

Trade intensity in the Russian stock market: dynamics, distribution and determinants

Stanislav Anatolyev*

Dmitry Shakin

New Economic School

Academy of National Economy, Russia

Abstract

We investigate the distribution and evolution of intertrade durations for frequently traded stocks at the Moscow Interbank Currency Exchange. We use a flexible econometric model based on ARMA and GARCH which, when coupled with a certain class of distributions that allow for skewness and slim-tailedness, adequately captures the characteristics of conditional distribution of durations for Russian stocks, and is able to generate high quality density forecasts. We also analyze what factors determine the dynamics of log-durations and in which way. The results in particular indicate that the Russian market is characterized by aggressive informed traders and timid liquidity traders, and that the participants react evenly to upward and downward short-run price trends.

Key words: High frequency data; Trading intensity; Intertrade durations; ACD model; ARMA–GARCH model; Market microstructure.

JEL classification numbers: C22, C41, G10, G15.

*Corresponding author. Address: Stanislav Anatolyev, New Economic School, Nakhimovsky Prospekt, 47, Moscow, 117418 Russia. E-mail: sanatoly@nes.ru. We would like to thank an anonymous referee of this journal for useful suggestions that helped greatly improve the presentation, Konstantin Tyurin for helpful discussions, and Georgy Kolodyazhny for kindly providing the data on Russian stocks.

1 Introduction

Modeling high frequency (HF) transaction data has been an intensive topic of research in recent decades. The seminal articles of Engle and Russell (1998) and Engle (2000) generated a burst of new interest in modeling the dynamics of duration and price processes and gaining understanding of market microstructure. In the applied literature, the most frequent analysis is of IBM stocks at the NYSE (e.g., Engle and Russell, 1998, Engle, 2000, McCulloch and Tsay, 2001, Zhang, Russell and Tsay, 2001), with different studies often using the same database. Some exploit data for other stocks at the NYSE (Dufour and Engle, 2000, Bauwens and Giot, 2000, Engle and Lunde, 2003). Exploration of other stock exchanges is much rarer and does not go unnoticed. Several papers (e.g., Ghysels, Gouriéroux and Jasiak, 2004, Drost and Werker, 2004) use data from the Paris Stock Exchange (Paris Bourse); Grammig and Wellner (2002) use data from the Frankfurt Stock Exchange; and Tyurin (2002) uses data from the Reuters D2000-2 foreign exchange electronic brokerage system. Papers that study high frequency features of emerging stock markets are rare but do exist. For instance, Hanousek and Podpiera (2003) explore the impact of informed trading on the bid-ask spread using HF data from the Czech stock market. To our knowledge, there is no econometric study analyzing HF data from the Russian market, although a couple of papers do look at the microstructure of Russian financial markets: Medvedev and Kolodyazhny (2001) investigate the behavior of nonresidents in the state bond market and find evidence of herding behavior; the same authors (2003), henceforth M&K, explore possibilities of profitable market making, and inspect relationships between brokers and their clients. In this paper, we analyze tick-by-tick data on trades in six frequently traded common stocks at the Moscow Interbank Currency Exchange (MICEx) focusing attention on the dynamics, distribution and determinants of intertrade durations.¹

A careful analysis of the dynamics of intertrade durations shows that the classical autoregressive conditional duration (ACD) model (Engle and Russell, 1998) and its logarithmic modifications (Bauwens and Giot, 2000) fit the Russian market intratrade durations data worse than a simple specification applied to log-durations and based on ARMA and GARCH, which we call *Log-ARMA-GARCH*. In addition, this econometric model is more flexible, has desirable theoretical properties, and at the same time is computa-

¹Other segments and aspects of Russian financial markets are studied in, among others, Fedorov and Sarkissian (2000), Anatolyev and Korepanov (2003), and Hayo and Kutan (2005).

tionally simpler than complicated models that have appeared recently. Using the class of so-called Skewed Generalized Error (SGE) distributions (Bali and Theodossiou, 2003) to model the conditional distribution within the Log-ARMA-GARCH framework allows one to adequately capture slight right skewness and slim-tailedness of log-durations of the MICEx stocks. In addition, density forecasting experiments (as in Bauwens, Giot, Grammig, and Veredas, 2004) show that the SGE Log-ARMA-GARCH model performs out-of-sample no worse than some preferable ACD specifications. We also analyze what factors determine the dynamics of log-durations and in which way. From this analysis, it appears that the Russian stock market is characterized by aggressive informed traders and timid uninformed traders, and that the participants react evenly to upward and downward short-run price trends. Throughout, along with the Russian stocks, we analyze the Alcatel stock at the Paris Bourse for comparison purposes. Our results show that the behavior of Russian traders is quite similar to that of traders in developed financial markets, but the SGE Log-ARMA-GARCH framework may not be so successful for developed markets as for the Russian market.

