will all lie between zero and one if

- a) All elements of A , $C-A$ and μ are non-negative
- b) $\max\{\text{column sums } A\} + \max\{\text{column sums } (C - A)\} + \text{column sum}(\mathbf{m}) \leq 1$

Proof : Appendix

When other variables are included in the model such as Z or more lags, it becomes very difficult to ensure that the probabilities lie in the unit interval. Hence it is attractive to use a link function $h(\cdot)$ to bound the probabilities. Just as for the static model, the logit specification (7) is a very simple and attractive link. However it becomes more difficult to investigate the dynamic properties of the ACM⁴.

Consider first the version of (6) with only one lag and no exogenous variables or lags of x by itself.

$$(14) \quad h(\mathbf{p}_i) = A(x_{i-1} - \mathbf{p}_{i-1}) + Ch(\mathbf{p}_{i-1}) + \mathbf{m}$$

The multistep forecasts of h can be obtained exactly as before. If all eigenvalues of C lie inside the unit circle, then

$$(15) \quad E_i(h(\mathbf{p}_{i+k})) = \bar{h} + C^{k-1}(h(\mathbf{p}_{i+1}) - \bar{h}), \quad \bar{h} = (I - C)^{-1} \mathbf{m}$$

Because h is a 1-1 mapping from probabilities to \mathbb{R}^{K-1} , (15) can be uniquely solved for the steady state probabilities, $\bar{\mathbf{p}}$. These probabilities have the property that if x and π are equal to the steady state probabilities in period i , they also will in the next period.

Furthermore, the average fraction of periods spent in each state will approach $\bar{\mathbf{p}}$. This conjecture follows from the ergodicity of h , which further implies that π is ergodic. Such results and the corresponding conditions must be developed more rigorously.

With more lags in (14), conditions can easily be found for a stationary solution for h and for π . If the innovations in (14) were multiplied by M , as in equation (6), completely similar results are available. However in the more general set-up of (6) it does not appear possible to find an explicit formula for the steady state probabilities although often they can be computed.

The log likelihood of the ACM model expressed as the sum of the conditionals is simply

$$(16) \quad L = \sum_{i=1}^N \sum_{k=1}^K (x_i^j \log(\mathbf{p}_i^j)) = \sum_{i=1}^N x_i' \log(\mathbf{p}_i).$$

⁴ Some special cases have been considered in the literature. If $K=2$, h is the log odds, and $p=r=0$ the model reduces to a q^{th} order linear logistic model first suggested by Cox (1971, 1981) and more recently discussed by Zegar and Qaqish.

If the data are irregularly spaced and the conditional distribution of price changes depends on the timing of transactions then joint estimation of (2) may be required.

3. Parameter Restrictions and Price Dynamics

Depending on the number of states and number of lags there are potentially a large number of parameters to be estimated. Economic intuition may guide us in imposing certain restrictions in the model specification. In particular, there is a certain type of symmetry that we might expect in the dynamic process of price movements. In particular, the marginal impact of the state “down 1 tick” on the conditional probability of a subsequent “up tick” may be the same as the marginal impact of the state “up 1 tick” on the conditional probability of a subsequent “down tick”. Similar relations might be expected to hold true for other states.

Consider the case for the linear probability model in (11). Without loss of generality, arrange the elements of x in the natural ordering implied by the transaction prices (i.e. lowest to highest). We omit the zero price movement state. As an example consider a simple 3 state model of transaction prices. One possible ordering is state 1 is a downward price movement, state 2 is a zero price movement and state 3 is an upward price movement. Now restrict our attention to the simple linear model specified in (11). If state 2 is the base state the symmetry intuition suggests the following parameter restrictions in an ACM model with $q=r=1$ and $p=0$.

$$(17) \quad \mathbf{m} = \begin{bmatrix} \mathbf{m}_1 \\ \mathbf{m}_1 \end{bmatrix} \quad A = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 \\ \mathbf{a}_2 & \mathbf{a}_1 \end{bmatrix} \quad C = \begin{bmatrix} \mathbf{c}_1 & \mathbf{c}_2 \\ \mathbf{c}_2 & \mathbf{c}_1 \end{bmatrix}$$

We see that α_1 characterizes the impact of a lagged downward (upward) price movement on the probability of another downward (upward) price movement. Similarly, α_2 characterizes the impact of a lagged downward (upward) price movement on the probability of another upward (downward) price movement. The parametrization of C implies a similar symmetry for the impact of the historic probability on the future distribution.

We might also expect higher order lags of A and C to have this symmetric response structure. We emphasize that these restrictions do not in any way imply that the conditional distribution will be symmetric. The shocks and their persistence will determine the shape of the distribution. It is only the marginal impact of the shocks and their decay rate that is assumed to be symmetric. The following definitions help to generalize these restrictions.

Definition 1: An $N \times N$ matrix Z is response symmetric if for the $N \times N$ matrix Q defined by

$$(18) \quad Q = \begin{bmatrix} 0 & & 1 \\ & \ddots & \\ 1 & & 0 \end{bmatrix}$$

$QZ = ZQ$. That is, Q and Z commute.

Definition 2: A vector z of length N is symmetric if $Qz = z$.

A generalization of the symmetric parameter restrictions in (11) to a $(K-1)$ state vector is then defined by the requirement that A and C are $(K-1) \times (K-1)$ response symmetric matrices and ω is a symmetric vector of dimension $K-1$.

The symmetry restrictions have an additional implication about the unconditional transaction price distribution described in the following theorem:

Theorem 1

Consider the a linear ACM(1,1,1) model defined using the identity link function. If the following conditions hold:

- i. $(B+C)$ has eigenvalues inside the unit circle,
- ii. w is a symmetric vector, and
- iii. B and C are response symmetric

then $p_{i+k} = E(x_{i+k}|I_i)$ converges to a symmetric vector.

Proof in appendix.

This theorem implies that as we forecast farther out and the impact of past shocks die out, the expected transaction price change approaches zero while the cumulative price change is potentially non-zero. For these very short time periods, the riskless rate is essentially zero so the Martingale assumption is plausible. In implementing these symmetry conditions it is convenient (but not necessary) to choose the zero price change state as the omitted state for purposes of estimation. Clearly this restriction reduces the number of parameters to be estimated by half. If this restriction is valid there are potentially large gains in efficiency by imposing them.

The intuition surrounding the symmetric response restrictions still holds for the logistic model. In particular, the log odds is parameterized as response symmetric when the logistic link function is used rather than the probabilities themselves as in the linear probability model. While the Theorem above is only proven for the linear ACM model, simulations as well as our intuition suggest that similar results hold for the logistic link function. The non-linearities associated with the logistic link function, however, greatly complicate the proof. These more complicated scenarios are currently being pursued.

A final model restriction that we consider in this paper is diagonal matrix specification for C_j . In this case, shocks to the log odds decay at a geometric rate determined by the diagonal elements of the C_j matrices. Thus the impact of new information is generously specified while the long run decay is more parsimoniously formulated.

4. A closer look at the ACM model with duration dependence.

Section 2 developed a flexible framework for modeling the dynamics of discrete price changes conditional on the filtration of price changes and the past distribution of price changes and Z . We now return to the joint distribution of arrival times and price changes. Following (1) we consider the joint distribution as the product of the marginal distribution of durations and the distribution of price changes conditional on not only the

filtration of arrival times and the historic distribution, but also on the contemporaneous duration.

