

Time varying trade intensities and the Deutsche Telekom IPO

Zeitvariable Handelsintensitäten und die Deutsche Telekom IPO

Reinhard Hujer, Joachim Grammig and Stefan Kokot
University of Frankfurt *

ABSTRACT

We apply the Threshold Autoregressive Conditional Duration Model (TACD) as proposed by Zhang, Russell, and Tsay (1999) to model the after market trading duration process associated with the initial public offering of the Deutsche Telekom AG share in November of 1996. Special emphasis is devoted to the empirical specification of intra-day seasonality and to the detection of non-stationarity and structural breaks in the trading process.

Keywords: Initial public offerings, financial transactions data, nonlinear time series, autoregressive conditional duration.

JEL classification: C22, C41, G32.

ZUSAMMENFASSUNG

Wir verwenden das Threshold Autoregressive Conditional Duration Model (TACD) von Zhang, Russell, and Tsay (1999) um den Handelsprozeß der Deutsche Telekom AG Aktie im Anschluß an das Going Public im November 1996 zu analysieren. Besonderer Wert wird auf die empirische Spezifikation der Intra-Tages Saisonalität gelegt, sowie auf die Erfassung von Nichtstationarität und Strukturbrüchen.

Schlüsselwörter: Initial Public Offerings, Finanztransaktionsdaten, Nicht-lineare Zeitreihenmodelle, Autoregressive Conditional Duration Modelle.

JEL Klassifikation: C22, C41, G32.

*Institute for Statistics and Econometrics, Faculty of Economics and Business Administration, Johann Wolfgang Goethe-University, Mertonstr. 17, 60054 Frankfurt. Phone: +49 69 798 22893 Fax: +49 69 798 23673. E-mail: hujer@wiwi.uni-frankfurt.de, grammig@wiwi.uni-frankfurt.de, kokot@wiwi.uni-frankfurt.de. We thank Luc Bauwens, Pierre Giot, Kai-Oliver Maurer and David Veredas for their close collaboration, Roland Straub and Sandra Vuletić for their valuable research assistance and Marco Caliendo, Dubravko Radić and Stephan Thomssen for usefull comments and their help in the preparation of this paper. The usual disclaimer applies.

1. INTRODUCTION

The class of Autoregressive Conditional Duration (ACD) models proposed by Engle and Russell (1997, 1998) is designed to analyze irregularly spaced time series data and to forecast time intervals between reoccurring events like trades or price changes. The model class is based on an econometric approach which combines elements of the GARCH class of models and econometric tools for the analysis of transition data [Lancaster (1990)]. ACD models have been almost exclusively applied to high frequency financial data stemming from electronic stock and foreign exchange markets. The ACD model is not only useful in forecasting durations between successive trade events, they can also be combined with GARCH-models and used to forecast volatility [see Ghysels and Jasiak (1998), Grammig and Wellner (1999) and Engle (2000)].

The application of the ACD approach in order to model transition processes in financial markets is motivated by the high autocorrelation of durations between successive transactions. Given the institutional framework of the trading process in a stock exchange market, a large part of the transaction duration dependence is due to intra-day seasonality factors such as the opening and closing of stock exchanges in different time zones, lunch breaks and other periodic events that occur during the regular trading interval of the market under consideration.

While there is a large quantity of empirical as well as theoretical contributions that deal with abnormal behaviour of stock prices of IPO's in the after market trading process, the bulk of the international empirical evidence on IPO's rests on low frequency, cross-sectional data sets, containing firm specific characteristics as well as data on market evolution during the IPO episode. Although high frequency data for a large number of financial markets have become accessible to researchers in the past decade, there are very few papers that deal with the IPO issue using these data sets.¹ This has also consequences for the current state of art of the economic theory, that is concerned with IPO issues.

A common feature to most of the models in the literature is that they concentrate on the identification of several determinants of these anomalies. In a way this concentration reflects the empirical features or 'stylized facts' of the IPO process that have been established until now et vice versa. The process of trading itself on the after-market has no explicit role in these models and accordingly most empirical contributions don't bother either.

This is a serious flaw, since in many recent contributions that are related to the microstructure theory of financial markets it has been recognized, that several institutional as well as informal features of the trading process will influence prices and volume. If this is true for 'established' assets in general, then it is very likely, that the performance of any IPO will not remain unaffected by these factors.

The remaining sections of our paper are organized as follows: In section

¹A notable exception to this practice is the paper by Ellis, Michaely, and O'Hara (2000).

two, we review previous empirical findings on the IPO process and give some background information on the institutional circumstances of the Telekom IPO. Specification and estimation of standard ACD models is treated in section 3. We also discuss a new nonlinear ACD model, that was only recently developed by Zhang, Russell, and Tsay (1999). ML-estimation as well as specification tests that are useful in assessing the fit of *threshold ACD* (TACD) models are outlined in this section as well. In section 4 we present empirical evidence on the characteristics of financial transaction durations which motivate the application of the ACD approach and discuss ways to account for intra-day seasonality. In this section we apply EACD, WACD and TACD models to our data set and test the performance of the TACD approach. We conclude in section five with a summary of our main findings and an outlook.

2. INITIAL PUBLIC OFFERINGS

2.1. Previous findings

There are three stylized facts of initial public offerings [henceforth abbreviated as IPO] that are confirmed in a large number of empirical studies for many countries. These findings may be considered as anomalies, because of the obvious incompatibility with the notion of market efficiency [see Ibbotson and Ritter (1995)]:

1. IPO shares are regularly underpriced, i.e. prices at the end of the first trading day tend to exceed the offering price by large.
2. IPO shares perform remarkably bad in the long run, i.e. investors buying IPOs on the secondary market earn substantial negative abnormal returns over longer horizons.
3. IPO volume and initial returns tend to cluster in periods of so called "hot issue markets", i.e. IPO-volume is high in periods of high underpricing. It seems as if there are cycles in both volume and average returns.

A summary of the international empirical evidence, mostly based on data sets for a cross section of IPOs, is provided in Loughran, Ritter, and Rydqvist (1994). Most of the studies confirm the short-run underpricing of IPOs, although the average level of initial returns differs considerably from country to country. Empirical evidence on German IPOs that is basically consistent with international findings on underpricing and long-run underperformance is provided by Uhler (1989), Wasserfallen and Wittleder (1994), Schuster (1996), Ljungqvist (1997), Stehle and Ehrhardt (1999) and Schlag and Wodrich (2000). These papers typically use data sets for a selection of companies containing price and volume information as well as a host of additional company specific characteristics.

Explanations for the observed patterns are given in several papers on the basis of information asymmetries among the three parties usually involved in an

IPO. *Private and institutional investors* typically have different degrees of uncertainty about the 'true' value of the issue. *Underwriters* have private information about investor's demand schedules and may also have private information that aids in the valuation of the share. *Issuers* usually have superior information (compared to investors and underwriters) about the 'true' value of the issue. There are several hypothesis offered for the initial underpricing and the long run underperformance-phenomena, that draw on this constellation. Apart from these asymmetric information stories, there are a number of alternative explanations for the IPO-anomalies, emphasizing other aspects of the IPO-process [see Ibbotson and Ritter (1995)].

A common feature to all of these models is that, since they concentrate on the identification of several determinants of the offer price and the influence of distinct types of private investors on the initial share allocation, the empirical implications that can be drawn from these models refer solely to the 'pre-market' or subscription phase of the IPO-process and on the initial, i.e. first day return. Typically, secondary markets are treated only as a device that is used by market participants to determine the 'true' value of IPO-shares. The trading process on the after-market has no explicit role in these stories.

Recent contributions to the branch of microstructure theory of financial markets² suggest, that several institutional as well as informal features of the trading process itself do exert a non negligible influence on the outcome of trading in terms of prices and volume. If this is true for 'established' assets in general, then it is very likely, that the performance of any IPO will not remain unaffected by these factors. An empirical investigation of the features of the trading process of a large IPO might therefore be helpful in several ways, e.g. in discovering new anomalies and establishing an extended set of stylized facts.