The paper is organized as follows. The description of the MICEx and the database we use is given in section 2. An inspection of the data reveals that, along with intraday deterministic patterns, there are also interday variations in intertrade durations, returns, and volumes. In section 2, we also show how we remove both sorts of nonstationarity from the data prior to the analysis, and we discuss the properties of adjusted durations. Section 3 is devoted to empirical analysis of evolution of intertrade durations, including modeling the conditional distribution and performing the density forecasting exercise. In section 4, we analyze factors that influence trading intensity, and we draw some conclusions on the microstructure of the Russian financial market. Finally, section 5 concludes.

2 The MICEx and data

2.1 Russian stock market and MICEx

The organized stock market in Russia is composed of several stock exchanges, two of which, MICEx and RTS (short for “Russian Trading System”), account for more than 95 percent of trade turnover, with the share of MICEx being near 80 percent.² In 2002,

²A brief introduction to the Russian stock market can be found in Ostrovsky (2003). Details on the MICEx are available in English at www.micex.com.

the volume of stocks of more than 100 Russian issuers transacted on the MICEEx reached about \$39 billion. Most of these stocks are traded very rarely, but several blue chips are traded at a frequency of from 200 to 6,000 transactions a day. The MICEEx is evidently quite active for an Eastern European market compared, for example, with the Czech stock market, where the daily turnover is \$19 million, with most liquid stocks being traded at 67 trades per day (Hanousek and Podpiera, 2003).

The electronic system at the MICEEx is organized as a typical order-driven market.³ The MICEEx takes advantage of technological innovations to enable its members to trade via the Internet. (In 2002, more than 70 percent of trades were made from 1,500 remote workstations in 50 different cities.) The players primarily represent Russian investors; the percentages of American and European investors are relatively small. Along with the “normal” trading regime with continuous matching of orders, there are special regimes of trades in the same security. Regimes and their brief descriptions are presented in Table 1. As at other exchanges, participants in trades are individual members (brokers) who trade either for themselves or on behalf of their customers, which include more than 500 banks and financial companies. M&K roughly estimate that 43 percent of brokers trade for themselves, 44 percent have up to ten customers, and the rest have more numerous clients. There are no official market makers for liquid stocks, although there may be appointed market makers for bonds and low liquidity stocks like specialists at the NYSE, where orders are almost never executed automatically. However, participants who set limit orders that are not immediately executed are *de facto* market makers, although none of them profit all the time solely from intermediation between buyers and sellers (see M&K). M&K roughly estimate that 16 percent of transactions are made by market makers (those who sell and buy the same number of shares during a day), and 47 percent are made by pure buyers or sellers (those who only sell or only buy during a day). M&K also draw the interesting conclusion that there are attractive opportunities for market making at the MICEEx, which allow earning higher returns with lower risk, but the investor body is quite heterogeneous so that only a portion of market makers can enjoy excessive gains.

2.2 MICEEx database

Our sample covers the period from August 12, 2002, to October 18, 2002, composed of 50 trading days, and includes six Russian corporations whose common stocks were

³An excellent explanation of an order-driven system can be found in Tyurin (2002).

most frequently traded during this period.⁴ Among these six companies, three are oil extractors (Lukoil, SurgutNG, and Yukos); MosEnergO is a large electricity provider; Rostelecom is a leading Russian telecommunications company; Sberbank is the largest Russian commercial bank. For some corporations, there are several issues of common stocks as well as preferred stocks and bonds; we focus on only one security for each corporation, the issue that is traded much more frequently than others. A brief description of stocks is presented in Table 2. The column “stock share” shows shares of each stock turnover in the total turnover in stocks at the MICEx during the period from August 19, to August 23, 2002. We concentrate on the “trading session,” or “normal,” trading regime. All other observations are excluded from the sample, i.e., trades that are recorded out of trading hours for the “normal” regime (which opens at 10:30 a.m. and closes at 6:45 p.m.), and trades that belong to the “negotiated deals,” “REPO,” or “incomplete lots” regimes. As a consequence, the Lukoil, SurgutNG, Rostelecom, and MosEnergO samples are left with about 600 to 750 transactions a day; for Yukos we have 400 transactions a day, and “only” 250 trades for Sberbank.

