Modelling the Time Between Trades in the After-Hours Electronic Equity Futures Market.*

Mardi Dungey*+ Nagaratnam Jeyasreedharan* and Tuo Li*
* University of Tasmania
+ CFAP, University of Cambridge

November-2010

Abstract

Equity futures markets now have extensive after-hours trading periods via electronic exchanges. This paper models the time between trades on the Globex platform for NASDAQ and S&P500 equity futures contracts. Unlike many spot markets, these contracts support autoregressive conditional duration models with persistence and threshold effects. The S&P500 threshold model displays typical low autocorrelation and strong clustering for short durations, but greater autocorrelation and low clustering for large duration transactions. Volume is an informative mark in both cases, with increased volume reducing the time before the next trade.

JEL Categories: G12, C22, C41, C52

Keywords: duration, high frequency data, electronic futures markets

*We are grateful for comments and advice from Olan Henry and Susan Thorp. Support for ARC Discovery Grant DP0664024 and a Faculty of Business, University of Tasmania Research Scholarship are gratefully acknowledged. Author contacts: Dungey; Mardi.Dungey@utas.edu.au, Jeyasreedharan; NJ.Sreedharan@utas.edu.au, Li; Tuo.Li@utas.edu.au.

1 Introduction

The trading environment in financial markets has changed rapidly in the past 10 years. Many instruments are increasingly traded on electronic exchanges and trading hours are extending beyond standard business hours. A particularly successful example of these innovations is the trade in equity futures contracts on the GLOBEX exchange. Equity futures contracts which trade on the open outcry Chicago Mercantile Exchange (CME) pit are now also generally available outside pit hours on the electronic market. Since 1993 the standard contract for the S&P500 has been available in this format, followed in the mid-1990s by the NASDAQ contract, and growth in volume has been relatively rapid. However, to date, the behavior of the after-hours market has been relatively little studied; Coppejans and Domowitz compare the electronic and open outcry markets and Dungey, Fakhrutdinova and Goodhart (2009) explore the volume and volatility characteristics of the NASDAQ and S&P500 futures contracts.

This paper makes three contributions. First, it models trade duration, that is the time between trades, in the out of hours equity futures markets for the NASDAQ and S&P500 indices. The time between trades provides information to the market, indicating the presence of news and potentially in the absence of trade that there is no new information, see Easley and O'Hara (1992). Trade duration has not previously been modelled for this market. Because the market is after-hours it has a peculiarly marked diurnal pattern, with relatively intense trade in the period immediately following the close of the open outcry market, lower volume and intensity in the Asian trading zone, an increase in activity and intensity in European trading hours and a dramatic increase in both intensity and volume immediately prior to the opening of the pit. Volume increases particularly with the 8:30am EST scheduled macroeconomic news announcement period in the US. Modelling trade duration in this market is thus a completely different proposition from previous empirical work on duration modelling, which typically involves spot equity market contracts; for example Engle and Russell (1998), Zhang, Russell and Tsay (2001).

Second, the data sample of this paper covers two years, a dramatic increase on the usual 3 month sample analyzed in existing papers on time between trades. A particular challenge is to fit a consistent model to this length of sample - given that Zhang, Russell and Tsay (2001) find evidence for 7 structural breaks in a 3 month data set. The final contribution is to include volume of trade as an additional mark in the modelling process, which makes a small, but significant, negative contribution to conditional duration. That is, an observed larger trade volume results in a smaller time to the next trade - which may be interpreted as either due to the arrival of public information resulting in market participants making portfolio adjustments, or alternatively in the absence of public information, that when market participants observe a high volume trade they interpret this as private information which encourages them to trade, thus increasing trade intensity.

The modelling framework of the paper is based on the ACD models proposed by Russell and Engle (1998) and subsequent extensions. ACD models account specifically for the observed serial correlation and clustering in trade duration, and are closely related in form to the GARCH framework. Like GARCH, the preferred lag structure in most applications strongly suggests an ACD(1,1) starting point, although various alternatives exist for the assumed error distribution; including the exponential, Weibull, generalized Gamma, Burr, generalized F and mixtures of distributions; see Russell and Engle (1998), Lunde (1999), Gramming and Maurer (2000), Hautsch (2002) and De Luca and Gallo (2004). The markets explored here require both an extension of the lag structure and accounting for non-linearities through a two regime threshold ACD model. Specifically, the trade duration in the NASDAQ futures data incorporates higher order lags, while the more intensely traded S&P500 contract is more effectively modelled with a threshold model, featuring differing levels of higher order lags for large duration observations.

The paper proceeds as follows. Section 2 provides a brief overview of the after-hours electronic equity futures market for the NASDAQ and S&P500 contracts, followed by the description of the sample period in Section 3. The ACD framework is outlined in Section 4. Section 5 documents the development of the final model via benchmark ACD(1,1) models, extensions to the lag order, the introduction of volume and threshold models. Section 6 concludes.