The empirical evidence on the microstructure of IPOs is very limited. To the best of our knowledge, the paper by Grammig, Hujer, Kokot, and Maurer (1998) is the only study that uses high-frequency data to model some features of the after market trading process for a German IPO. However, empirical evidence on market microstructure characteristics of the IPO-process for an American stock exchange has recently been provided in a very interesting, explorative study by Ellis, Michaely, and O'Hara (1999). They focus on the role played by the main underwriter in the after market trading process.

Since there is at the moment no economic theory available, that addresses explicitly the microstructure of the IPO after market trading process, our study serves primarily explorative purposes. Besides, the environment of a large IPO constitutes a challenging test of the capabilities of the ACD-model to describe and forecast outcomes of the trading process on a continuous, order driven asset market.

²See e.g. Goodhart and O'Hara (1997) or Madhavan (2000).

2.2. Institutional background of the Deutsche Telekom IPO and the IBIS trading system

The Deutsche Telekom IPO in November 1996 was the first step of the largest ever privatization project in Germany which had a significant effect on the investment culture of Germany's private households. With the Deutsche Bank, the Dresdner Bank, and Goldman Sachs as global coordinators and unprecedented marketing efforts, the first tranche of 713 Mio. shares, about 25 % of the total shares of the formerly 100 % state owned German telecommunication monopolist, was offered to private investors. Special incentive programs for private households, including price reductions and bonus issues, led to a 5-6 fold oversubscription.

23 millions of the offered shares were distributed among the Deutsche Telekom employees. 67 % of the remaining shares were bought by German private (174 Mio. shares) and institutional investors (254 Mio. shares). International investors were located in the USA (14 % of shares), the UK (8 % of shares), other Europe (6 % of shares) and Asia (5 % of shares). The issue price was 28.5 DM (18.89 \$) per share resulting in a total proceeds of the issue of 20 trillion DM. Regular trading started on Monday, the 18th of November on the Frankfurt and New York stock exchanges. In Tokyo, Telekom trading started one day later.

Due to the expected extraordinary trading intensity, the trading hours in the German electronic trading system (IBIS) were prolonged in the first Telekom trading week, from 8.30 a.m. - 5.00 p.m. to 08.30 a.m. - 7:00 p.m. (CET).³ The Deutsche Telekom was immediately included in the DAX, the top 30 blue chip index for German stocks (replacing the Metallgesellschaft), making the Deutsche Telekom the number 9 of the 30 blue chips in the DAX in terms of market capitalization.

In order to apply the ACD model, we use IBIS - Integrated Stock Exchange Trading and Information System - transaction data provided by the the Karlsruher Kapitalmarktdatenbank (Lüdecke (1996)). IBIS is an electronic trading system of the Deutsche Börse AG where members of the system trade securities in an interactive double auction framework.

Each trader has the opportunity to display on his/her IBIS computer screen the complete system order book for the securities of interest and can choose ("hit") the bids and the asks defined by price and volumes. The order book is displayed with price ascending bids and price descending asks including the volume of each bid and ask. Members of the system can freely enter and delete their bids and asks in the IBIS order book. The bids and asks entered in the system are binding. The IBIS system does not provide an automatic matching of bids and asks, hence a transaction must always be initiated by a trader who is willing to hit a bid or ask.

³In the first week after the IPO, roughly 13% of all trades of the Telekom share on the IBIS system occurred in the prolonged trading phase between 5.00 p.m. and 7.00 p.m.

A trader willing to buy from the order book is restricted to hit the cheapest ask unless the volume the trader wants to buy is smaller than the volume offered by the cheapest ask and this offer to sell the security contains a special mark that the supplier is willing to trade only the volume entered in the system and not less. In this case the initiating buyer is allowed to hit the next (more expensive) ask in the order book. A buying initiator can also select more than one ask from the list, if he chooses them in the price ascending order. A trader is thus able to buy the whole ask side of the market. For a seller willing to hit the bids in the electronic order book, the same logic applies.

3. THE CLASS OF ACD MODELS

3.1. Some Basic Models

ACD models are designed to account for autocorrelation often observed in time series of arrival times between successive occurrences of certain events associated with the trading process. The definition of the trading event depends on the aim of the study. Examples include the time between successive trades, the time until a price change occurs or until a prespecified number of shares or level of turnover has been traded.⁴

Let $x_i = t_i - t_{i-1}$ be the time interval between any two successive trading events and denote the conditional mean of x_i as

$$(1) \quad E(x_i | \Omega_{i-1}) = \psi_i(\Omega_{i-1}; \theta_\psi) \equiv \psi_i,$$

where Ω_{i-1} is the information set at the beginning time of the spell and θ_ψ is the corresponding set of parameters that determine the conditional mean function. In this framework all of the time dependence of the duration process is captured by the conditional mean function ψ_i . The ACD model is defined by some parameterization of this conditional mean and the following decomposition

$$(2) \quad x_i = \psi_i \cdot \varepsilon_i,$$

where the stochastic process $\{\varepsilon_i\}$ is independent of ψ_i and i.i.d. with a non-degenerate density function $g(\cdot; \theta_\varepsilon)$ with parameters⁵ θ_ε and support on the positive real line, since $\varepsilon_i > 0$. Note that the unconditional expectation of ε_i is equal to unity. This follows from a straightforward application of the law of iterated expectations to equation (2).

Engle and Russell (1998) introduced the EACD and the WACD models. The EACD model arises when $g(\cdot; \theta_\varepsilon)$ is an exponential distribution given by

$$(3) \quad g(\varepsilon_i; \lambda) = \frac{1}{\lambda} \cdot \exp\left(-\frac{\varepsilon_i}{\lambda}\right),$$

⁴Naturally, the *price*, *volume* and *turnover* duration processes arise from the trade durations series by dropping intervening observations from the sample, thus yielding a 'thinned' [Engle and Russell (1998)] or 'weighted' [Gouriéroux and Jasiak (2000)] duration process.

⁵The parameters of the conditional mean θ_ψ and of the conditional density θ_ε are assumed to be variation free, i.e. if $\theta_\psi \in \Theta_\psi$ and $\theta_\varepsilon \in \Theta_\varepsilon$, then $(\theta_\psi, \theta_\varepsilon) \in \Theta_\psi \times \Theta_\varepsilon$.

with parameter λ equal to one. In this case it can be easily shown that the duration x_i has a conditional exponential distribution with parameter $\lambda_i = \psi_i$. The WACD model is obtained, when instead of the standard exponential a Weibull (γ, ϕ) distribution is assumed for the standardized durations

$$(4) \quad g(\varepsilon_i; \gamma, \phi) = \frac{\gamma}{\phi} \cdot \left(\frac{\varepsilon_i}{\phi}\right)^{\gamma-1} \cdot \exp\left[-\left(\frac{\varepsilon_i}{\phi}\right)^\gamma\right],$$

where $\gamma, \phi > 0$ and $E(\varepsilon_i) = \phi \cdot \Gamma\left(\frac{1+\gamma}{\gamma}\right) = 1$. It follows then that the conditional distribution of x_i is also Weibull with parameters γ and $\phi_i = \frac{\psi_i}{\Gamma\left(\frac{1+\gamma}{\gamma}\right)}$. Since the exponential distribution is just a special case of the Weibull when $\gamma = 1$, the EACD model is nested in the WACD. One drawback of these models is that the distributional assumption implies that the conditional intensity function is either constant as in the EACD or a monotonic function of the length of the ongoing spell as in the WACD. This is at odds with observed behaviour in many financial applications, where inverse U-shaped patterns are typically recognized.