Equation (1) can be viewed as a competing risks model. Classic multiple failure time data with competing risks models used in the analysis of unemployment spells, strikes, or medical studies generally consist of large cross section and short time series dimensions. The joint model of arrival times and discrete price changes developed in this paper is a competing risk model for time series data.

The hazard function characterizes the instantaneous probability of exiting to state k at time $t_{i-1} + t$ conditional on the i^{th} transaction not having occurred by time $t_{i-1} + t$. Expressed as a function of the duration t , the hazard function for state k can as follows:

$$(19) \quad \mathbf{q}_k(t|I_{i-1}) = \lim_{dt \rightarrow 0} \frac{\Pr(t \leq T < t + dt, Y = k | T > t, I_{i-1})}{dt}$$

So for small dt , $\mathbf{q}_k(t|I_{i-1})dt$ is the probability the i^{th} transaction is at price k and occurs by time $\tau+dt$ given survival to time τ . Let $\tilde{\mathbf{p}}_i$ denote the full K dimensional vector where the k^{th} element is the conditional probability that the k^{th} state is realized. That is, if the K^{th} state was omitted then.

$$(20) \quad \tilde{\mathbf{p}}(t) \equiv \begin{bmatrix} \mathbf{p}_i(t) \\ 1 - \mathbf{1}'\mathbf{p}_i(t) \end{bmatrix}$$

where $\mathbf{1}$ is a $(K-1)$ vector of ones. Noting that the conditional probability that the i^{th} event has not occurred by time $t_{i-1} + t$ is obtained from the marginal distribution of t the K dimensional vector of transition intensities can be expressed as:

$$(21) \quad \mathbf{q}(t|I_{i-1}) = h(t|I_{i-1})\tilde{\mathbf{p}}_i(t)$$

where

$$(22) \quad h(t|I_{i-1}) = \frac{q(t|I_{i-1})}{1 - \int_0^t q(s|I_{i-1})ds}$$

$h(t|I_{i-1})$ is the hazard function⁵ that characterizes the waiting times and q is given in (2). Loosely speaking, $h(t|I_{i-1})$ describes the conditional arrival rate of traders. When the hazard function is determined by the ACD specification of Engle and Russell (1998) we refer to this model as an ACM(p,q,r)-ACD(s,t).

The transition intensities in (21) provide an interesting perspective of the relationship between the traditional distribution of price changes measured in calendar time and transaction rates. First, suppose that the price changes were i.i.d. Then roughly speaking the volatility as measured in calendar time would be proportional to the number of transactions that occurred in that time interval which is proportional to the transaction rate. Hence even if time didn't affect the distribution of price changes from transaction to

⁵ See Kalbfleish and Prentice or, more recently, Lancaster for references on duration models.

transaction we would expect classical calendar time measures of volatility to be positively related to the trading frequency as observed in empirical studies such as Jones, Kaul and Lipson (1994) or McInish and Wood (1991).

A primary feature of this paper, however, is that the distribution of price changes from transaction to transaction is not likely to be i.i.d., but rather depends on, among other things, the waiting time between transactions as (21) suggests. Hence, trading rates have two potential impacts on the distribution of price changes as measured in calendar time. On the one hand the transaction price process evolves at a stochastic rate. On the other, the arrival rate of traders has an impact on the transaction by transaction distribution of prices. The proposed model captures both features.

Engle and Russell (1998) show that the ACD model is in the class of accelerated failure time models. In particular, if I_0 is the baseline hazard then the hazard function can be expressed as:

$$(23) \quad h(\mathbf{t}|I_{i-1}) = \frac{1}{\mathbf{y}_i} I_0\left(\frac{\mathbf{t}}{\mathbf{y}_i}\right)$$

The arrival rate of traders as characterized by the expected i^{th} waiting time ψ_i affect time flow in two ways. The rate at which time progresses through the baseline hazard varies with the inverse of ψ_i . Additionally, the level of the baseline hazard is inversely related to ψ_i . If the arrival rate of traders controls the flow of time then it would be reasonable that π depends not only on τ but also on ψ_i as in⁶:

$$(24) \quad \mathbf{q}(\mathbf{t}|I_{i-1}) = \frac{1}{\mathbf{y}_i} I_0\left(\frac{\mathbf{t}}{\mathbf{y}_i}\right) \tilde{\mathbf{p}}_i\left(\frac{\mathbf{t}}{\mathbf{y}_i}\right)$$

Now, the flow of calendar time is proportional to the arrival rate of traders.

To examine various relationships between the price distribution and arrival rates define Δp and Δp^2 , be K dimensional vectors with k^{th} elements given by the price change if state k occurs and the square of that price change if state k occurs respectively. Then the expectation of the transaction price change at time $t_{i-1} + \mathbf{t}$ over the next instant (the instantaneous conditional mean) is given by

$$(25) \quad \mathbf{m}(\mathbf{t}) = \Delta p' \mathbf{p}_i(\mathbf{t}) h(\mathbf{t}|I_{i-1})$$

where $t = t_{i-1} + \mathbf{t}$

Similarly, the instantaneous expected volatility is given by

$$(26) \quad \mathbf{s}^2(\mathbf{t}) = \Delta p^2' \mathbf{p}_i(\mathbf{t}) h(\mathbf{t}|I_{i-1})$$

The unconditional mean and squared transaction price change over the i^{th} duration can be obtained by integrating \mathbf{t} out of (25) and (26) respectively:

⁶ Since most link functions introduce a nonlinear relationship between probabilities and the conditioning variables the exact form of the dependence of the probabilities on ψ in (24) may be difficult to impose.

$$(27) \quad \mathbf{m}_i = \int_0^{\infty} \Delta p' \mathbf{p}_i(s) z(s|I_{i-1}) ds$$

$$(28) \quad \mathbf{s}_i^2 = \int_0^{\infty} \Delta p^2 \mathbf{p}_i(s) z(s|I_{i-1}) ds.$$

The relationship between prices and trading rates can be examined over more than one transaction but calendar time results (such as volatility per unit time) will in general require simulations.

5. The IBM Transaction Data

This section of the paper applies the ACM model to transaction data for IBM. The data were extracted from the Trades Orders Reports and Quotes (TORQ) data set constructed by J. Hasbrouck and the NYSE. 58,944 transactions were recorded for IBM over the 3 months of trading on the consolidated market from November 1990 through January 1991. The average transaction price for the sample is \$111.04 with a standard deviation of \$2.80. A histogram of the transaction price changes is presented in figure 1. We see that 69% of the transaction prices are unchanged from their previous value. The distribution is relatively symmetric with 14.0% and 14.2% up one tick and down one tick respectively. Up and down two ticks occurred with almost identical frequency at 1.0%. Up and down by more than two ticks occurred with frequency 0.3% and 0.4% respectively.

Of the 58,944 transactions there are only 53,857 unique times. Of the transactions occurring at non-unique trading times, 87% corresponded to a zero price movement. This suggests that these transactions may reflect large orders that were broken up into smaller pieces. It is not clear that each piece should be considered a separate order, hence the zero second durations were considered to be a single transaction and were deleted from the data set. In the case where the prices differ, the transaction price for that time is taken as the first transaction price observed in the sequence of zeros.

Following Engle and Russell (1998) the first half hour of the trading day is omitted. This is to avoid modeling the opening of the market which is characterized by a call auction followed by heavy activity. The dynamics are likely to be quite different over this period. The entire first half hour is deleted since the opening auction transactions are not recorded at the same time each morning.