2 The After-Hours Electronic Equity Futures Market

The standard equity futures contracts for the NASDAQ and S&P500 traded on the CME are contracts for $250 times the equity index price with 0.10 ticks. Both contracts trade in the CME open outcry pit between the hours of 8:30CST to 15:15 CST and on the electronic GLOBEX exchange after-hours. The after-hours trading period currently begins at 17:00pm CST on Sunday evenings (corresponding to the opening of trade in the Japanese trading day) and continues until 8:15 CST Monday morning. For the remainder of the working week the contract begins trade at 15:30CST after the closure of the pit, and continues to trade until 8:15CST the next morning, with the exception of Fridays where there is no electronic trade following the closure of the open outcry pit on Friday afternoon. The electronic exchange closes for maintenance everyday between 16:30CST and 17:00CST, and on public holidays trades reduced hours.

There is no overlap in trade of the open outcry pit and the electronic trading of this contract. The two platforms trade the same product, thus making it possible for market participants to change their portfolio holdings in these indices almost 24 hours per day. Although there is no electronic trading in the standard contract during the open-outcry market, E-mini contracts which are one-fifth of the standard contract size and only available electronically do trade 24 hours (other than the half-hour shutdown for maintenance).[1]

Total volume accounted for by electronic trade in this market has been growing rapidly in recent years; Figure 1 shows that total volume traded in the electronic market has grown from 200 million in 2002 to more than 2 billion in 2007, although this includes the consolidation of the CME and CBOT trades into the total volume in 2007. The sample period in this period covers the years 2004 to 2006, and hence excludes this structural change.

It is not at first evident how the 3 contract forms (standard futures, electronic, E-mini) for the same instrument coexist. However, the standard contract trades

[1] E-minis were introduced for the S&P500 in 1997 and for the NASDAQ in 1999. These two types of contracts soon become the two fastest-growing products in CME history; see CME group website.http://www.cmegroup.com/globex/resources/history-of-globex.html

electronically only when the pit floor is closed and is five times larger than the E-mini product which trades for virtually 24 hours. The smaller contract is designed to appeal to retail investors. Additionally, trading on the electronic platforms is more expensive than trading in the pit via transaction fees, although precise details of the transaction fees vary by market participant and are not readily publicly available. As the E-minis trade during the pit period of 8:30CST to 15:15CST Hasbrouck (2003) and Coppejans and Domowitz (1999) have compared the relative efficiency of the E-mini and open outcry market - finding that the open outcry market is more efficient at absorbing local information. However, this comparison is made more difficult by the difference in the size and transaction fees of the contracts. Trading in the pit and on the electronic platform for the standard contract do not overlap - rather in combination they complete the trading day, so their relative efficiency can not be easily compared.

Dungey, Fakhrutdinova and Goodhart (2009) describe volume and price impact for the after-hours standard equity futures contracts for the S&P500 and NASDAQ indices. They find that the period of highest average volume in the day occurs immediately prior to the opening of the open outcry pit, peaking around 7:30CST, which corresponds to the time of prescheduled macroeconomic news releases in the US at 8:30EST. They find that price impact for the S&P500 contracts is lowest in the high volume period immediately prior to the opening of the open outcry pit, and higher in general during the European and Asian trading hours, but for the NASDAQ is highest immediately post-close of the open outcry market. This may suggest that the relatively low volume traded on the NASDAQ compared with the S&P500 has made the gains from anonymous electronic trading lower than those for the highly liquid S&P500, reducing the attractiveness of trade in the post-close period for this instrument.

3 The Data Sample

Information on the transactions on the GLOBEX electronic exchange for the NASDAQ and S&P500 futures contracts were obtained from the CME for the period from July 1, 2004 to September 30, 2006. The data comprise 213,332 tick observations for the NASDAQ and 1,053,524 for the S&P500. After cleaning the data set to remove negative durations and aggregating volumes for transactions with the same time stamp to be treated as a single transaction, following Engle and Russell (1998), the sample data comprises 149,314 observations on the NASDAQ and 684,010 observations in the S&P500. The data display a distinct diurnal pattern, and it is customary in this literature to remove this prior to estimation. Using a piecewise linear spline with 17 knots representing hourly intervals during the after-hours trade period covered the data are diurnalised using a multiplicative specification of the diurnality, in a manner similar to that proposed in Engle and Russell (1998).

Table 1 contains basic descriptive statistics of the diurnally adjusted data sample. In both indices there is evidence of relatively large higher order moments, strongly rejecting normality. Figures 2 and 3 show the average adjusted daily duration and volume pattern for the NASDAQ and S&P500 data beginning from midnight CST each day. Trade at midnight CST is equivalent to the Asian trading day, and the durations are relatively high. Duration then decreases until 8:15 CST when the GLOBEX market ceases shortly before the open of the pit trading session. During the morning electronic trade duration drops first during the European trading day and most dramatically around the 7:30 CST period (corresponding to the usual announcement time for pre scheduled US macroeconomic news). As discussed in the previous section, diurnal volume in these markets peaks at this time.