Less restrictive ACD models arise, if the distributional assumptions are further modified. Recent extensions include the use of the Burr, the gamma and the generalized gamma distribution, leading to the BACD-model of Grammig and Maurer (2000), the GACD of Lunde (1999a) and the GG-ACD model used in Lunde (1999b). Both, the EACD and the WACD models are nested within the BACD and the GG-ACD framework. The latter also nests the GACD. In addition to the likelihood approach, GMM-type estimators have been used to estimate the parameters of interest, see Grammig and Wellner (1999).

In the simplest case of the ACD model, the parameterization of the conditional mean ψ_i is completely analogous to the parameterization of the conditional variance in a GARCH model. Thus, the ACD(m, q) model arises when the conditional mean is given by

$$(5) \quad \psi_i = \omega + \sum_{j=1}^m \alpha_j x_{i-j} + \sum_{j=1}^q \beta_j \psi_{i-j},$$

which can be expressed as an ARMA ($\max(m, q), q$)-process in the following manner

$$(6) \quad x_i = \omega + \sum_{j=1}^{\max(m, q)} (\alpha_j + \beta_j) x_{i-j} - \sum_{j=1}^q \beta_j \eta_{i-j} + \eta_i,$$

where $\eta_i \equiv x_i - \psi_i$ is a (non-gaussian) martingale difference sequence. Expressions for the unconditional mean and the variance of the durations can easily be determined from the ARMA-representation. Since durations are non-negative, the parameters have to obey some constraints. The usual restriction is that $\omega \geq 0, \alpha_j \geq 0, \beta_j \geq 0, \forall j$, although this can be relaxed to some extent as shown by Nelson and Cao (1992) in the GARCH context.

An alternative way to get rid of these unpleasant non-negativity restrictions, is to use a specification for the conditional mean ψ_i that closely resembles the EGARCH specifications of Nelson (1991). The following two specifications have been suggested by Bauwens and Giot (1997) for the ACD model

$$(7) \quad \psi_i = \exp \left(\omega + \sum_{j=1}^q \alpha_j \ln(x_{i-j}) + \sum_{j=1}^p \beta_j \ln(\psi_{i-j}) \right),$$

$$(8) \quad \psi_i = \exp \left(\omega + \sum_{j=1}^q \alpha_j \varepsilon_{i-j} + \sum_{j=1}^p \beta_j \ln(\psi_{i-j}) \right).$$

As in Nelson's EGARCH model, analytical expressions for the unconditional moments of x_i in the logarithmic ACD specification are quite cumbersome, see Bauwens and Giot (2000) for details.

3.2. The threshold ACD model

The threshold ACD model has recently been suggested by Zhang, Russell, and Tsay (1999) to overcome some of the observed deficiencies of the above mentioned models. This model is in the spirit of threshold autoregressive models that have been introduced by Tong and Lim (1980). In the recent literature it has been recognized, that the dynamics of the duration process may be governed by a very complex, possibly non-linear stochastic system [see e.g. Gouriéroux and Jasiak (1999)]. The basic idea in the TACD framework is that such complex systems may be approximated by piece-wise linear subsystems. The process is allowed to 'visit' each subsystem or regime more than once, thus permitting some interesting cycling patterns.

Consider the positive valued stochastic process of event durations $\{x_i\}$. In the TACD model $\{x_i\}$ may be generated by a different data generating process depending on the observed value of a threshold variable Z_{i-d} , where the delay parameter d is a positive integer. If there are S different regimes in total, the TACD model is given by the usual decomposition of the duration as in equation

$$(9) \quad x_i = \psi_i \cdot \varepsilon_i,$$

and a regime dependent specification for the conditional mean

$$(10) \quad \psi_i = \begin{cases} \psi_i^{(1)} & \text{if } Z_{i-d} \in R_1 \\ \psi_i^{(2)} & \text{if } Z_{i-d} \in R_2 \\ \dots & \dots \\ \psi_i^{(S)} & \text{if } Z_{i-d} \in R_S \end{cases},$$

where $R_s = [r_{s-1}, r_s)$ and the threshold values r_s are such that $0 = r_0 < r_1 < \dots < r_S = \infty$ for $s = 1, 2, \dots, S$. The standardized duration series ε_i is assumed to be independently, but not necessarily identically distributed across different regimes. Instead we assume that $g(\cdot; \theta_\varepsilon^{(s)})$ is a distribution with mean equal to one, support in \mathbb{R}_0^+ and a parameter vector $\theta_\varepsilon^{(s)}$ that is regime specific as well.

The conditional mean function $\psi_i^{(s)}$ is analogous to the GARCH type specification in (5), except for the dependence of the parameters on the regime selected by the threshold variable Z_{i-d} . Thus, the conditional mean function is given by

$$(11) \quad \psi_i^{(s)} = \omega^{(s)} + \sum_{j=1}^m \alpha_j^{(s)} \cdot x_{i-j} + \sum_{j=1}^q \beta_j^{(s)} \cdot \psi_{i-j}.$$

A quite general functional form for the threshold variable Z_{i-d} is given by

$$(12) \quad Z_{i-d} = Z(\Omega_{i-d}; \theta_Z),$$

where the threshold variable is allowed to dependent on the whole information set at time $i-d$ and a set of parameters denoted by θ_Z . In empirical applications one would like to restrict this dependence. A *self exciting threshold model* is obtained, when Z_{i-d} is allowed to depend on past realized values of the dependent series $\{x_i\}$. A very convenient specification in this case is to let the threshold variable be equal to a lagged value of the process, as in

$$(13) \quad Z_{i-d} = x_{i-d}.$$

In this case the support of the threshold variable, as defined by the union of the regime sets $R = \bigcup_{i=1}^S R_i$ is just a partition of the positive real line \mathbb{R}_0^+ . In order to estimate the parameter vector $\theta = (\theta_\psi^{(1)}, \dots, \theta_\psi^{(S)}, \theta_\varepsilon^{(1)}, \dots, \theta_\varepsilon^{(S)})'$ by ML techniques we have to specify a distribution for $\{\varepsilon_i\}$. Zhang, Russell, and Tsay (1999) employ a generalized gamma distribution, with density function given by

$$(14) \quad g(\varepsilon_i; \delta, \vartheta, \varphi) = \frac{\delta \cdot (\varepsilon_i)^{\delta\vartheta-1}}{\varphi^{\delta\vartheta} \cdot \Gamma(\vartheta)} \cdot \exp\left[-\left(\frac{\varepsilon_i}{\varphi}\right)^\delta\right],$$

which has moments equal to

$$(15) \quad E(\varepsilon_i^m) = \varphi^m \cdot \frac{\Gamma(\vartheta + \frac{m}{\delta})}{\Gamma(\vartheta)}.$$

A natural way to proceed is to parameterize the generalized gamma distribution in terms of the parameter φ , so that the conditional mean function is given by

$$(16) \quad \psi_i^{(s)} = \varphi_i^{(s)} \cdot \frac{\Gamma(\vartheta^{(s)} + \frac{1}{\delta^{(s)}})}{\Gamma(\vartheta^{(s)})}.$$

Thus, the log likelihood function for a TACD (m, q) model, conditional on the set of threshold values and an appropriate number of initial observations to start the recursion, can be written as follows

$$(17) \quad \ln \mathcal{L}(\theta | x_1, \dots, x_N; r_1, \dots, r_S) = \sum_{i=1}^N \sum_{s=1}^S I(Z_{i-d} \in R_s) \left[\ln(\delta^{(s)}) + (\vartheta^{(s)} \cdot \delta^{(s)} - 1) \ln\left(\frac{x_i}{\varphi_i^{(s)}}\right) - \ln\left(\varphi_i^{(s)} \Gamma(\vartheta^{(s)})\right) - \left(\frac{x_i}{\varphi_i^{(s)}}\right)^{\delta^{(s)}} \right],$$

where $I(\cdot)$ denotes the indicator function, that is equal to one if the condition in brackets is satisfied and zero otherwise.