Finally, the data set has 46,047 remaining observations with 64.3% corresponding to zero price movement and 15.8% corresponding to 1 tick and down and one tick up each. Finally, up 2 ticks and down 2 ticks correspond to frequencies 1.3% and 1.4 % respectively. All other price movements greater than 2 ticks have a combined frequency of 1.4%.

In order to keep the number of parameters manageable and to avoid problems of data sparseness we choose a five state model. Following the discussion of response symmetric matrices our model will be identified by normalizing with respect to the zero price change. The state vector is defined as follows:

$$(29) \quad x_i = \begin{cases} [1,0,0,0]' & \text{if } \Delta p_i < -.125 \\ [0,1,0,0]' & \text{if } -.125 \leq \Delta p_i < 0 \\ [0,0,0,0]' & \text{if } \Delta p_i = 0 \\ [0,0,1,0]' & \text{if } 0 < \Delta p_i \leq .125 \\ [0,0,0,1]' & \text{if } \Delta p_i > .125 \end{cases}$$

Hence state 1 occurs if the transaction price changes by more than one tick down. State 2 occurs if the price moves by just one tick down. State 3 occurs for a zero price move and states 4 and 5 occur when the price increases by 1 tick and more than one tick respectively.

A natural measure of intertemporal dependence is based on the intertemporal cross correlations of the vector x_i . In order to present the (cross) correlation structure in a user friendly way we adopt a method proposed by Tiao and Box(1981). The intertemporal cross correlations are presented in matrix form with the numbers replaced by the symbols “+”, “-“, and “.”. A dot indicates that the (cross) correlation is not significant at the 1% level. A plus and minus indicate a positive and negative significant (cross) correlation respectively. The 99% confidence intervals⁷ are calculated using the approximation of $2.58 * n^{-1/2}$.

Denoting the sample mean of x_i by \bar{x} , the m^{th} sample cross correlation matrix is calculated by

$$(30) \quad P_m = R_0^{-1} R_m$$

where $R_m = \frac{1}{N - (m + 1)} \sum_{i=m+1}^N (x_i - \bar{x})(x_{i-m} - \bar{x})'$

Figure 2 presents the Box Tiao representation for lags 1 through 15. The r,s element of the m^{th} matrix gives the correlation of state r with state s lagged m periods. The sample cross correlations for $m=1$ are easily interpreted in the context of bid ask bounce. The upper right and lower left quadrants represent price reversals. The positive signs are reflective of bid ask bounce. The upper left quadrant and lower right quadrant correspond to price continuations. For example, the (1,4) element suggests that the probability of moving down two ticks (state 1) is positively correlated with the event up two ticks last period (state 4). Examining row 2 we find that the probability of moving down 1 tick (state 2) is negatively correlated with the event down 1 last period (suggested by the (2,2) element) and positively correlated with the events up 1 last period and up 2 last period (suggested by the (2,3) and (2,4) elements respectively).

Moving to lags beyond the first we see the 4 plus signs in the center of the matrix suggest that states 2 and 4 (corresponding to down 1 and up 1 tick respectively) are not correlated with past occurrence of the 2 tick price movements but are correlated with each other. This is indicative of the bid ask bounce as the price “bounces” back and forth between buy and sell orders for many transactions at a time.

⁷ Due to the very large number of observations we use a 99% confidence level.

Finally, only the diagonal elements and the elements corresponding to the correlation between the two extreme states of up and down 2 remain positive and significant out through lag 15. The extreme states appear to exhibit the strongest intertemporal correlation since they are significant (and positive) for every lag. This raises the intuitively appealing possibility that it is the occurrence of the extreme states that carry the most information about the future of the price distribution.

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. This is exactly what we would expect to see if the symmetry restrictions suggested in section 3 are correct.

6. Model Estimates for IBM Transaction Price Data

In this section we estimate various ACM models using a logistic link function. There are several reasons that we chose the logistic link function over other possibilities. Russell (1996) found that the linear ACM model suggested in (11) does not satisfy the conditions stated in Lemma 1 that ensure all the probabilities lie between zero and one. The logistic model will ensure that all probabilities lie in $[0,1]$. Also, the logistic ACM model has the nice interpretation that the log odds follows an ARMA type structure.

Choosing the logistic link model is only the first step. It is clear that a very rich class of models are given in (6). Furthermore, the models are estimated using numerical maximization techniques of the likelihood function which can be time consuming for the large sample sizes and potentially large numbers of parameters to be estimated. Hence we choose simple to general model selection procedure. Initially we restrict our attention to “pure” ACM(p,q,r) model, that is, models that only depend on the history of the price process. Later we consider the affects of the contemporaneous duration and then other predetermined variables such as volume and spreads. Hence we begin by estimating models of the form

$$(31) \quad h_i(\mathbf{p}_i) = \mathbf{m} + \sum_{j=1}^p A_j V_{i-j}^{-1/2} (x_{i-j} - \mathbf{p}_{i-j}) + \sum_{j=1}^q B_j x_{i-j} + \sum_{j=1}^r C_j h(\mathbf{p}_{i-j}).$$

Here, V_i is the $(K-1) \times (K-1)$ diagonal matrix with the (k,k) element given by the k^{th} element of $\text{diag}(\mathbf{p}_i(1-\mathbf{p}_i)')$. Initially we set $p=q=r=2$. We maintain the 5 state model of up 2 or more ticks, up 1 tick, no change, down 1 tick and down 2 or more ticks. The state vector is defined by (29) and we implement the symmetry conditions discussed in section 3. Hence the model is identified by normalizing the log odds of the zero price change to unity and omitting that state. Imposing the symmetric response restrictions discussed in section 3 and restricting the matrix C to be diagonal yields the following structures:

$$(32) \quad \mathbf{m} = \begin{bmatrix} \mathbf{m}_1 \\ \mathbf{m}_2 \\ \mathbf{m} \\ \mathbf{m}_2 \end{bmatrix} \quad \mathbf{A} = \begin{bmatrix} \mathbf{a}_{1,1} & \mathbf{a}_{1,2} & \mathbf{a}_{1,4} & \mathbf{a}_{1,5} \\ \mathbf{a}_{2,1} & \mathbf{a}_{2,2} & \mathbf{a}_{2,4} & \mathbf{a}_{2,5} \\ \mathbf{a}_{2,5} & \mathbf{a}_{2,4} & \mathbf{a}_{2,2} & \mathbf{a}_{2,1} \\ \mathbf{a}_{1,5} & \mathbf{a}_{1,4} & \mathbf{a}_{1,2} & \mathbf{a}_{1,1} \end{bmatrix} \quad \mathbf{B} = \begin{bmatrix} \mathbf{b}_{1,1} & \mathbf{b}_{1,2} & \mathbf{b}_{1,4} & \mathbf{b}_{1,5} \\ \mathbf{b}_{2,1} & \mathbf{b}_{2,2} & \mathbf{b}_{2,4} & \mathbf{b}_{2,5} \\ \mathbf{b}_{2,5} & \mathbf{b}_{2,4} & \mathbf{b}_{2,2} & \mathbf{b}_{2,1} \\ \mathbf{b}_{1,5} & \mathbf{b}_{1,4} & \mathbf{b}_{1,2} & \mathbf{b}_{1,1} \end{bmatrix} \quad \mathbf{C} = \begin{bmatrix} \mathbf{c}_{1,1} & 0 & 0 & 0 \\ 0 & \mathbf{c}_{2,2} & 0 & 0 \\ 0 & 0 & \mathbf{c}_{2,2} & 0 \\ 0 & 0 & 0 & \mathbf{c}_{1,1} \end{bmatrix}$$

The dynamic structure of the data associated with two consecutive trades from the closing transaction one evening to the opening transaction the next morning is unlikely to be the same as the dynamic structure associated with two consecutive trades within the same day. Hence, we reinitialize variables to their unconditional means at the beginning of each day. Furthermore, as in Engle and Russell (1998) we omit the opening trades since they are not generated by the same trading mechanism. This is done by omitting the first half hour of recorded trades each morning from 9:30 to 10:00.