Each record in our database corresponds to one transaction for which the following three characteristics are available:

1. Time T_t (with the precision of one second) when the transaction occurred;
2. Price P_t at which the transaction occurred;
3. Number of shares V_t bought or sold.

The database does not contain quote data. Using the available characteristics we generate the following three variables:

1. Duration between consecutive trades: $\hat{d}_t = T_t - T_{t-1}$;
2. Log return: $\hat{r}_t = \ln P_t - \ln P_{t-1}$;
3. Log volume: $\hat{v}_t = \ln V_t$.

When generating the variables we pay special attention to simultaneous transactions. Such transactions occur in two ways. First, simultaneous transactions are recorded when a new buy (sell) order matches several sell (buy) orders and is in sufficient amount to execute all these matching orders. It is natural to treat such trades as one big transaction. A characteristic of simultaneous transactions of this type is that their prices are either non-increasing or non-decreasing. The second way is that the duration between consecutive

⁴The most heavily traded stock at the MICEx is that of Unified Energy System of Russia (UES), with an average of about 6,000 transactions a day. We do not take these data into analysis because of a huge number of simultaneous transactions (see below).

trades is so small that the precision of one second is not sufficient for their discrimination and they are marked with the same time stamp. Because we cannot distinguish between the two sources, we use the following procedure. If there are several transactions recorded at the same time, we aggregate transactions with non-increasing or non-decreasing prices, using for these aggregated transactions a weighted average price and aggregate volume, and remove all remaining transactions with this time stamp.⁵

The Alcatel sample analyzed throughout for comparison purposes originates in the Paris Bourse, a similarly organized market equipped with an electronic order matching device, and has been used in a number of studies; see Jasiak (1999), Ghysels, Gouriéroux and Jasiak (2004), Drost and Werker (2004) for a detailed description.⁶ These data cover 43 days in July and August 1996. The aggregation of simultaneous transactions is performed as described above.

2.3 Seasonal data adjustment

Most of the previous literature studying high frequency transactions data ignores inter-day variation in variables. We find that on different days the transaction activity and price volatility at the MICEx fluctuate noticeably, which may be caused by day-of-week deterministic patterns (on average, the durations are larger on Mondays and Fridays than on other weekdays) and interday variations in the number of market players. Figure 1 depicts plots of average durations across trading days for Lukoil and Alcatel.⁷ For return volatilities and volumes, the effects of interday variations (not shown) are also present, more clearly for returns than for volumes. One way to deal with these fluctuations is to create a dummy variable for each day, akin to fixed time effects in the panel data analysis, or to apply the model to each day separately. Another approach is taken in the hierarchical model of McCulloch and Tsay (2001) where the parameters are modeled as random variables drawn from some distribution taking new values every day, akin to random time effects in the panel data analysis. We instead remove interday variations from the data

⁵This deletion is relatively innocuous because the number of such situations is small for all stocks (except very heavily traded UES which we excluded from the sample). Of course, a final user of our analysis may not necessarily be most interested in data aggregated in this way because the aggregation can potentially change some of its time series properties.

⁶This dataset is publicly available from the *Journal of Business & Economic Statistics* data website.

⁷Note that trading intensity for Alcatel seems to be subject to medium-run (i.e. a few week long) ups and downs, while that for Lukoil apparently is not.

from the outset in the following way:

$$\begin{aligned}\tilde{d}_t &= \frac{\hat{d}_t}{\bar{\hat{d}}_s}, \\ \tilde{r}_t &= \frac{\hat{r}_t}{\sqrt{\bar{\hat{r}}_s^2}}, \\ \tilde{v}_t &= \hat{v}_t - \bar{\hat{v}}_s,\end{aligned}$$

where $\bar{\hat{d}}_s$ is the average duration for day s if observation t belongs to day s , $\bar{\hat{r}}_s^2$ is the average squared return for day s , and $\bar{\hat{v}}_s$ is the average log volume for day s . For durations and returns we use a multiplicative form of adjustment as durations and squared returns are naturally positive. However, for log volumes we use additive adjustment as this variable may take any sign; this adjustment implies the multiplicative form for volumes in levels as well.

In addition to interday variations in the data, there are specific intraday patterns of a deterministic nature in durations, returns, and volumes. It is conventional in the literature to remove such patterns prior to estimation by using various nonparametric methods. Most authors (e.g., Engle, 2000, Grammig and Wellner, 2002) use cubic or piecewise linear splines, while some authors estimate regressions based on kernel methods (e.g. Zhang, Russell and Tsay, 2001). Even more rarely they include diurnal dummy variables (as in Dufour and Engle, 2000), or remove the seasonal component by a linear regression on time (as in Ghysels, Gouriéroux and Jasiak, 2004) or by simple averaging over a moving window and linear interpolation (as in Bauwens and Giot, 2000). To identify and estimate U-shaped patterns one may also use parametric specifications for their shape (as in Aradhyula and Ergün, 2004, Panas, 2005).