Immediately following the closure of the floor market at 15:15 CST trading is relatively intense in the electronic market, and volume is again relatively high. Dungey, Fakhrutdinova and Goodhart (2009) associate this higher trading volume with a desire on the part of market participants to settle their end of day positions in the anonymity of the electronic market as opposed to the open outcry pit, despite the higher costs of trading the same contract on the electronic market. After this point trade duration begins to climb again as the market becomes less active entering the Asian trading zone. Overall, the figures indicate the existence of a negative relationship between volume and duration. This feature will be incorporated into the formal model of duration in Section 5.4.

4 ACD Models

Define the (irregular) time between consecutive trades in a single market as $x_i = t_i - t_{i-1}$, where t_i represents the time of the current trade and t_{i-1} is the immediately previous trade. Assuming that the trade duration, x_i, evolves according to the process

$$x_i = \psi_i \varepsilon_i, \tag{1}$$

where $\psi_i \equiv E(x_i | x_{i-1,} \ldots, x_0)$ represents conditional expected duration and ε_i is an error process, the autoregressive and clustering aspects of duration are captured through specification of the conditional expected duration as

$$\psi_i = \omega + \sum_{j=0}^{p} \gamma_j x_{i-j} + \sum_{k=0}^{q} \omega_k \psi_{i-k}, \tag{2}$$

where ω, γ_j and ω_k are parameters, and p and q represent the lag orders, denoted as an ACD(p, q), see Engle and Russell (1998).

A number of alternatives have been considered for the error distribution ε_i, including the Exponential (EACD), Weibull (WACD), generalized Gamma distribution (GACD), Burr and generalized F; see Engle and Russell (1998), Lunde (1999) Gramming and Maurer (2000) and Hautsch (2002). De Luca and Gallo (2004) use a mixture of two distributions.

This paper concentrates on comparisons of the EACD, WACD and GACD forms of the model. In each case the duration, x_i, is restricted to be non-negative The probability density function

$$f(x) = \frac{\alpha}{\beta^{\alpha} \Gamma(\kappa)} x^{\kappa\alpha - 1} e^{(-x/\beta)^{\alpha}}, \tag{3}$$

represents the generalized Gamma distribution with two shape parameters, α and κ and scale parameter β, which in the case of $\kappa = 1$ is equivalent to the Weibull distribution and when $\alpha = \kappa = 1$ is the exponential distribution. Each of these functions possesses high concentration at shorter durations and a long right tail for longer durations.

A number of alternative specifications to the conditional duration given in equation (2) also exist. Expressing equation (2) in log form rules out negative durations which have occurred in other applications with the addition of further explanatory variables to the conditional duration model; Bauwens and Giot (2000), but are not an issue in the current application. Jasiak (1998) introduced the fractionally integrated ACD model (FIACD), to account for long memory, while Zhang, Russell and Tsay (2001) introduced the threshold ACD model (TACD), where different conditional means, error distributions and persistence may occur in each regime. In a two regime threshold model the conditional duration equation (2) is replaced by

$$\psi_i = \begin{cases} \omega^{(1)} + \sum_{j=1}^{p_1} \gamma_j^{(1)} x_{i-j} + \sum_{k=1}^{q_1} \omega_k^{(1)} \psi_{i-k}, \text{ if } 0 < x_i \leq r_1 \\ \omega^{(2)} + \sum_{j=1}^{p_2} \gamma_j^{(2)} x_{i-j} + \sum_{k=1}^{q_2} \omega_k^{(2)} \psi_{i-k}, \text{ if } r_1 < x_i < \infty \end{cases} \tag{4}$$

which is notated as TACD$(p_1, q_1 : p_{2,}q_2)$ where p_1 and q_1 represent lag orders in the first regime, and p_2 and q_2 represent lag orders in the second regime and r_1 is some exogenously chosen cut off point delineating the regimes. Other recent alternatives include Markov Switching ACD models, as in Hujer et al (2002); mixtures of distributions applied to price durations in De Luca and Gallo (2004) and trade durations in Hujer and Vuletić (2007), stochastic volatility duration models such as Ghysels, Gourieroux and Jasiak (2004) and the simultaneous modelling of price and trade duration in Engle and Russell (2005).

The next section presents the results of applying the ACD model with varying error assumptions and threshold ACD specifications to the NASDAQ and S&P500 equities futures data. Parameter estimates are undertaken using maximum likelihood based on the log-likelihood functions for the individual models using RATS version 7.

5 Empirical Results

The majority of the existing literature has fitted ACD$(1, 1)$ models with alternative distributional assumptions. EACD(1,1), WACD(1,1) and GACD(1,1) models are fitted to the two data series in the next section, followed by extensions to higher lag orders and then the potential role of volume traded in providing further information. Finally, evidence of non-linearity in the S&P500 results lead to the estimation of a threshold ACD model for this data.