The models are estimated by maximum likelihood using the Berndt, Hall, Hall and Hausman (1974) (BHHH) algorithm. Numerical derivatives are necessary because the analytic expression for the scores is defined recursively as function of past partial derivatives similar to the GARCH class of models for volatility studied by Bollerslev (1985).

6.1 Parameter Estimates for the simple ACM(p,q,r) model

In the interest of saving space parameter estimates for only selected models will be presented. We first estimate an ACM(2,2,2) model. Since the dimension of h and x is equal to the number of states less one the coefficient matrices A , B , and C are 4×4 . ω is a vector with dimension 4. The symmetry condition implies that we only need to estimate $(K-1)/2 + (p+q+r)(K)(K-1)/2$. With $K=5$ and $p=q=r=2$ this corresponds to 62 parameters. Imposing the diagonal restriction on C_j suggested in section 3 the number of parameters to be estimated is reduced to 46.

The LM test for an additional lag of each term yields a test statistic of 70.53. The test statistic is calculated by taking the R^2 from the first iteration of the BHHH algorithm with the initial values of the parameters set to the maximum likelihood estimates of the restricted model⁸. Due to the very large sample size we use a 1% critical value. With 22 degrees of freedom the 1% critical value is 40.29 hence the null hypothesis is easily rejected in favor increasing the order of the model.

The LM test associated with the null of an ACM(3,3,3) against the alternative of an ACM(4,4,4) is not rejected. The test statistic is 20.72 with a corresponding p-value of about 40%. We present the estimated parameters of the ACM(3,3,3) model in table 1. In the interest of saving space only the upper half of the matrices are presented.

States 1 and 5 are the extreme states of down and up two ticks or more respectively. States 2 and 4 correspond to down one and up one tick respectively. Generally all the parameters are significant at the 5% level with only a few parameters corresponding to states 2 and 4 not significant.

As a further diagnostic check, we turn our attention to the K dimensional vector of residuals defined by

⁸ See Berndt, Hall, Hall, and Hausman (1974) for a more complete description.

$$(33) \quad v_i^* = x_i - \hat{p}_i$$

where \hat{p}_i denotes the estimated conditional expectation of x_i .

Standardized residuals are then obtained by pre-multiplying x_i by the Cholesky factorization of the conditional variance covariance matrix associated with x_i :

$$(34) \quad v_i = U_i v_i^* \text{ where } U_i V_i U_i = I$$

V_i is the conditional variance covariance matrix of x_i at time t_i given by (4).

Correct specification and true parameter values imply that

$$(35) \quad E(v_i | I_{i-1}) = 0 \quad \text{and} \quad E(v_i v_i' | I_{i-1}) = I$$

The sample cross correlations associated with the standardized residuals are calculated by

$$(36) \quad P_m = \frac{1}{N - (m + 1)} \sum_{i=m+1}^N v_i v_{i-m}'$$

The cross correlations are presented in figure 4. The first and second order cross correlations still have several significant elements. 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

$$(37) \quad Q = N \sum_{m=1}^M \text{Trace}(P_m P_m')$$

The test statistic has a χ^2 distribution with $(K-1)^2 * M$ degrees of freedom.

The test statistic based on the first 15 sample cross correlations is 423.0. The 1% critical value is 293.1 so the null is rejected⁹. The Q-statistic based on the original series, however is 23324.5. So while the test suggests remaining intertemporal correlation the model has accounted for a great deal of the intertemporal correlation. Additional lags do not significantly improve the statistic.

Figure 3 is a Tiao Box plot for the cross correlations of the standardized residuals. The long sets of positive cross correlations are not apparent in this series. Furthermore, the vast majority of the correlation matrices contain no significant correlations. We see that for the first two lags there are 5 significant correlations suggesting we might consider a more elaborate model. Rather than pursue further dynamic lag structures investigate modeling strategies that allow the contemporaneous duration to affect the conditional distribution of price changes.

6.2 Parameter Estimates for Models with Duration Dependent Probabilities.

⁹ We are concerned about the validity of this test statistic. In particular, since the conditional probabilities of the extreme states (up and down 2 or more ticks) are frequently very small (on the order of 10^{-5}) the standardization by premultiplying by the inverse of the cholesky factor may be problematic. In the univariate case it would be as if we were dividing a number that is occasionally 1 but often near zero by (roughly) the square root of 10^{-5} . This test statistic or perhaps another one that is better suited is the subject of current research.

In this section we expand the simple ACM(p,q,r) model to allow for duration dependence. We allow durations to enter both in terms of the realized duration and the expectation of the duration. We use the ACM(3,3,3) specification discussed in the previous section and add several variables. We allow the contemporaneous duration and the expected duration obtained from the ACD model to enter linearly into the log odds specification of (6). Additionally, we put in the expected duration as a measure of the current market activity as well as the logarithm of the ratio of the duration and the expected duration. Since we are using a logistic link function, any variable entering in the logarithm implies a relationship between the percent change in that variable and the percent change in the log odds.

$$(38) \quad h(\mathbf{p}_i) = \mathbf{m} + \sum_{j=2}^3 A_j V_{i-j}^{-1/2} (x_{i-j} - \mathbf{p}_{i-j}) + \sum_{j=1}^3 B_j x_{i-j} + \sum_{j=1}^3 C_j h(\mathbf{p}_{i-j}) + \ln(\mathbf{t}_i) g_1 + \mathbf{t}_i g_2 + (\mathbf{t}_i / \mathbf{y}_i) g_3 + \mathbf{y}_i g_4$$

where τ_i is the waiting time associated with the i^{th} transaction. That is, $\tau_i = t_i - t_{i-1}$. \mathbf{y}_i is the conditional expectation of the i^{th} waiting time. Given the success in Engle and Russell (1998) we model this conditional expectation with the exponential ACD(2,2) model expressed as

$$(39) \quad \mathbf{y}_i = E(\mathbf{t} | t_{i-1}, t_{i-2}, \dots) = \mathbf{w} + \mathbf{a}_1 \mathbf{t}_{i-1} + \mathbf{a}_2 \mathbf{t}_{i-2} + \mathbf{b}_1 \mathbf{y}_{i-1} + \mathbf{b}_2 \mathbf{y}_{i-2}$$

The parameter vectors $g_1, g_2, g_3,$ and g_4 are restricted to be symmetric in the sense of definition 1 stated above. Since we don't have any reason to suspect that the contemporaneous duration or its expectation should have an asymmetric impact on the distribution of price changes, this symmetry restriction appears to be a reasonable starting point.