3.3. Specification tests based on parametric density forecasts

Despite the recent boom of empirical analyses of financial duration processes, the literature has so far devoted little attention to testing the specification of the econometric model. It is common to perform simple diagnostic tests to check whether the standardized durations are independent and identically distributed (iid). While most papers use the Ljung-Box statistic to test for serial correlation, only a few test whether the distribution of the durations is correctly specified.

Bauwens, Giot, Grammig, and Veredas (2000) use a method advanced by Diebold, Gunther, and Tay (1997) to test the forecast performance of ACD models. Denote by $\{\hat{f}_i(x_i | \Omega_{i-1})\}$ a sequence of one-step-ahead density forecasts evaluated using parameter estimates from some parametric model and by $\{f_i(x_i | \Omega_{i-1})\}$ the sequence of densities corresponding to the true, but unobservable data generating process of x_i . The $\{\hat{f}_i(x_i | \Omega_{i-1})\}$ sequence is given by the conditional density of the duration series $\{x_i\}$ and thus depends directly on the density $g(\cdot; \theta_\varepsilon)$ specified for the standardized duration series $\{\varepsilon_i\}$ and the estimates of the associated parameter vector θ .

Diebold, Gunther, and Tay (1997) show that a forecast based on the correctly specified density will be preferred by all forecast users regardless of the form of their loss functions. This suggests that forecast performance can be evaluated by assessing whether the forecasting densities are correct, i.e. whether

$$(18) \quad H_0 : \quad \{\hat{f}_i(x_i | \Omega_{i-1})\} = \{f_i(x_i | \Omega_{i-1})\}.$$

Since the true distribution $f_i(x_i | \Omega_{i-1})$ is never observed, the sequence of conditional empirical distribution functions defined by

$$(19) \quad \hat{z}_i = \int_{-\infty}^{x_i} \hat{f}_i(u | \Omega_{i-1}) du$$

is used as a test statistic. As shown by Rosenblatt (1952), under the null hypothesis the distribution of the sequence of probability transforms $\{\hat{z}_i\}$ is iid $U(0, 1)$, so that any test for uniformity of the $\{\hat{z}_i\}$ sequence can be used to assess the forecast performance of the model under consideration. The test can also be used to compare the performance of several non-nested model specifications.

The recommendation of Diebold, Gunther, and Tay (1997) is to supplement statistical tests for iid uniformity by graphical tools. Departures from uniformity can easily be detected using a histogram plot based on the $\{\hat{z}_i\}$ sequence. A straightforward χ^2 goodness-of-fit test can be computed by exploiting the statistical properties of the histogram under the null hypothesis of uniformity. Furthermore the autocorrelogram of the $\{\hat{z}_i\}$ sequence can help to identify potential deficiencies of a model to account for the dynamics of the duration process.

4. EMPIRICAL RESULTS

4.1. Our dataset

The IBIS data provided by the Karlsruher Kapitalmarkt-Datenbank (KKMDB) contains records on each security traded in the IBIS system where an observation includes the transaction price, the transaction volume, the securities ID and the date and time of the transaction with an accuracy in hundredths of seconds. In a first step we select all observations of transactions of the Deutsche Telekom security from Nov. 18, 1996 until Dec. 20, 1996 and exclude the first and last transaction of each day (otherwise the transaction duration, would include the non trading time between successive days) and transactions during the prolonged trading hours of the first trading week.

The next step in our sample selection procedure involves the treatment of *split transactions*. If a trader has initiated a buy or sell by hitting more than one of the asks or bids within a very short time period, the database contains the sub-transactions (each order that was hit) as a separate trade, since more than two orders were involved on one side of the market and separate settlements may be needed. Therefore, a sequence of very short (but always non-zero) durations will be observed, that economically all belong to one trade. Thus the question arises, how to deal with these observations.

In our data set we have no information, that would allow us to identify these split transactions exactly. Therefore we consolidate trades according to the following algorithm: If we observe two successive transactions with a duration of less than one second, these are recognized as part of a split transaction. If the following trade occurs also within a second it must satisfy an additional condition: The sequence of transaction prices for the three transactions must be either non-increasing or non-decreasing.

Non-increasing prices would imply that a trader has initiated a split transaction on the bid side of the order book, i.e. he sold to the bid side with a

falling (or constant) prices sequence in the sub-transactions. Non-decreasing prices occur, when a trader initiated a split transaction on the ask side of the order book at increasing (or constant) ask prices. The consolidation stops if either the duration between two successive trades is longer than one second or the continuity condition on the price sequence is not met. The time stamp of the first transaction included is assigned to the consolidated trade, while the volume is the sum of the sub-transaction volumes and the price is the volume weighted average of the price sequence of the sub-transactions.⁶

We have also removed the observations immediately before and after a system breakdown on Dec. 13, 1996 between 09:00 h and 10:30 h in order to avoid introducing outliers to our data set. This sample selection resulted in 12038 Deutsche Telekom transaction durations. Sample descriptive statistics for the resulting time series are included in Table I.

Trade durations are known to be serially correlated. Figure 1 displays the sample autocorrelation function. The shape of the ACF suggests, that the series of transaction durations exhibits a strong form of time dependence, with autocorrelations that are relatively small, but significantly different from zero for very long lags. These findings are confirmed by the results of the Ljung-Box test statistics which reject the null hypothesis for a white noise process at virtually any lag we considered.⁷ On the other hand, both these findings do not suggest, that the series is non-stationary in general.

Other characteristics of the transaction time series become visible, if we take a look at average durations for different trading days and time intervals. The left panel of Figure 2 depicts the evolution of daily average trade durations for our sampling interval of 25 trading days. Since each trading day has the same length (8.5 hours), the estimate of the daily mean duration is approximately equal to the length of the trading day divided by the number of the trades.⁸ Thus, comparing mean durations for different trading days gives a crude measure of the overall level of market activity as indicated by the variation of the daily number of transactions.⁹

Figure 2 reveals, that daily averages of trade durations are not constant during the sample period, but rather change in time. Although at first glance the series of daily mean durations seems to exhibit trending behaviour, a closer

⁶In the resulting consolidated sample we still have 21 trades with durations of less than one second. All of these trades occurred after a consolidated trade and were not included by our algorithm, because they did not meet the continuity condition stated above.

⁷A more comprehensive set of summary statistics including the results of the Ljung-Box statistics as well as many other details is available by the authors upon request.

⁸The daily mean duration would be exactly equal to the trading day length divided by trades, if the time span between the opening of the exchange and the first trade as well as the time between the last trade and market closure would be treated as (censored) spells. It is common practice in the ACD literature to exclude all censored spells.

⁹The 95% confidence interval displayed in this and the following figures for daily and total sample averages are computed assuming a normal distribution for these sample statistics. This can be justified asymptotically if the appropriate conditions of the central limit theorem for serially dependent observations as stated e.g. in proposition 7.8 in Hamilton (1994), p. 193 hold.

Sample ACF




FIGURE 1: Sample autocorrelation function for Deutsche Telekom AG trade durations from Nov. 18, until Dec., 20, 1996.

inspection reveals that there seem to be two relatively long periods (Nov. 19 - Nov. 29 and Dec., 2 - Dec., 20) with an approximately constant average and random deviations around this average. These two subperiods are accompanied by three trading days, that differ strongly. The first two trading days (Nov., 18 & 19) have very high trading activity due to the IPO and are clearly not comparable to the rest of the sample. On Dec., 18 trading activity was extraordinarily low. The sample statistics for the four subsamples contained in Table I show, that durations appear to be more homogenous in each of the subperiods, as measured e.g. by the lower degree of overdispersion in each subsample.