Here we estimate means \tilde{d}_t , \tilde{r}_t^2 and \tilde{v}_t conditional on time-of-the-day T_t by fitting local linear regressions⁸. Moreover, we do not assume that the intraday patterns are the same across different days of the week but do estimation for each of them separately. This is also different from most of previous studies, some of which allow but do not detect weakly seasonal components, e.g., Bauwens and Giot (2000), although Grammig and Wellner (2002) and Bauwens, Giot, Grammig, and Veredas (2004) also condition both on time-of-day and on day-of-week. An extreme position is taken in Meitz and Teräsvirta (2003), where splines are fit separately for each day of the sample.

⁸The gaussian kernel is employed, and bandwidths are manually selected. The global bandwidth equals 2000 seconds for durations and return volatilities, and 3000 seconds for log-volumes.

The estimates of intraday patterns for durations of Lukoil stocks are presented in Figure 2. The pattern is consistent with previous studies: the durations are longer at mid-day and shorter in the morning and in the evening. Figure 3 plotting intraday patterns together with confidence bands⁹ for two weekdays confirms that day-of-week differences are statistically significant: the two confidence bands do not overlap except during short time intervals. We use the estimates of intraday patterns to generate final adjusted versions of our variables:

$$\begin{aligned} d_t &= \frac{\tilde{d}_t}{\hat{E}[\tilde{d}_t|T_t]}, \\ r_t &= \frac{\tilde{r}_t}{\sqrt{\hat{E}[\tilde{r}_t^2|T_t]}}, \\ v_t &= \tilde{v}_t - \hat{E}[\tilde{v}_t|T_t]. \end{aligned}$$

Table 3 reports descriptive statistics for all stocks under consideration, and Figures 4 and 5 show nonparametric kernel density estimates¹⁰ of (adjusted) durations and log-durations for Lukoil and, for comparison purposes, Alcatel data. One sees a bit lower variability of durations and their autocorrelation for Alcatel than for Russian stocks, as well as a higher degree of kurtosis. Distributions of durations for Lukoil and Alcatel look similarly. A much starker contrast can be observed from the statistics and graphs for log-durations. While the distribution of log-durations for a typical Russian stock resembles normal (although slightly slimmer-tailed and right-skewed), those for Alcatel exhibit much higher skewness. Some of the skewness is certainly due to the fact that very short durations (between 0 and 1 seconds) were rounded to 0 or 1 when recorded in the database. Finally, note that the autocorrelation coefficients are higher for log-durations than for durations. In our view, this is an indication that an econometric model that tries to capture dynamics of log-durations promises to be more successful than a similar model for durations in levels.

⁹Let $g_1(T_t) = E[y_t|T_t]$ and $g_2(T_t) = E[y_t^2|T_t]$. Confidence intervals for $g_1(T_t)$ are constructed as $\hat{g}_1(T_t) \mp 1.96\sqrt{(\hat{g}_2(T_t) - \hat{g}_1(T_t)^2) \int k(u)^2 du / (nh\hat{f}(T_t))}$, where \hat{g}_1 and \hat{g}_2 are local linear regression estimates of g_1 and g_2 , \hat{f} is kernel density estimate of T_t , k is kernel function, h is bandwidth, n is sample size.

¹⁰The gaussian kernel is employed, except for durations on the segment $[0,1]$, in which case the Beta kernel (Chen 2000) is used because of boundedness of the support from below, with the smoothing parameter $b = 0.089$ and the global bandwidth of $h = 0.254$ as suggested in Chen (2000).

3 Time-series models for durations

3.1 ACD and Log-ARMA models

The class of autoregressive conditional duration (ACD) models was proposed by Engle and Russell (1998). The idea of an ACD model is to capture the persistence in durations, a stylized feature of high frequency transactions data, by means of an autoregressive model for conditional means. Let d_t be the duration between consecutive trades occurring at times T_{t-1} and T_t , which is assumed to be strictly greater than zero. This duration is factored into the predictable and unpredictable components:

$$d_t = E[d_t|I_{t-1}]\varepsilon_t \equiv \pi_t\varepsilon_t, \quad (1)$$

where I_{t-1} is the information set embedding all previous durations and possibly other variables, and ε_t is a shock having mean of unity conditional on I_{t-1} . Engle and Russell (1998) use the following functional form of evolution of π_t :

$$\pi_t = \omega + \sum_{i=1}^p \phi_i d_{t-i} + \sum_{i=1}^q \psi_i \pi_{t-i}. \quad (2)$$