The expected contemporaneous duration enters the conditional likelihood of the transaction price changes so the durations cannot be considered weakly exogenous in the sense of Engle Hendry and Richard (1987). Hence we efficiently estimate by maximum likelihood using the joint distribution of price changes and arrival times. The durations are first deseasonalized using a two step procedure suggested in Engle and Russell (1995). After the deseasonalization, the durations have an unconditional expectation of unity.

In the interest of saving space, we only present the estimated parameters for the duration and expected duration terms in table 2. An LM test for the addition of these variables strongly rejects the null of the ACM(3,3,3) in favor of this expanded model. The test statistic is 156.4 with a critical value of 20.9. The economic impact appears to be small given the level of the estimates.

To get a better understanding of the impact of these variables on the conditional distribution of price changes, we set all the explanatory variables equal to their sample means and plot the conditional distribution of price changes as a function of the realized duration in figure 4 and the expected duration in figure 5. The probability is on the vertical axis and the normalized duration is on the horizontal axis. The normalized durations should be interpreted as the fraction above or below the mean duration by time of day. The realized duration appears to have only a slight impact on the price distribution. We see that very short durations suggest relatively smaller probabilities of 1 tick price moves. Although it is difficult to see in the graph, the probabilities associated with two tick moves are slowly falling a total of 10% as the duration ranges from .25 to 5.

The expected duration, however, has a very noticeable impact on the price distribution. Very rapid transaction rates (short expected durations) are associated with a higher probabilities of price movements. This is obvious in the one tick probabilities and the two tick probabilities increase by 8.6% as the expected duration ranges from .25 to 5.

Figure 6 plots the expected squared price change associated with the distributions in the previous plots. The normalized durations take on a larger range of values in sample than do their expectation so the scaling on these plots are not the same. The volatility plotted against the realized duration is a concave function. Very short and very long durations imply smaller volatility. Short durations may be associated with large trades that have been broken up into smaller pieces and perhaps should not be considered as separate transactions.

Figure 7 is a plot of volatility versus the expected duration. We see that volatility is a monotonically decreasing function of the expected duration. A slower market is associated with lower volatility all else equal. This is very much in agreement with predictions from Easley and O'Hara (1992) who suggest that more frequent transaction rates are due to a larger fraction of informed traders. In a rational expectations environment the specialist knows this and will make prices more sensitive to order flow when transactions are frequent.

7. Models with Other Weakly Exogenous Variables and an Impulse Response Study

At the heart of modern theoretical market microstructure is the question of how new information is incorporated into asset prices. If all relevant information were publicly available and all agents agreed on the impact this information should have on the price then prices would adjust immediately to any new information. On the other hand, if not all agents have equal access to the information, or disagree about the impact of the information then information may not have a full and immediate impact. In a rational expectations setting with better informed agents trading strategically the specialist or other traders may learn by observing trading characteristics of the transaction process. Models by Easley and O'Hara (1987) suggest that better informed agents should trade larger volume so as to capitalize on short lived information. Easley and O'Hara (1992) suggest that the timing of transactions should also influence the price process. More traders implies a higher ratio of privately informed traders in the market hence prices should adjust more quickly when transaction rates are high. Numerous other studies suggest that the specialist will widen the spread if informed trading is likely, hence wide spreads may be correlated with more rapid price adjustment¹⁰.

Empirical investigation of these theories can be carried out by including variables such as the spread and volume in the ACM model. One way of doing this is to include weakly exogenous variables in the form of predetermined variables. Hence we now consider a model that includes lagged volume and spreads. A simple model might include these lagged variables linearly in the log odds specification just as was done for the duration variables in the previous section as in:

¹⁰ See O'Hara 1995 for a very good summary of theoretical microstructure models.

$$(39) \quad h(\mathbf{p}_i) = \mathbf{m} + \sum_{j=2}^3 A_j V_{i-j}^{-1/2/} (x_{i-j} - \mathbf{p}_{i-j}) + \sum_{j=1}^3 B_j x_{i-j} + \sum_{j=1}^3 C_j h(\mathbf{p}_{i-j}) \\ + g_1 \ln(\mathbf{t}_i) + g_2 \mathbf{t}_i + g_3 \mathbf{t}_i / \mathbf{y}_i + g_4 \mathbf{y}_i + g_5 \ln(vol_{i-1}) + g_6 spd_{i-1}$$

where $\ln(vol)$ is logged volume, spd is the spread calculated as the percent difference between the bid and the ask price, and \mathbf{y} is the conditional expectation of the duration as defined in (39). V_{i-1} is the diagonal matrix of conditional variances as defined for (31) and g_5 and g_6 are response symmetric parameter vectors. The response symmetric restriction seems reasonable since we have little reason to expect that large spreads, volume or arrival rates should affect the probability of an up tick differently from the probability of a down tick¹¹.

This linear specification is simple and we might expect that we need a richer specification. For example, a large (two tick) price movement with large volume may have a different affect on the conditional distribution of the next price change than large volume with a small (1 or zero tick) price movement. Similarly, we might think the impact of a price movement on the subsequent conditional distribution of price changes might depend on the spread, or on the transaction rate¹². One way of allowing for these possibilities is to interact the volume, spread, and expected duration with the state vector. Another intuitively appealing possibility is to interact these variables with the residual $V^{-1/2/}(x_{i-j} - \mathbf{p}_{i-j})$. This way, the marginal impact of the ‘‘surprise’’ on the log odds will be a linear function of the spread, (logged) volume, and the expected duration. We restrict the parameter matrix for these vectors to be response symmetric¹³. In summary, we estimate the following ACM model:

$$(40) \quad h(\mathbf{p}_i) = \mathbf{m} + \sum_{j=2}^3 A_j V_{i-j}^{-1/2/} (x_{i-j} - \mathbf{p}_{i-j}) + \sum_{j=1}^3 B_j x_{i-j} + \sum_{j=1}^3 C_j h(\mathbf{p}_{i-j}) \\ + g_1 \ln(\mathbf{t}_i) + g_2 \mathbf{t}_i + g_3 \mathbf{t}_i / \mathbf{y}_i + g_4 \mathbf{y}_i + g_5 vol_{i-1} + g_6 spd_{i-1} \\ + (\ln(vol_{i-1})G_1 + spd_{i-1}G_2 + \mathbf{y}_{i-1}G_3) V_{i-1}^{-1/2/} (x_{i-1} - \mathbf{p}_{i-1})$$

The LM test for the addition of the interacted terms and the linear volume and spread strongly rejects the null of the model presented in section 6.2. The test statistic is 206 with 28 degrees of freedom and a p-value of .0000. The parameter estimates are presented in table 3. The coefficients on the non-interacted spreads are positive suggesting that wider spreads, all else equal, increase the probability of non-zero price movements. The coefficient on the non-interacted volume is less intuitive suggesting that,

¹¹ Diamond and Verrecchia, 1987 suggest that short selling constraints could induce a negative correlation between trading frequency and price movements. This is beyond the scope of this paper however.

¹² The microstructure models mentioned above, for example, suggest that larger spreads, larger volume transacted should be correlated with informed trading. In a rational expectations setting, the specialist will make price movements more sensitive to order flow when volume, spreads, or transaction rates are higher.

¹³ It is important at this point to recall that the symmetric response does not imply that the marginal impact of, for example, up one tick in large volume to be the same for states 1 and 5 and 2 and 4. Rather it restricts the marginal impact of, for example, a large volume 2 tick *up* price movement on the probability of a subsequent *down* 1 tick to be the same as the marginal impact of a large volume 2 tick *down* price on the probability of a subsequent *up* 1 tick.

all else equal, large lagged volume decreases the probability of a price movement. Of course this is just examining the marginal impact of the linear component of logged volume and spreads. To get a more complete picture we have to consider both the linear component as well as the interacted terms.