We interpret these findings as occurrences of structural breaks in the trading process and therefore decided to split the sample in four subsamples defined by the aforementioned dates. The subsamples are ordered according to declining trading activity, with subsample A (Nov., 18 & 19) being the period with the highest activity and subsample D (Dec., 18) with the lowest. The subsamples B (Nov., 20 - 29) and C (Dec., 2 - 20, excluding Dec., 18) have intermediate activity levels.

The basic finding, that market activity, as measured by the average number of trades (or quotes) per time interval, tends to vary systematically with the time of day has been well documented for a number of financial markets and across many countries. Similar time of the day patterns have been reported for other characteristics of the trading process like returns, bid-ask spreads, volume and volatility. The right panel of Figure 2 depicts the number of transactions

TABLE I: Descriptive statistics for raw data

Statistic	Total	Sample A	Sample B	Sample C	Sample D
N	12038	2657	4828	4407	146
Mean	61,43	17,55	50,55	94,88	209,51
Standard deviation	110,14	27,69	76,56	143,07	282,88
Overdispersion	1,79	1,58	1,51	1,51	1,35
Skewness	5,11	5,18	3,47	4,11	2,39
Kurtosis	48,58	48,98	18,18	32,44	6,31
Minimum	0,13	0,28	0,18	0,13	0,28
25 %- Quartile	6,87	4,26	7,44	11,30	26,81
Median	20,63	8,04	21,17	40,25	113,00
75 %- Quartile	68,87	19,13	61,18	121,83	243,15
Maximum	2356,10	477,16	892,53	2356,10	1633,58

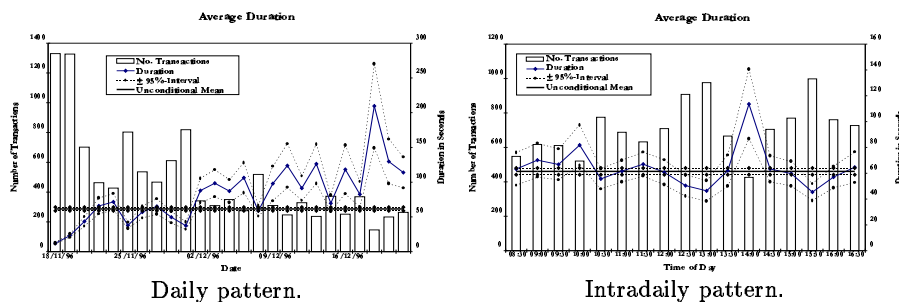


FIGURE 2: Daily and intradaily average of trade durations and number of transactions for Deutsche Telekom AG from Nov. 18, until Dec., 20, 1996.

and the average durations during successive half hour intervals as functions of the time of day for the whole sample.

Most of the trading appears to take place between 10.30 a.m. and 1.30 p.m. which is exactly the interval during which the Frankfurt stock exchange (FSE) operates regularly. In the interval beginning at 2.00 p.m. a significant increase of average trade durations is observed. This slowdown may be attributed to the closure of the FSE and the associated lunch time break. Another outburst of trade activity occurs after 3.30 p.m., when the NYSE starts to operate regularly. From then on, trading intensity seems to decrease until the end of the trading day at 5.00 p.m.

The variation of the average duration throughout the trading day for each of the four subsamples are displayed in Figure 3. The basic shape of the time of day function that was visible in the total sample figures is retained, although it operates on a different level for each of the subsamples.¹⁰

¹⁰Note that we omitted the 95% confidence intervals for the time of day function for sample D, because they exhibited very unstable behaviour due to the low number of observations.

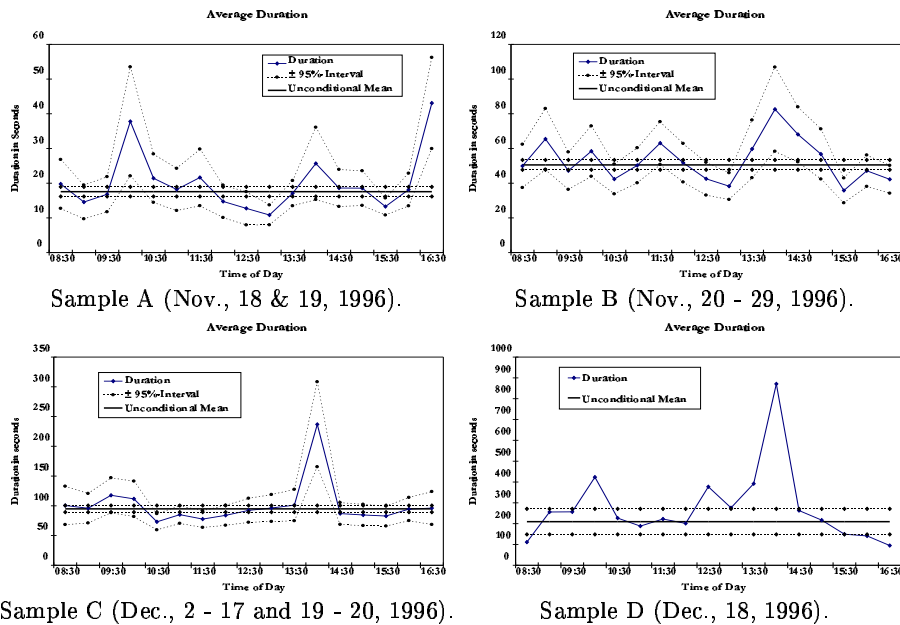


FIGURE 3: Intradaily average of trade durations for subperiods.

4.2. Modeling the deterministic intra day seasonality

The resolution of new information into prices on a continuous stock exchange does not happen all at once, but usually takes some time. In high frequency data sets one therefore often finds, that events in the trading process tend to occur in clusters. This implies that the waiting times between events during the intra-day trading and quoting process exhibit a significant autocorrelation. This is exactly what the ACD-model tries to depict. Some of these events, for example the opening and closing of major exchanges or lunch breaks, occur repeatedly in a deterministic manner thus giving rise to a periodic component. Hence, a part of the duration persistence is due to this intra-day seasonality.

As revealed by the sample statistics in the previous section, the trade duration series for the Deutsche Telekom AG does exhibit considerable time of day effects, and thus the question arises how to deal with them. Engle and Russell (1997) propose to decompose the duration series into a deterministic function of the time of day, $\Phi(t_i)$ and a stochastic component \tilde{x}_i , so that we have

$$(20) \quad x_i = \tilde{x}_i \cdot \Phi(t_i).$$

In this paper, we will apply the two step method proposed by Engle and Russell (1997) in which the time of day function function is estimated separately from the other model parameters.¹¹ Dividing each duration in the sample by the

¹¹ Simultaneous ML estimation would also be possible. Engle and Russell (1998) report that both procedures give similar results if sufficient data is available.

appropriate time of day function value, a sequence of deseasonalized durations is obtained, $\tilde{x}_i = \frac{x_i}{\Phi(t_i)}$.

In order to estimate the time of day function we employ the semi-nonparametric (SNP) estimator introduced by Gallant (1981). The basic approach here is to approximate some unknown functional $y = f(x)$ using a fourier series expansion accomodated by polynomials in the regressor variables. Gallant calls this a Fourier flexible form (FFF) approximation and suggests to use a polynomial of degree 2. Estimation in the univariate case is carried out by fitting a regression curve of the type

$$(21) \quad f(x_i; \theta_x) = \sum_{j=0}^J (\mu_j \cdot x_i^j) + \sum_{p=1}^P [\gamma_p \cdot \cos(p \cdot x_i) + \delta_p \cdot \sin(p \cdot x_i)],$$

where $0 \leq x_i < 2\pi$ and θ_x is the corresponding vector of parameters.

This type of estimator is especially well suited for our purposes, since it can reproduce the non-linear shape of the time of day function. Also it assumes that the regressor variable has bounded support, which is true in our application where the trading interval is limited to 8.5 hours per day.