5.1 ACD(1 1) specifications

Consider first the results for estimates of the EACD(1,1), WACD(1,1) and GACD(1,1) models for the NASDAQ data reported in Table 2. The Ljung-Box statistics for each model are relatively high, ranging between 280 and 372 for the Q(20) statistic , although this reflects the large sample size in addition to potential problems with the fit of the model. The parameter estimates in the GACD(1,1) and WACD(1,1) also provide some evidence as to which model best describes the data. There is considerably more variation in the parameter estimates for autocorrelation and clustering across the specifications than obtained by De Luca and Gallo (2004) in their comparison of ACD(1,1) models for price durations across different distributional assumptions. The parameter estimates for κ and α reported in the final column of Table 2 do not support the EACD ($\alpha = \kappa = 1$) or WACD ($\kappa = 1$) specification..

The parameter values themselves support a relatively low autocorrelation component to the conditional duration equation, with γ_1 less than 0.22. The clustering component, given by the parameter ω_1 is stronger at around 0.8 in each estimation. The general form of low autocorrelation and high clustering parameter estimates are common to existing literature estimating ACD models for IBM equities in Engle and Russell (1998) and Disney stocks in Hautsch (2006). The shape parameter κ, from the GACD(1,1) estimation supports a mixture of greater than 1 Weibull distributions, while the α parameter suggests a smaller influence from the exponential distributions. Thus far, the results for the NASDAQ data support a GACD(1,1) specification on the basis of the non-unit values of α and κ, although measures of fit suggest that a less complex distributional assumption provides a slightly better fit to the data.

The S&P500 data has a far greater trading intensity than the NASDAQ data as described in Section 3, and the Ljung-Box coefficients are an order of magnitude higher than those reported for the NASDAQ. The estimated value of α in the WACD specification rejects the null hypothesis of $\alpha = 1$, which would support an EACD specification. In this case the GACD(1,1) model failed to converge, producing extremely high estimates of κ, suggesting that there are problems remaining with the specification. The next section explores generalizations of these baseline specifications to examine the most likely means of improving the models.

5.2 Higher order lags

Although many applications do find that ACD(1,1) models with varying distributional assumptions provide the best characterizations of their data, a small number of papers have favoured higher order lag lengths, (Engle and Russell, 2005; Zhang et al., 2001). To explore the WACD and GACD specifications for the NASDAQ and the WACD specification for the S&P500 are considered with extended lag lengths. A similar process is not applied to the EACD models as none of the more general specifications reported in Section 5.1 support an exponential distributional assumption.

The best results for the NASDAQ are a WACD(5,5) and GACD(3,3) which are reported in Table 3. It is evident that the WACD(5,5) has reduced the Ljung-Box statistics considerably over the results reported in Table 2, and the sum of the estimated coefficients, $\sum_{j=1}^{5}\left(\gamma_j+\omega_j\right)\approx 0.9997$, indicates persistence in the adjusted durations. The unconditional mean adjusted duration for this specification is given by $E(\psi_i)=\omega/\left(1-\sum_{j=1}^{5}\left(\gamma_j+\omega_j\right)\right)\approx 3.0814$ seconds. It is notable that there is a drop in the value of the estimate of ω by two orders of magnitude compared with the WACD(1,1) specification from Table 2, but the shape parameter, α is unchanged to two decimal places.

The GACD(3,3) specification contains some problematic outcomes. The Ljung-Box statistics are not reduced over the GACD(1,1) specification, and importantly the sum of the ω_j and λ_j parameters, $\sum_{j=1}^{3}\left(\gamma_j+\omega_j\right)\approx 1.0000$, a case not encompassed by the GACD specification. The shape parameter values for α and κ are not greatly changed from the GACD(1,1) specification. Of the two longer lag lengths investigated for the NASDAQ model the WACD(5,5) seems the more satisfactory.

Specifications incrementing the lag lengths in the S&P500 WACD(1,1) model fail to converge providing further evidence of the difficulties in fitting the S&P500 data.

5.3 The role of volume

As lag length adjustments have not made a substantial improvement to the model specifications, this section turns to the inclusion of other marks in the process; specifically, whether volume transacted has any extra information over the simple duration information. Bauwens and Veredas (2004) documented evidence of a significant relationship, but were restricted to daily volume proxies in their analysis. A further stream of literature considers the combination of trade time and price durations, but given the difficulties with the unsigned price data in this sample, which introduces problems of bid-ask bounce requiring an approximating algorithm and associated uncertainty, this is left for future work.

Figures 2 and 3 suggest a negative relationship between volume and trade duration, an increase in volume transacted is associated with a decrease in trade duration, consistent with trade volume possessing information in this market, and that lack of trade indicates a lack of new information. The conditional duration equation (2) is augmented with the transacted volume information using the WACD(1,1) models reported in Table 2 as the baseline models.