Before considering the full picture it is interesting to note how the probability of a price reversal is affected by the interacted terms. For each matrix G , $g_{i,1}$ and $g_{i,2}$ for $i=1,2$ gives the marginal impact of the interacted term on the probability of a price continuation. $g_{i,3}$, and $g_{i,4}$ denote the marginal impact of the interacted term on the probability of a price reversal. With the exception of a single insignificant parameter we find that shorter durations, larger volume, and wider spreads all decrease the conditional probability of a price reversal. This suggests that the price change is more likely to be permanent when transaction rates are high, spreads are wide, or volume is large. Of course this affect is only on the one step conditional distribution. To examine the long run or permanent impact we would need to consider multiple step forecasts. Analytical solutions to the impulse response functions are not available so we now consider a simple simulation study.

While it is feasible to construct forecasts of the entire price distribution we restrict our attention to the conditional mean here. In particular, we are interested in examining the expected cumulative price change following a particular sequence of price movements. Here we consider the sequence of price movements down 1 tick followed by another down 1 tick. We examine how volume, spreads, and transaction rates impact the long run expected cumulative price change of these initial two price movements.

To this end, we consider simulations where the expected duration, the spread, and the logged volume are all set to their median values. Four simulations are run. The first considers the long run impact of two consecutive down ticks when all variables are set to their median values. Simulations are then run setting each variable, one at a time, equal to its 90th percentile value for the 2 consecutive down ticks only and then back to the median value. Hence the 90th percentile values are only used for the two consecutive down ticks, not for the subsequent steps in the simulation. For the initial conditions of $h()$ and x we use the in sample values. Since we have 46,047 observations we use 46,047 iterations for each of the four simulations.

Figure 8 presents the expected cumulative price change for all four simulations. The first two price changes are always two consecutive down ticks (12.5 cents each) for a total of -25 cents. The price changes appear to stabilize quickly so we consider the cumulative price change after 50 transactions to be the long run impact. The long run price impact when spreads, volume, and durations are all set to their median values is just under 15 cents. That is the entire first tick is expected to be permanent and about 15 percent of the second price move is expected to be permanent. The long run price impact for the high transaction rate and wide spreads are slightly larger. For the large volume simulation we see that over 40 percent of the second price move is permanent in expectation.

A more convenient way to examine the results is to look at how the expected cumulative sums for the 90th percentiles differ from the cumulative sums for the median simulation. These results are presented in figure 9. We see that two consecutive down ticks when the spread is wide or transaction rates has a larger expected permanent impact

decreasing the price by .85 and .7 of a cent more respectively. 90th percentile volume has the largest expected permanent impact on the price which is 3 cents greater than the impact when all variables are set to their median values.

These simulations suggest that spreads, volume, and transaction rates can all affect the expectation of the permanent impact of a price movement. Large volume, however, has a greater impact on the expectation of the permanent impact of a price movement than large spreads or high transaction rates. Future research might consider how robust these results are to different transaction sequences as well as different parameterizations and perhaps various quantiles.

8. Conclusion

This paper views financial transactions data from the context of a marked point process. That is, traders arrive at irregular time intervals. The time of each trade has several characteristics such as volume, spreads, or transaction prices. A model is proposed for the joint distribution of arrival times and prices conditional the filtration of arrival times, prices, and potentially other weakly exogenous variables.

The majority of the price changes fall on just 5 values so discreteness is a dominant feature of the data. Decomposing the joint likelihood of arrival times into the product of the conditional distribution of price changes given arrival times and the marginal distribution of arrival times we propose a new model for discrete valued time series data. The model admits a rich dynamic structure which is necessary for the financial transactions data analyzed. The model can be viewed in the context of generalized linear models with an ARMA type structure.

Symmetry restrictions are suggested that greatly reduces the number of parameters to be estimated and give the model some intuitive properties; namely forecast distribution converges to a symmetric distribution as the forecast horizon becomes large. These results are rigorously proved for the linear ACM model while simulations and our intuition suggest these results must also hold for the nonlinear logistic models estimated here. We continue to pursue these results identity as well as the asymptotic properties of the estimator for link functions other than the identity link.

Maximum likelihood estimates given for several models for IBM transactions data. A simple ACM(3,3,3) model suffices based on LM tests. More interestingly, models for the joint distribution of arrival times and price changes suggest that the transaction price variance is small for the shortest and longest durations between trades. We also find that the variance of the transaction price is negatively related to the expected waiting time. This is consistent with predictions from Easley and O'Hara (1992) where active markets are indicative of a larger than normal fraction of informed traders.

A model that includes volume and spreads is also considered. We find that the probability of price moves increases as the spread widens. Simulations suggest that the full affect of a transaction is not realized for many trades. As an example we ask "What is the long run impact of two consecutive transactions that move the price down one tick each?" We find that while spreads, and trading rates can affect the expected long run

impact volume appears to be the most important. We view these simulation results as a possible starting point for more robust studies.

Appendix

Proof of lemma 1: All probabilities will be non-negative under condition a) since they will be the sum of three non-negative terms. The omitted state will have positive probability if the sum of the π in (11) is less than unity. This also insures that each element of π is less than unity. The column with the greatest sum gives the maximum that Ax can be. The weighted average of the columns of $C-A$ will be less than the maximum column. If these two numbers plus the sum of μ is less than 1 this is sufficient that the probability of the omitted state is non-negative.

Proof of Theorem 1

We have the linear ACM(1,1,1) model defined as:

$$(*) \quad \mathbf{p}_i = A(x_{i-1} - \mathbf{p}_{i-1}) + Bx_{i-1} + C\mathbf{p}_{i-1} + \mathbf{m}$$

For the K state system let the $(K-1)$ dimension vector $\bar{\mathbf{p}} = E(x)$. Then taking expectations on both sides of equation (*) and rearranging terms yields

$$(1') \quad (\mathbf{I} - (B + C))\bar{\mathbf{p}} = \mathbf{m}$$

Premultiplying both sides of (1') by Q yields

$$(2') \quad \mathbf{Q}(\mathbf{I} - (B + C))\bar{\mathbf{p}} = \mathbf{Q}\mathbf{m}$$

It is easily verified that if B and C are both response symmetric then $(B+C)$ is response symmetric. Since $(B+C)$ is response symmetric and \mathbf{m} is symmetric

$$(3') \quad (\mathbf{I} - (B + C))\mathbf{Q}\bar{\mathbf{p}} = \mathbf{m}$$

follows from (2').

Since $(\mathbf{I} - (B + C))$ is of full rank equations (1') and (3') imply that $\mathbf{Q}\bar{\pi} = \bar{\pi}$.

Hence, $\bar{\pi}$ is symmetric.