Asymptotic normality and consistency of SNP-estimators for several types of data generating process with iid errors have been established in Eastwood (1991), and Andrews (1991). Whether these results hold for time dependent processes in general and the class of ACD-processes in particular is not known at the time being and thus remains open to further research. Nevertheless the technique has been applied by several researchers to estimate seasonal components in a GARCH-model.¹²

For our purposes the SNP-approach has to be modified as follows: Since the FFF is designed to approximate functions over the domain $[0, 2\pi]$, regressor variables have to be shifted and rescaled in order to meet the obligation. Let t_{\min} be the time of day at which trading begins and t_{\max} the corresponding closing time, then a suitable FFF approximation of the time of day function for trade durations may be expressed as follows

$$(22) \quad \tilde{x}_i = \sum_{j=0}^J [\mu_j \cdot h(t_{i-1})^j] + \sum_{p=1}^P [\gamma_p \cdot \cos(p \cdot h(t_{i-1})) + \delta_p \cdot \sin(p \cdot h(t_{i-1}))],$$

where the normalizing function $h(t)$ is given by

$$(23) \quad h(t) = 2\pi \cdot \frac{t - t_{\min}}{t_{\max} - t_{\min}}.$$

¹²See e.g. Andersen and Bollerslev (1997).

Our choice of the trimming parameters J and P was guided by the principle of parsimony, so that we selected the model that best matched the basic shape of the time of day function reported in the previous section with the lowest degree of parameters to be estimated.¹³ In our implementation of the SNP-estimator we found that including $J = 4$ polynomials of t generally yielded the best results. Lower order polynomials did not catch the shape of the time of day function, irrespective of the number of Fourier terms included, while higher order polynomials exhibited instable behaviour near the boundaries of the trading interval.

With respect to the choice of the Fourier lag truncation parameter P , we re-estimated the model with the given polynomial degree several times, successively increasing P . Doing so, we found that the shape changed relatively little for low order Fourier lags, but again exhibited instable behaviour during the middle part of the trading day for higher order lags and accordingly we choose the model with the lowest value for the Akaike (AIC) and Schwartz (SBC) model selection criteria. The preferred model included $J = 2$ Fourier lags. All in all it seems as though the polynomial terms do a good job for the fit at the boundaries of the trading interval, while the Fourier terms catch the pattern of variation in the middle part of the day.


Although we estimated a separate time of day function for each subperiod, we restricted the trimming parameters J and P to be the same for all subperiods, because of the obvious similarity of the basic shape of the function as indicated by the results in the previous section. The intradaily seasonality patterns for the four subsamples are depicted in Figure 4. Note that the shapes of the four time of day functions are generally in good accordance with the corresponding intradaily average durations.

This is confirmed by an inspection of the evolution of daily and intradaily average durations in Figure 5. Daily mean durations vary in a basically non-systematic manner around the unconditional mean of the \tilde{x}_i series which is equal to 1, while the intradaily averages all contain the unconditional mean within their 95%-confidence intervals. Also note that the overdispersion of the \tilde{x}_i series is equal to 1.50, which is about the same order of magnitude as in the first three subperiods, which contain the bulk of the data, while the overdispersion in the total sample had a larger value of 1.79 [see Table II].

Furthermore, an inspection of the autocorrelation function of the \tilde{x}_i series [Figure 6] reveals, that the time dependence dropped considerably as compared to the original duration series x_i . This is confirmed by the significant drop of the values of the Ljung-Box statistic at all lags.


We conclude, that our empirical specification for the time of day function was in general successful in removing both, the intradaily seasonality pattern as well as the apparent level shifts in the original series due to structural breaks occurring at several dates in our sample period. Nevertheless, the resulting

¹³This practice was inspired by a similar model selection strategy applied by Andersen and Bollerslev (1997).



Sample A (Nov., 18 & 19, 1996).

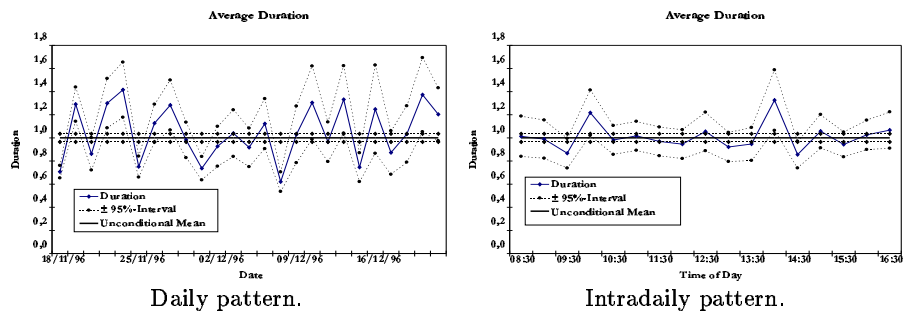
Sample B (Nov., 20 - 29, 1996).



Sample C (Dec., 2 - 17 and 19 - 20, 1996).

Sample D (Dec., 18, 1996).

FIGURE 4: Estimated intradaily seasonality pattern for subperiods.



Daily pattern.

Intradaily pattern.

FIGURE 5: Daily and intradaily averages of deseasonalized trade durations from Nov. 18, until Dec., 20, 1996.

TABLE II: Descriptive statistics for deseasonalized durations, total sample

Statistic	Deseasonalized
Number of observations	12038
Mean	1,00
Standard deviation	1,50
Overdispersion	1,50
Skewness	4,46
Kurtosis	40,92
Minimum	0,0
25 %-Quartile	0,2
Median	0,5
75 %-Quartile	1,2
Maximum	31,5

deseasonalized duration time series still exhibits considerably strong time dependence, which we will tend to analyze using the TACD model in the next section.

4.3. Estimation of the TACD model

In this section we will present estimation results for the TACD-model introduced by Zhang, Russell, and Tsay (1999) for the deseasonalized time series of transaction durations $\{\tilde{x}_i\}$ and compare it's forecast performance to the EACD and the WACD models of Engle and Russell (1998) using the method of Diebold, Gunther, and Tay (1997). All models are estimated employing the log ACD-specification of equation (7). We restrict our attention to the estimation of parsimonious ACD(1,1) specifications for all models.

We chose to estimate a three regime TACD(1,1) specification using \tilde{x}_{i-1} as the threshold variable. In accordance to Zhang, Russell, and Tsay (1999) we specified a generalized gamma distribution for the standardized duration series $\{\varepsilon_i\}$, allowing the shape parameter ϑ to vary across regimes but restrict the other shape parameter δ to be constant. The dynamics for the conditional mean are thus given by

$$(24) \quad \psi_i = \begin{cases} \exp(\omega^{(1)} + \alpha^{(1)} \ln(\tilde{x}_{i-1}) + \beta^{(1)} \ln(\psi_{i-1})) & \text{if } 0 < \tilde{x}_{i-1} \leq r_1 \\ \exp(\omega^{(2)} + \alpha^{(2)} \ln(\tilde{x}_{i-1}) + \beta^{(2)} \ln(\psi_{i-1})) & \text{if } r_1 < \tilde{x}_{i-1} \leq r_2 \\ \exp(\omega^{(3)} + \alpha^{(3)} \ln(\tilde{x}_{i-1}) + \beta^{(3)} \ln(\psi_{i-1})) & \text{if } r_2 < \tilde{x}_{i-1} < \infty \end{cases}$$

ML-estimates were obtained by maximizing the log likelihood function conditional on the regime parameters (r_1, r_2) , using the BHHH algorithm.¹⁴ In

¹⁴The estimates were computed using the ACD Gauss-code library *DTYFACD* which has