Table 4 reports the results for the WACD(1,1) models for the NASDAQ and S&P500 datasets augmented with volume information. In each case the volume parameter is negative and statistically significant at the 1% level. This result is consistent with the hypothesis that higher volume transacted indicates some form of information entering the market and shortening trade durations. There are two possible mechanisms for this outcome. In the first case public information may be causing market participants to reassess their positions and increasing trade intensity. In the second case, market participants observe increased trade volume and interpret it as an indicator of private information, and are hence encouraged to trade themselves, thus increasing trade intensities. Comparing the results with those reported in Table 2 there are few changes in the other parameter estimates. In particular, the shape parameter α, is little changed in either case. However, the Ljung-Box statistics have been improved by the inclusion of the additional volume mark.

5.4 Threshold effects

While the NASDAQ data has been modelled in a way which may be considered acceptable, there remain considerable problems with the S&P500 data. As shown in Section 3 there are some indications of different tail behaviors for large durations. Accounting for the possibility that these larger durations behave significantly differently to the bulk of the durations through a threshold model can significantly improve the model estimates; for example Zhang, Russell and Tsay (2001) for 3 months worth of IBM data examined in Engle and Russell (1998) by introducing non-linearities.

Table 5 reports the parameter estimates for a two regime threshold model with Weibull distribution TWACD(4,1:4,1) including the volume mark process as a further explanatory variable. That is, the complete model estimated is:

$$x_i = \psi_i \varepsilon_i \tag{5}$$

$$\psi_i = \begin{cases} \omega^{(1)} + \sum_{j=1}^{4} \gamma_j^{(1)} x_{i-1} + \omega_1^{(1)} \psi_{i-k} + v_i^{(1)}, & \text{if } 0 < x_i \leq r_1 \\ \omega^{(2)} + \sum_{j=1}^{4} \gamma_j^{(2)} x_{i-j} + \omega_1^{(2)} \psi_{i-k} + v_i^{(2)}, & \text{if } r_1 < x_i < \infty \end{cases} \tag{6}$$

where the regime cutoff, r_1 is set to be 19 seconds.[2] Note that this is a relatively large duration compared with the average adjusted duration of 1 second. Some 622 standardized duration observations exceed the cutoff point.

The results in Table 5 show a remarkable improvement in the performance of the model compared with the WACD(1,1) for the S&P500, with the Ljung-Box statistics dropping by a factor of 5, to levels commensurate with the models estimated for the NASDAQ data in earlier sections. The model supports the existence of two regimes, with quite distinct characteristics. In the first regime the mean adjusted duration is relatively small at 0.0312. The coefficient $\gamma_4^{(1)}$ is insignificant at 10% so that dropping that coefficient makes the preferred form a TWACD(3,1:4,1). The sum of the γ_j and ω_j coefficients in this first regime is 0.95, indicating considerable persistence. The volume coefficient, $v^{(1)}$ is negative and significant, so that as previously increased volume results in decreased trade duration.

[2] Grid searching and Q-Q plots were used to examine sensitivity to the choice of cutoff, with 19 seconds producing the best fit.

In the second regime, however, a number of important differences are evident. Firstly, the mean duration, $\omega^{(2)}$ is increased 10 fold over the first regime, although this estimate is statistically insignificant. The role of volume with these longer duration transactions is also negative but is increased by over 6 times that of the first regime. The sum of the $\gamma_j^{(2)}$ and $\omega_1^{(2)}$ coefficients is greater than 1, due mainly to the estimate of $\omega_1^{(2)}$ indicating an extremely high degree of persistence in these right tail duration observations, a feature of the data which is not well handled by the standard model specifications.

The threshold ACD model provides a much improved description of the S&P500 data than previous simpler specifications. There is a clear need to account for non-linearities in this dataset and a future research agenda would be to explore the use of mixture models such as De Luca and Gallo (2004) and Hujer and Vuletić (2007) and stochastic volatility duration models such as Ghysels, Gourieroux and Jasiak (2004) which hold promise of more flexibly incorporating the possibility of different regimes in the data.

6 Conclusion

This paper provides, to the best of the authors' knowledge, the first attempt to model the time between trade durations of an electronic after-hours equity futures market. The contributions of the paper are the application to the previously unexploited after-hours electronically traded data, the use of a much longer data sample than previously explored in models of trade duration, and the use of volume as an informative mark. The preferred modelling framework is found to include relatively long lag lengths and two regimes separating long durations from the remainder of the data.

The markets studied comprise data from the standard NASDAQ and S&P500 equity futures contracts traded on the Chicago Mercantile Exchange using data from the GLOBEX electronic trading platform during periods when the open out-cry market for this contract is closed. The empirical results show that the trade duration of the equity market future contracts for the NASDAQ are characterized by relatively low autocorrelation and strong clustering, regardless of the distributional assumptions employed. In the S&P500 data, the majority of the distribution also exhibits low correlation and high clustering, but large duration observations require a separate specification characterized by higher autocorrelation and little clustering. The results show that the addition of volume information to the ACD model captures a statistically significant negative relationship between the trade duration and volume, consistent with either of two possibilities. The first of these possibilities it that public news results in large volume and high trade intensity as market participants adjust portfolios, and the second is that in the absence of public information, market participants interpret large volume trades as indicative of private information which feeds back to encourage further trading activity.