If all the eigenvalues of $(B+C)$ lie in the unit circle then the usual dynamic analysis implies that

$$\lim_{k \rightarrow \infty} \mathbf{p}_{i+k} \rightarrow \bar{\mathbf{p}} = (\mathbf{I} - (B + C))^{-1} \mathbf{m}$$

QED

References

1. Admati, Anat R. and Paul Pfleiderer, 1988, A theory of Intraday Patterns: Volume and Price Variability, *The Review of Financial Studies* 1, 3-40.
2. Berndt, E., B. Hall, R. Hall, J. Hausman,(1974), Estimation and Inference in Nonlinear Structural Models, *Annals of Economic and Social Measurement*, 3,pp 653-665
3. Cox, D. R., 1970 *The Analysis of Binary Data*. London: Chapman and Hall
4. Cox, D. R., 1981 *Statistical Analysis of Time Series: Some Recent Developments*. Scandinavian Journal of Statistics
5. Easley, D., O'Hara, M., (1987), Price, Trade Size, and Information in Securities Markets. *Journal of Financial Economics* 19, 69-90
6. Easley and O'Hara, 1992, Time and the Process of Security Price Adjustment. *The Journal of Finance* 47, 69-90
7. Engle, Robert 1996, *The Econometrics of Ultra-High Frequency Data*, University of California, San Diego unpublished manuscript
8. Engle, Robert, D. Hendry, and Richard 1983, Exogeneity *Econometrica*
9. Engle, Robert and J. Russell, 1997, Forecasting the Frequency of Changes in Quoted Foreign Exchange Prices with the Autoregressive Conditional Duration Model, *Journal of Empirical Finance*
10. Engle, Robert and J. Russell, 1998, Autoregressive Conditional Duration: A New Model for Irregularly Spaced Data, Forthcoming in *Econometrica*
11. Engle, Robert and J. Russell, 1995, Autoregressive Conditional Duration: A New Model for Irregularly Spaced Data, University of California, San Diego Working Paper Series
12. Hasbrouck, J., 1991, Measuring the Information Content of Stock Trades, *The Journal of Finance* 46, 1, 179-207
13. Hasbrouck, J., Analysis of Transaction Price Data, Forthcoming in *The Handbook of Statistics*
14. Hausman, J., A. Lo, and C. MacKinlay, 1992, An Ordered Probit Analysis of Transaction Stock Prices, *Journal of Financial Economics*

15. Jones, C., Kaul, G., Lipson, M., 1994, Transactions, volume and volatility. *Review of Financial Studies* 7, 631-651
16. Kyle, Albert, 1985, Continuous Time Auctions and Insider Trading, *Econometrica* 53, 1315-1336
17. Kalbfleisch, J., and R. Prentice, 1980, *The Statistical Analysis of Failure Time Data*, John Wiley & Sons.
18. Lee, C., and M. Ready, 1991, *The Journal of Finance*, V46 733-746
19. Lancaster, T., 1990, *The Econometric Analysis of Transition Data* Cambridge University Press
20. MacDonald, I., Zucchini, Walter, 1997, *Hidden Markov and Other Models for Discrete-valued Time Series*. Chapman & Hall
21. McNish, T.H., Wood, R.A., 1991, Hourly Returns, Volume, Trade Size and Number of Trades. *Journal of Financial Research* 1, 458-491
22. O'Hara, M., 1995, *Market Microstructure Theory*, Basil Blackwell Inc.
23. Russell, J. 1996, *Econometric Analysis of High Frequency Transactions Data Using a New Class of Accelerated Failure Time Models with Applications to Financial Transaction Data*, Dissertation University of California, San Diego.
24. Shephard, Niel, 1995, "Generalized Linear Autoregressions", unpublished manuscript, Nuffield College, Oxford
25. Tiao, G and Box, G, 1981, "Modeling Multiple Time Series with Applications", *Journal of the American Statistical Association* 76
26. Zegar, Scott, and B. Qaqish, 1988, "Markov Regression Models for Time Series: A Quasi Likelihood Approach" *Biometrics*, 44 December

Table 1: Parameter estimates for ACM(3,3,3)-ACD(2,2) model

		State 1&5		State 2&4				State 1&5		State 2&4	
A_1	μ_1	-0.04086 (-4.59)		μ_2	-0.09232 (-4.40)	B_1	β_{11}	0.418793 (10.32)	β_{21}	0.111428 (3.06)	
	α_{11}	-2.2074 (-8.62)		α_{21}	-1.93178 (-10.80)		β_{12}	0.394214 (6.73)	β_{22}	0.303763 (8.18)	
	α_{12}	-1.76827 (-8.50)		α_{22}	-2.24194 (-22.04)		β_{13}	0.010252 (0.32)	β_{23}	-0.44707 (-19.98)	
	α_{13}	2.216657 (17.18)		α_{23}	3.08897 (46.33)		β_{14}	0.250657 (10.75)	β_{24}	0.019691 (0.88)	
A_2	α_{14}	3.557097 (19.17)		α_{24}	2.122032 (15.15)	B_2	β_{11}	-0.68123 (-9.97)	β_{21}	-0.18547 (-3.34)	
	α_{11}	1.897501 (4.17)		α_{21}	1.06814 (3.19)		β_{12}	-0.3261 (-3.21)	β_{22}	-0.17269 (-2.62)	
	α_{12}	0.77655 (2.13)		α_{22}	1.083212 (3.96)		β_{13}	-0.15693 (-2.90)	β_{23}	0.399623 (8.72)	
	α_{13}	-1.98998 (-8.60)		α_{23}	-2.40172 (-9.56)		β_{14}	-0.31318 (-6.51)	β_{24}	-0.01818 (-0.54)	
A_3	α_{14}	-3.61868 (-12.54)		α_{24}	-1.35097 (-5.93)	B_3	β_{11}	0.273505 (7.56)	β_{21}	0.074599 (3.01)	
	α_{11}	0.29167 (1.08)		α_{21}	0.465187 (2.50)		β_{12}	-0.0405 (-0.77)	β_{22}	-0.0657 (-1.76)	
	α_{12}	0.916703 (4.37)		α_{22}	0.71281 (3.70)		β_{13}	0.159475 (4.68)	β_{23}	-0.02221 (-0.70)	
	α_{13}	-0.19488 (1.16)		α_{23}	-0.16672 (-0.85)		β_{14}	0.07612 (2.26)	β_{24}	0.002392 (0.14)	
	α_{14}	0.203016 (1.06)		α_{24}	-0.3136 (-2.23)	C_1	χ_{11}	1.647887 (26.14)	χ_{22}	1.453127 (16.56)	
						C_2	χ_{11}	-0.66094 (-6.80)	χ_{22}	-0.38867 (-2.82)	
						C_3	χ_{11}	0.002585 (.066)	χ_{22}	-0.11071 (-1.84)	

Where
$$h_i(\mathbf{p}_i) = \mathbf{m} + \sum_{j=1}^p A_j V_{i-j}^{-1/2} (x_{i-j} - \mathbf{p}_{i-j}) + \sum_{j=1}^q B_j x_{i-j} + \sum_{j=1}^r C_j h(\mathbf{p}_{i-j})$$

$$\mathbf{m} = \begin{bmatrix} m_1 \\ m_2 \\ m \\ m_2 \end{bmatrix} \quad A = \begin{bmatrix} a_{1,1} & a_{1,2} & a_{1,4} & a_{1,5} \\ a_{2,1} & a_{2,2} & a_{2,4} & a_{2,5} \\ a_{25} & a_{2,4} & a_{2,2} & a_{2,1} \\ a_{1,5} & a_{1,4} & a_{1,2} & a_{1,1} \end{bmatrix} \quad B = \begin{bmatrix} b_{1,1} & b_{1,2} & b_{1,4} & b_{1,5} \\ b_{2,1} & b_{2,2} & b_{2,4} & b_{2,5} \\ b_{2,5} & b_{2,4} & b_{2,2} & b_{2,1} \\ b_{1,5} & b_{1,4} & b_{1,2} & b_{1,1} \end{bmatrix} \quad C = \begin{bmatrix} c_{1,1} & 0 & 0 & 0 \\ 0 & c_{2,2} & 0 & 0 \\ 0 & 0 & c_{2,2} & 0 \\ 0 & 0 & 0 & c_{1,1} \end{bmatrix}$$