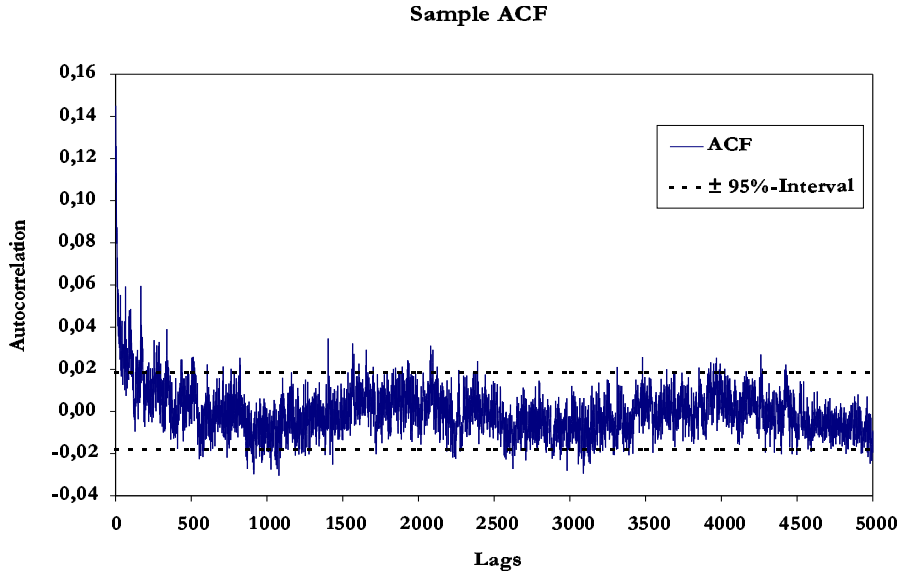


FIGURE 6: Sample autocorrelation function for deseasonalized trade durations from Nov. 18, until Dec., 20, 1996.

order to find the values of (r_1, r_2) , that maximize the sample log likelihood we perform a grid search. Estimation results are summarized in Table III.¹⁵

Note that all parameter values, except for the constant in the fast trading regime, $\omega^{(1)}$, are individually significant at the 5% level. The estimates reveal high persistence in the first and second regime as indicated by the sum of the parameters α and β parameters which are equal to 0,980 and 0,986. This is in contrast to the slow trading regime 3, which has a relatively low level of persistence. Note that the shape of the conditional density function also varies across regimes, as indicated by the estimates for the ϑ parameters.

Turning to the evaluation of the forecast performance, we will compare the TACD model to the ordinary EACD and WACD. The parameter estimates for these two models are presented in Table IV. Both models indicate a medium level of persistence, as compared to the three regimes TACD estimates. Also note that all parameters are individually significant at the 5% level. On the other hand, both models have significant lower values for the χ^2 -statistic, that we employ to test for uniformity on the interval $[0, 1]$ of the empirical cumulative forecast density.

As suggested by Diebold, Gunther, and Tay (1997), a closer look at histogram plots can help to identify possible causes of misspecification. The histogram plots for the three models in Figure 7 confirm that in general the TACD

been developed by L. Bauwens, P. Giot, J. Grammig, and D. Veredas.

¹⁵The χ^2 statistic referred to in Table III is the test statistic for the goodness of fit of the in-sample forecast density to the uniform distribution.

TABLE III: Estimation results for the TACD model

Parameter	Estimate	t-statistic
$\omega^{(1)}$	0,061	0,516
$\omega^{(2)}$	0,057	2,788
$\omega^{(3)}$	0,041	2,757
$\alpha^{(1)}$	0,069	2,073
$\alpha^{(2)}$	0,102	6,422
$\alpha^{(3)}$	0,217	11,088
$\beta^{(1)}$	0,911	34,854
$\beta^{(2)}$	0,884	49,006
$\beta^{(3)}$	0,566	12,667
$\vartheta^{(1)}$	5,332	8,119
$\vartheta^{(2)}$	7,462	8,166
$\vartheta^{(3)}$	7,028	8,121
δ	0,289	15,669
r_1	0,067	
r_2	0,668	
N	12038	
$\ln \mathcal{L}$	-10375,9	
χ^2	420,5	

TABLE IV: Estimation results for the EACD and WACD model

Parameter	EACD		WACD	
	Estimate	t-statistic	Estimate	t-statistic
ω	0,071	12,788	0,072	13,158
α	0,832	49,742	0,824	48,863
β	0,095	13,061	0,100	13,545
γ			0,812	154,435
N	12038		12038	
$\ln \mathcal{L}$	-11350,0		-10826,8	
χ^2	1386,5		583,5	

model performs best. The EACD model obviously has great problems to forecast trade durations accurately at the tails of the distribution as well as in the middle part. Both long and short durations are predicted by the EACD more often than they occur in reality, while the frequency of durations in the third and most of the fourth quartile of the distribution (excluding only the 95% quintile) is strongly underestimated.

A similar result is revealed by the histogram plot for the WACD model. Although it does a better job than the EACD in forecasting long durations, this comes at the cost of a very poor performance in the lowest 5% quintile. The TACD model yields far better predictions than the other two models over the whole support of the distribution, except for the first 5% quintile. Although the TACD model is rejected by the χ^2 statistic at conventional significance levels, the histogram plots show, that deviations from uniformity are far less severe than in the EACD and WACD cases. With the exception of the first and last 5% bins, all deviations from the 95% confidence bands plotted in Figure 7 are only minor in magnitude. Thus, while the TACD model clearly performs better than standard ACD models, there still seems to be some room for improvements.

5. CONCLUSIONS

We aimed to model the after market trading process of the Deutsche Telekom AG using transaction data for the first five weeks after it's initial public offering on Nov., 18 1996. We carefully investigated the overall evolution of the trading process during this time span, finding evidence for strong time dependence, several level shifts, which we interpreted as structural breaks and intradaily activity patterns, that are caused by periodic openings and closures of other asset markets during the regular trading interval of the IBIS trading system.

We employed the recently introduced class of threshold ACD models to analyze the time dependence of the trading process and compared it's performance to other established ACD models, including the exponential and the Weibull model. Although a three regime specification of the TACD model clearly outperformed the other models in terms of forecast accuracy, none of the models employed in this paper did pass the test for uniformity of the cumulative forecast density in the strong sense. This result is probably not unexpected, if one bears in mind that trade duration data is known to be very difficult to model accurately (unlike volume- or price-weighted financial duration processes)¹⁶, the more so when the share under consideration is going through the IPO-process. Our results indicate that the dynamics of the IPO after market trading process are subject to inherent non-linearities and that future research along these lines is necessary to understand it properly.

¹⁶See Bauwens, Giot, Grammig, and Veredas (2000).

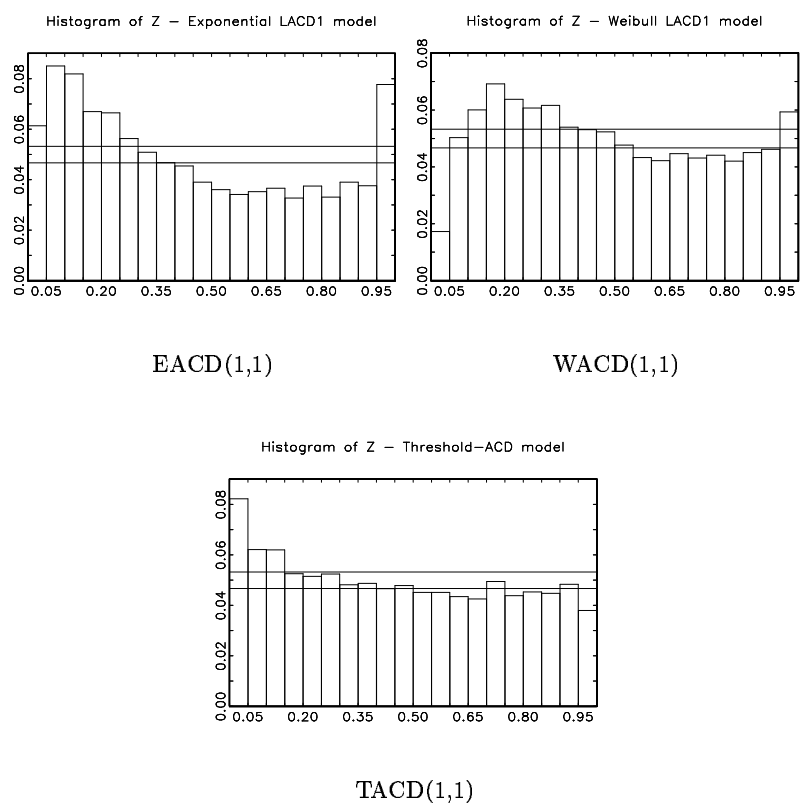


FIGURE 7: Histogram plots of the cumulative forecast density for the in-sample fit of the EACD, WACD and TACD model.