References

[1] Bauwens, 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.

[2] Bauwens,L and Verdas, D. (2004), "The Stochastic Conditional Duration Model: A Latent Variable Model for the Analysis of Financial Durations", *Journal of Econometrics*, Vol 119 (2), 381-482.

[3] Coppejans, M. and Domowitz (1999), "Pricing Behavior in an Off-hours Computerized Market", *Journal of Empirical Finance*, 6, 583-607.

[4] De Luca, G. and Gallo, G. (2004), "Mixture Processes for Financial Intradaily Durations", *Studies in Nonlinear Dynamics and Econometrics,* 8, 1-12.

[5] Dungey, M., Fakhrutdinova, L. and Goodhart, C. (2009), "After-hour Trading in Equity Futures Markets", *Journal of Futures Markets*, 29, 114-136.

[6] Easley, D. and O'Hara, M. (1992), "Time and the Process of Security Price Adjustment", *Journal of Finance*, 47, 567-607.

[7] Engle, R.F. and Russell, J. (1997), "Forecasting the Frequency of Changes in Quoted Foreign Exchange Prices with the Autoregressive Conditional Duration Model", *Journal of Empirical Finance*, 4 187-212

[8] Engle, R.F. and Russell, J. (1998), "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data", *Econometrica*, 66 1127-1162.

[9] Engle, R.F. and Russell, J. (2005), "A Discrete-State Continuous-Time Model of Financial Transactions Prices and Times: The Autoregressive Conditional Multinomial Autoregressive Conditional Duration Model", *Journal of Business & Economic Statistics,April* 2005, 23,2.

[10] Ghysels, E., Gourieroux, C. and Jasiak, J. (2004), "Stochastic Volatility Duration Models.", *Journal of Econometrics* 119,413-433.

[11] Grammig, J. and Maurer, K.-O. (2000), "Non-monotonic Hazard Functions and the Autoregressive Conditional Duration Model", *Econometrics Journal*, 3, 16–38.

[12] Hautsch, N. (2002), "Modelling Intraday Trading Activity Using Box-Cox-ACD Models", No 02-05, *CoFE Discussion Paper*, Center of Finance and Econometrics, University of Konstanz

[13] Hautsch, N. (2006) "Testing the Conditional Mean Function of Autoregressive Conditional Duration Models", Discussion Paper, 2006-06, Finance Research Unit, Department of Economics, University of Copenhagen, and Discussion Paper 2008-06, CRC 649, Berlin.

[14] Hasbrouck, J. (2003). "Intraday Price Formation in the US Equity Index Markets". *Journal of Finance*, 58, 2375-2400.

[15] Hujer, R. , Vuletic, S. and Kokot, S., "The Markov Switching ACD Model (2002)". Johann Wolfgang Goethe-University, Finance & Accounting Working Paper No. 90

[16] Hujer, R. and Vuletić, S. (2007), "Econometric Analysis of Financial Trade Processes by Discrete Mixture Duration Models", *Journal of Economic Dynamics and Control*, 31, 635-667.

[17] Jasiak, J. (1998), "Persistence in Intertrade Durations", *Finance*, 19, 166-195.

[18] Lunde, A. (2000) "A Generalized Gamma Autoregressive Conditional Duration Model", Discussion paper, Aarlborg University.

[19] Zhang, M, Russell, J. and Tsay, R. (2001) "A Nonlinear Augoregressive Conditional Model with Applications to Financial Transactions Data", *Journal of Econometrics*, 104, 179-207.

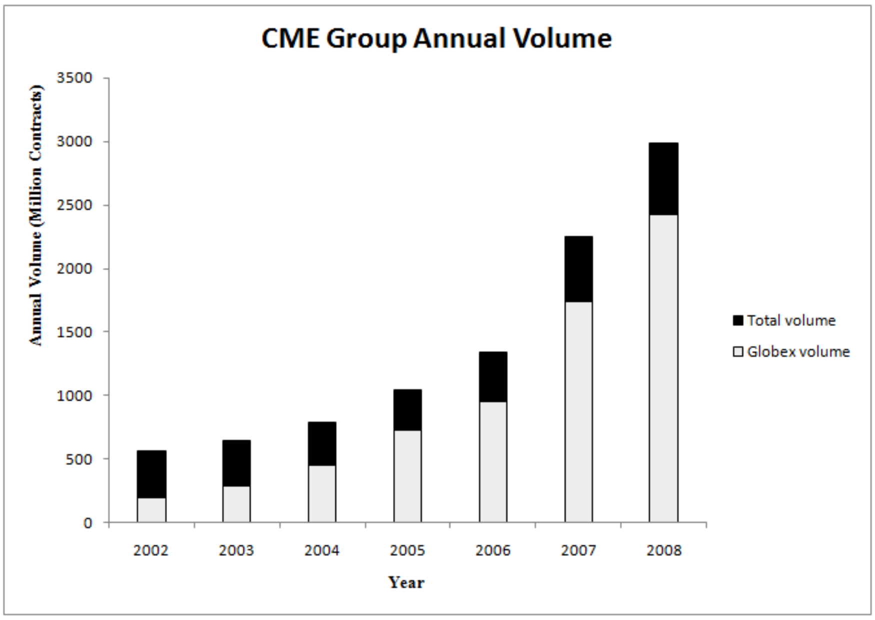

Annual Volume of Trade on the CME.