Table 2: Parameter Estimates for ACM(3,3,3)-ACD(2,2) with Duration Dependence
 (Only the Duration parameters entering the ACM model and ACD parameters are presented here)

ACM Duration Parameters			ACD(2,2) Parameters	
Variable	state 1 and 5	state 2 and 4	Parameter	Estimate
Log(τ_i)	-0.02473 (-3.634)	0.11942 (2.04)	ω	0.002322 (5.69)
τ_i	-0.00964 (-1.67)	-0.01917 (-1.67)	α_1	0.08963 (21.98)
Log(τ_i/ψ_i)	0.023822 (2.84)	0.10008 (1.68)	α_2	-0.06416 (-15.92)
ψ_i	0.023775 (2.54)	-0.14384 (-2.45)	β_1	1.46603 (23.41)
			β_2	-0.493341 (-8.40)

Table 3: Parameter Estimates for ACM(3,3,3)-ACD(2,2) with Duration Dependence, Volume and Spreads*

		State 1&5		State 2&4				State 1&5		State 2&4	
A_1	μ_1	-1.80653 (-6.66)	μ_2	-0.46063 (-4.51)	B_1	β_{11}	-0.0167 (-.11)	β_{21}	0.085262 (.85)		
	α_{11}	-1.70949 (-5.55)	α_{21}	-1.84436 (-7.65)		β_{12}	-0.27813 (-1.54)	β_{22}	-0.07619 (-.86)		
	α_{12}	-1.77127 (-6.36)	α_{22}	-2.00716 (17.45)		β_{13}	0.659049 (8.31)	β_{23}	0.313396 (6.91)		
	α_{13}	1.559506 (11.03)	α_{23}	2.17318 (30.04)		β_{14}	0.464961 (7.65)	β_{24}	0.197755 (2.98)		
A_2	α_{14}	3.556916 (17.71)	α_{24}	1.785919 (12.28)	B_2	β_{11}	-0.07464 (-1.93)	β_{21}	-0.0414 (-1.96)		
	α_{11}	-1.09218 (-4.25)	α_{21}	-0.94197 (-6.43)		β_{12}	0.137849 (2.09)	β_{22}	0.102307 (3.29)		
	α_{12}	-1.08803 (-4.88)	α_{22}	-1.06685 (-10.01)		β_{13}	0.027287 (.49)	β_{23}	0.003749 (.14)		
	α_{13}	0.329903 (1.72)	α_{23}	0.880543 (9.33)		β_{14}	-0.08102 (-2.38)	β_{24}	0.019464 (1.14)		
A_3	α_{14}	1.404484 (5.01)	α_{24}	1.261819 (9.86)	B_3	β_{11}	-0.09089 (-2.87)	β_{21}	-0.06459 (-3.62)		
	α_{11}	0.29725 (1.37)	α_{21}	-0.02613 (-.22)		β_{12}	0.077485 (1.42)	β_{22}	-0.02559 (-.92)		
	α_{12}	-0.18437 (-.97)	α_{22}	-0.10285 (-1.06)		β_{13}	-0.00507 (-.10)	β_{23}	0.012738 (.53)		
	α_{13}	0.089902 (.53)	α_{23}	0.315516 (3.89)		β_{14}	-0.0055 (-.16)	β_{24}	0.009082 (.50)		
G_{dur}	α_{14}	0.379977 (1.66)	α_{24}	0.270749 (2.40)	C_1	χ_{11}	0.354822 (.35)	χ_{22}	0.151845 (5.62)		
	γ_{11}	-0.09719 (-1.18)	γ_{21}	-0.13817 (-2.36)		C_2	χ_{11}	0.122868 (.12)	χ_{22}	0.120893 (4.58)	
	γ_{12}	0.096407 (1.60)	γ_{22}	-0.19138 (-5.29)			C_3	χ_{11}	0.048712 (.04)	χ_{22}	0.045823 (2.18)
	γ_{13}	0.148172 (4.35)	γ_{23}	0.038825 (2.06)	g_1	$\text{Log}(\tau_i)$		-0.57092 (-5.03)	$\text{Log}(\tau_i)$	0.141683 (2.35)	
γ_{14}	0.125698 (4.40)	γ_{24}	0.136902 (4.41)	τ_i		0.02548 (.89)	τ_i	-0.02125 (-1.81)			
G_{vol}	γ_{11}	0.047691 (2.80)	γ_{21}	0.012134 (.91)		$\text{Log}(\tau_i/\psi_i)$	0.722964 (6.08)	$\text{Log}(\tau_i/\psi_i)$	0.101407 (1.64)		
	γ_{12}	0.087872 (4.19)	γ_{22}	0.061953 (7.04)		ψ_i	0.102405 (.82)	ψ_i	-0.23997 (-3.89)		
	γ_{13}	-0.08392 (-9.02)	γ_{23}	-0.07629 (-14.84)	Vol_{i-1}	-0.06329 (-3.57)	Vol_{i-1}	-0.06221 (-8.00)			
	γ_{14}	-0.04223 (-6.00)	γ_{24}	-0.03458 (-4.41)	Spd_i	0.468306 (6.99)	Spd_i	0.20852 (5.75)			
G_{spd}	γ_{11}	0.045999 (.49)	γ_{21}	0.127581 (1.68)	g_6						
	γ_{12}	0.016184 (.21)	γ_{22}	0.223685 (5.33)							
	γ_{13}	-0.05747 (-1.23)	γ_{23}	-0.04436 (-1.49)							
	γ_{14}	-0.11685 (-3.31)	γ_{24}	-0.08697 (-2.05)							

$$\begin{aligned}
 h(\mathbf{p}_i) = & A \sum_{j=2}^3 A_j V_{i-j}^{-1/2} (x_{i-j} - \mathbf{p}_{i-j}) + \sum_{j=1}^3 B_j x_{i-j} + \sum_{j=1}^3 C_j h(\mathbf{p}_{i-j}) \\
 & + g_1 \ln(t_i) + g_2 t_i + g_3 t_i / y_i + g_4 y_i + g_5 \text{vol}_{i-1} + g_6 \text{spd}_{i-1} \\
 & + (\ln(\text{vol}_{i-1}) G_1 + \text{spd}_{i-1} G_2 + y_{i-1} G_3) V_{i-j}^{-1/2} (x_{i-1} - \mathbf{p}_{i-1})
 \end{aligned}$$

* The ACD(2,2) parameters are very similar to those presented in table 2 and are therefore not presented here.

Figure 1: Histogram of Transaction Prices




Figure 4: Distribution of Price Changes as a Function of Duration




Figure 5. Distribution of Price Changes as a Function of Expected Duration




Figure 6. Variance of Price Distribution as a Function of Duration




Figure 7. Variance of Price Distribution as a Function of Expected Duration




Figure 8. Expected Cumulative Price Change following Two Sequential Down Ticks




Figure 9. Expected Difference from Median Cumulative Price Change Following Two Consecutive Down Ticks.