REFERENCES

- ANDERSEN, T. G., AND T. BOLLERSLEV (1997): "Intraday periodicity and volatility persistence in financial markets," *Journal of Empirical Finance*, 4(2-3), 115–158.
- ANDREWS, D. W. K. (1991): "Asymptotic normality of series estimators for nonparametric and semiparametric regression models," *Econometrica*, 59(2), 307–345.
- BAUWENS, L., AND P. GIOT (1997): "The logarithmic ACD model: An application to market microstructure and NASDAQ," Discussion paper, Université Catholique de Louvain.
- BAUWENS, L., AND P. GIOT (2000): "The moments of first-order Log-ACD models," Discussion paper, Université Catholique de Louvain.
- BAUWENS, L., P. GIOT, J. GRAMMIG, AND D. VEREDAS (2000): "A comparison of financial duration models via density forecasts," Discussion paper, Université Catholique de Louvain and University of Frankfurt.
- DIEBOLD, F. X., T. A. GUNTHER, AND A. S. TAY (1997): "Evaluating Density Forecasts," Discussion paper, University of Pennsylvania, Discussion Paper.
- EASTWOOD, B. J. (1991): "Asymptotic normality and consistency of seminonparametric regression estimators using an upwards F test truncation rule," *Journal of Econometrics*, 20(1), 151–181.
- ELLIS, K., R. MICHAELY, AND M. O'HARA (1999): "When the underwriter is the market maker: An examination of trading in the IPO aftermarket," *Forthcoming in Journal of Finance*.
- (2000): "The accuracy of trade classification rules: Evidence from Nasdaq," Working Paper, Johnson Graduate School of Management, Cornell University.
- ENGLE, R. F. (2000): "The econometrics of ultra high frequency data," *Econometrica*, 68(1), 1–22.
- ENGLE, R. F., AND J. R. RUSSELL (1997): "Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model," *Journal of Empirical Finance*, 4(2-3), 187–212.
- (1998): "Autoregressive conditional duration: A new model for irregularly spaced transaction data," *Econometrica*, 66(5), 1127–1162.
- GALLANT, A. R. (1981): "On the bias in flexible functional forms and an essentially unbiased form," *Journal of Econometrics*, 20(2), 285–323.

- GHYSELS, E., AND J. JASIAK (1998): “GARCH for irregularly spaced financial data: The ACD-GARCH model,” *Studies in Nonlinear Economics and Econometrics*, 2, 133–149.
- GOODHART, C. A. E., AND M. O’HARA (1997): “High frequency data in financial markets: Issues and applications,” *Journal of Empirical Finance*, 4(2-3), 73–114.
- GOURIÉROUX, C., AND J. JASIAK (1999): “Nonlinear Innovations and Impulse Response,” Discussion paper, CREST and York University, Discussion Paper.
- (2000): “Financial Econometrics,” Manuscript, York University, Toronto.
- GRAMMIG, J., R. HUJER, S. KOKOT, AND K.-O. MAURER (1998): “Modeling the Deutsche Telekom IPO using a new ACD specification: An application of the Burr-ACD model using high frequency IBIS data,” SFB373 Discussion Paper 55, Humboldt-Universität zu Berlin.
- GRAMMIG, J., AND K.-O. MAURER (2000): “Non-monotonic hazard functions and the autoregressive conditional duration model,” *Forthcoming in Econometrics Journal*.
- GRAMMIG, J., AND M. WELLNER (1999): “Modeling the interdependence of volatility and inter-transaction duration processes,” SFB373 Discussion Paper 21, Humboldt-Universität zu Berlin.
- HAMILTON, J. (1994): *Time series analysis*. Princeton University Press, Princeton.
- IBBOTSON, R. G., AND J. R. RITTER (1995): “Initial public offerings,” in *Handbooks in Operations Research and Management Science Vol. 9: Finance*, ed. by R. A. Jarrow, V. Maksimovich, and W. T. Ziemba, pp. 993–1016. Elsevier, Amsterdam.
- LANCASTER, T. (1990): *The Econometric Analysis of Transition Data*. Cambridge University Press, Cambridge.
- LJUNGQVIST, A. P. (1997): “Pricing initial public offerings: Further evidence from Germany,” *European Economic Review*, 41(7), 1309–1320.
- LOUGHRAN, T., J. R. RITTER, AND K. RYDQVIST (1994): “Initial public offerings: International insights,” *Pacific-Basin Finance Journal*, 2(2), 165–199.
- LÜDECKE, T. (1996): “The Karlsruher Kapitalmarktdatenbank (KKMDB): The Kiss-Tape,” Discussion Paper 191, Institut für Entscheidungstheorie und Unternehmensforschung, University of Karlsruhe.

- LUNDE, A. (1999a): “A conjugate gamma model for durations in transactions data,” Discussion paper, Aalborg University.
- (1999b): “A Generalized Gamma Autoregressive Conditional Duration model,” Discussion paper, Aalborg University.
- MADHAVAN, A. (2000): “Market microstructure: A survey,” *Journal of Financial Markets*, 3(3), 205–258.
- NELSON, D. B. (1991): “Conditional heteroskedasticity in asset returns: A new approach,” *Econometrica*, 59(2), 347–370.
- NELSON, D. B., AND C. Q. CAO (1992): “Inequality constraints in the univariate GARCH model,” *Journal of Business & Economic Statistics*, 10(2), 229–235.
- ROSENBLATT, M. (1952): “Remarks on a multivariate transformation,” *Annals of Mathematical Statistics*, 23(3), 470–472.
- SCHLAG, C., AND A. WODRICH (2000): “Has there always been underpricing and long-run underperformance? - IPOs in Germany before world war I,” Department of Economics and Business Administration, University of Frankfurt.
- SCHUSTER, J. A. (1996): “Underpricing and crisis - IPO performance in Germany,” Discussion paper 0252, Financial Markets Group, London School of Economics.
- STEHLE, R., AND O. EHRHARDT (1999): “Renditen bei Börseneinführungen am deutschen Kapitalmarkt,” *Zeitschrift für Betriebswirtschaft*, 69(12), 1395–1422.
- TONG, H., AND K. S. LIM (1980): “Threshold autoregression, limit cycles and cyclical data (with discussion),” *Journal of the Royal Statistical Society, Series B*, 42(3), 245–292.
- UHLIR, H. (1989): “Der Gang an die Börse und das Underpricing-Phänomen,” *Zeitschrift für Bankrecht und Bankwirtschaft*, 1(1), 2–16.
- WASSERFALLEN, W., AND C. WITTELEDER (1994): “Pricing initial public offerings: Evidence from Germany,” *European Economic Review*, 38(7), 1505–1517.
- ZHANG, M. Y., J. R. RUSSELL, AND R. S. TSAY (1999): “A nonlinear autoregressive conditional duration model with applications to financial transaction data,” Discussion paper, Graduate School of Business, University of Chicago.