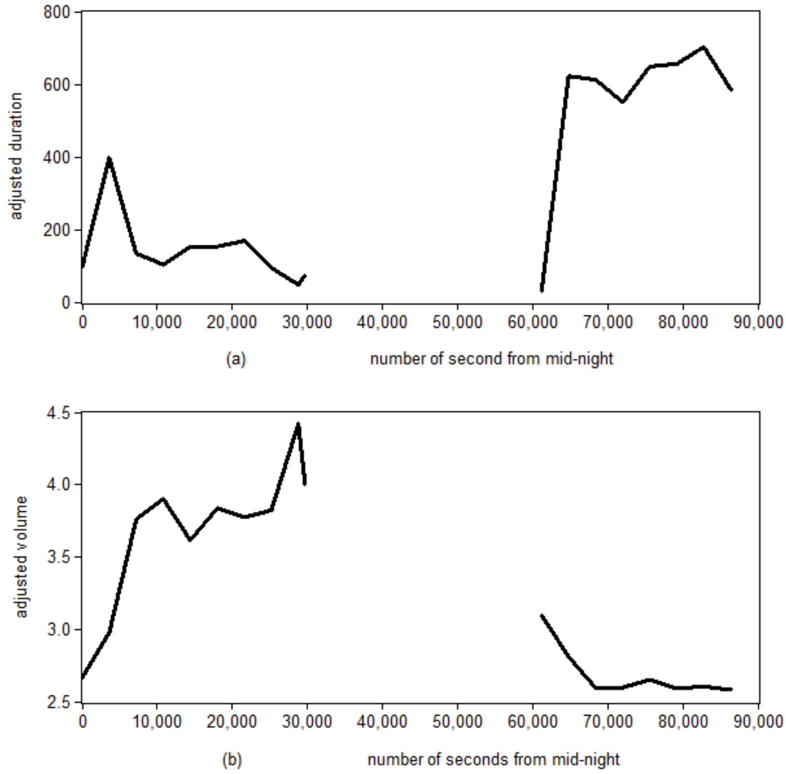

daily (a) durations and (b) volume pattern for NASDAQ

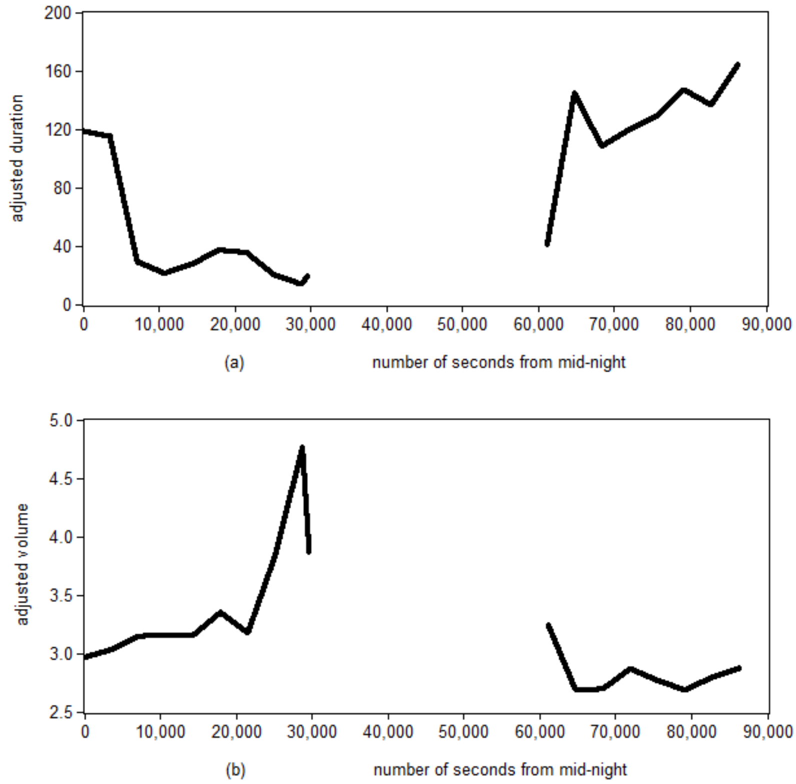

Daily (a) duration and (b) volume pattern for S&P500 data

Table 1:
Descriptive Statistics for adjusted durations in the NASDAQ and S&P500.

	NASDAQ	S&P500
number of observations	149314	684010
mean	0.9994	0.9973
max	66.3147	70.5960
min	0.0014	0.0061
variance	4.1081	3.6010
skewness	5.7514	5.4368
kurtosis	67.4259	62.4087
Jaque-Bera (p-value)	0.0000	0.0000

Table 2:
Parameter Estimates for ACD(1,1) models of the NASDAQ and S&P500 with different distributional assumptions; standard errors(), all parameters are significant at the 1% level.

Parameter	EACD(1,1)	WACD(1,1)	GACD(1,1)
	NASDAQ		
ω	0.0196	0.0307	0.0570
	(0.0002)	(0.0018)	(0.0033)
γ_1	0.1158	0.1561	0.2121
	(0.0006)	(0.0041)	(0.0064)
ω_1	0.8720	0.8256	0.7779
	(0.0006)	(0.0051)	(0.0071)
α	-	0.5466	0.2714
		(0.0008)	(0.0052)
κ	-	-	3.4202
			(0.1165)
Ljung-Box Q(10)	239.0057	254.7124	258.3202
Ljung-Box Q(20)	372.6006	294.7861	280.1999
AIC	1.7020	0.8851	0.8694
SBC	1.7022	0.8854	0.8697
	S&P500		
ω	0.0125	0.0182	-
	(0.0001)	(0.0003)	
γ_1	0.0834	0.0938	-
	(0.0002)	(0.0007)	
ω_1	0.9074	0.8881	-
	(0.0003)	(0.0007)	
α	-	0.6668	-
	-	(0.0007)	
κ	-	-	-
Ljung-Box Q(10)	1400.8710	1123.8400	-
Ljung-Box Q(20)	1971.5400	1489.7580	-
AIC	1.7293	1.3711	-
SBC	1.7294	1.3712	-

Table 3:
Parameter Estimates for WACD(5,5) and GACD(3,3) models of the NASDAQ; standard errors are given in parentheses, all parameters are significant at the 1% level.

Parameter	WACD(1,1)	GACD(3,3)
ω	0.0009	0.0030
	(0.0001)	(0.0003)
γ_1	0.2321	0.2977
	(0.0009)	(0.0063)
γ_2	-0.2510	-0.3261
	(0.0003)	(0.0081)
γ_3	0.0188	0.0465
	(0.0003)	(0.0022)
γ_4	0.0196	-
	(0.0007)	
γ_5	-0.0100	-
	(0.0006)	
ω_1	1.5840	1.5498
	(0.0001)	(0.0035)
ω_2	-0.5463	-0.5238
	(0.0002)	(0.0043)
ω_3	-0.0262	-0.0435
	(0.0002)	(0.0009)
α	0.5485	0.2784
	(0.0009)	(0.0005)
κ	-	3.2759
		(0.0038)
Ljung-Box Q(10)	183.7966	293.5161
Ljung-Box Q(20)	191.1481	306.5236
AIC	0.8799	0.7370
SBC	0.8807	0.7386

Table 4:
Parameter Estimates for WACD(1,1) models of the NASDAQ and S&P500 with volume; standard errors are given in parentheses, all parameters are significant at the 1% level.

Parameter	NASDAQ	S&P500
ω	0.0359	0.0227
	(0.0014)	(0.0031)
γ_1	0.1541	0.0938
	(0.0045)	(0.0068)
ω_1	0.8255	0.8857
	(0.0049)	(0.0072)
α	0.5470	0.6673
	(0.0007)	(0.0007)
v	-0.0011	-0.0007
	(0.0000)	(0.0000)
Ljung-Box Q(10)	246.6640	1093.8550
Ljung-Box Q(20)	284.1930	1422.3710
AIC	0.8841	1.3701
SBC	0.8844	1.3702

Table 5:
Parameter Estimates for ThresholdWACD(4,1:4,1) model of the S&P500;
standard errors are given in parentheses.

Parameter	estimate	standard error
$\omega^{(1)}$	0.0312	(0.0003)
$\gamma_1^{(1)}$	0.1653	(0.0018)
$\gamma_2^{(1)}$	-0.0486	(0.0022)
$\gamma_3^{(1)}$	-0.0173	(0.0019)
$\gamma_4^{(1)}$	-0.0016	(0.0013)
$\omega_1^{(1)}$	0.8744	(0.0007)
$\alpha^{(1)}$	0.6670	(0.0007)
$\upsilon^{(1)}$	-0.0009	(0.0000)
$\omega^{(2)}$	0.3718	(0.2322)
$\gamma_1^{(2)}$	0.0838	(0.0094)
$\gamma_2^{(2)}$	-0.1279	(0.0175)
$\gamma_3^{(2)}$	-0.0473	(0.0106)
$\gamma_4^{(2)}$	-0.0572	(0.0131)
$\omega_1^{(2)}$	1.3734	(0.0762)
$\alpha^{(2)}$	-0.0038	(0.0158)
$\upsilon^{(2)}$	0.5827	(0.0203)
Ljung-Box Q(10)	225.5850	
Ljung-Box Q(20)	261.6080	
AIC	1.3678	
SBC	1.3681	