This specification is called ACD(p, q). To complete the parametric model, one specifies the distribution for ε_t , which can be done in a variety of ways. Engle and Russell (1998) used the exponential and Weibull distributions. These choices, however, imply an unrealistic monotonic conditional hazard function, which led Lunde (1999) to propose the generalized gamma distribution, and Grammig and Maurer (2000) the more exotic Burr distribution. When the density is specified, one can proceed with the maximum likelihood (ML) estimation. When the above distributions do not fit the data well, one may treat the ACD model as semiparametric, and run quasi-maximum likelihood (QML) estimation presuming the exponential or (ordinary) gamma density for ε_t , with robust computation of standard errors. The ACD model also allows a researcher to investigate the influence of external factors by introducing additional regressors in equation (2).

An arguably more suitable parameterization for dynamics of π_t may be two versions of the LogACD(p, q) model (Bauwens and Giot, 2000), where

$$\pi_t = \exp \left(\omega + \sum_{i=1}^p \phi_i \ln d_{t-i} + \sum_{i=1}^q \psi_i \ln \pi_{t-i} \right) \quad (3)$$

or

$$\pi_t = \exp \left(\omega + \sum_{i=1}^p \phi_i \varepsilon_{t-i} + \sum_{i=1}^q \psi_i \ln \pi_{t-i} \right). \quad (4)$$

As Bauwens and Giot (2000) argue, these specifications have the advantage that the conditional mean π_t is always positive, so there is no need to impose non-negativity constraints either on the coefficients ϕ_i and ψ_i , or on the coefficients belonging to additional structural variables on the right sides of (3) or (4). In addition, Bauwens and Giot (2000) find that (4) fits their data better than the ACD specification (2). Subsequently, we denote the specifications (3) and (4) by LogACD_1 and LogACD_2 , respectively.

Although ACD models are able to capture high persistence in the duration process well, the major shortcoming of models of this type is that the conditional mean and variance are tied to each other:

$$V[d_t|I_{t-1}] = \pi_t^2 V[\varepsilon_t],$$

where $V[\varepsilon_t]$ is a constant depending on the error distribution. To separate out the persistence in mean and that in variance, Ghysels, Gouriéroux and Jasiak (2004) propose a stochastic volatility duration (SVD) model. They argue that their model is successful in capturing these two different sources of persistence. However, the likelihood function has a complicated form due to the presence of latent factors, and thus simulation-based techniques have to be applied for estimation and inference.

Now consider the following $\text{ARMA}(p, q)$ representation for log durations:

$$\ln d_t = \alpha + \sum_{i=1}^p \beta_i \ln d_{t-i} + \sum_{i=1}^q \gamma_i \zeta_{t-i} + \zeta_t, \quad (5)$$

where $\zeta_t = \sigma_t \eta_t$ follows $\text{GARCH}(r, s)$:

$$\eta_t \sim \text{IID}(0, 1), \quad (6)$$

$$\sigma_t^2 = \mu + \sum_{i=1}^r \lambda_i \zeta_{t-i}^2 + \sum_{i=1}^s \nu_i \sigma_{t-i}^2. \quad (7)$$

We will call this specification the $\text{Log-ARMA}(p, q)\text{-GARCH}(r, s)$ model, or Log-ARMA for brevity. As before, we have two options to complete the model. One is to specify the distribution of η_t , and proceed with ML estimation. The second option is to view the model as semiparametric and run QML using the normal density.

That the Log-ARMA specification (5)–(7) has not been seriously used in the previous literature¹¹ we attribute to the fact that it does not explicitly acknowledge the nature of

¹¹In fact, McCulloch and Tsay (2001) consider a very simple version of the Log-ARMA specification in their model for price changes and durations. Their equation for a price duration (i.e., a duration between trades with a nonzero price change) is a homoscedastic autoregression for log-durations with an additional explanatory variable, the lagged price change. Surprisingly, McCulloch and Tsay (2001) find no or weak dynamic dependence in time durations for IBM data.

d_t as a duration variable.¹² At the same time, this specification is natural to consider, it is flexible, and its properties are familiar; in Appendix it is shown how the Log-ARMA and ACD-type specifications are algebraically related. In addition, the Log-ARMA model fits the data we are using much better than ACD-type models, while the degree of parsimony is comparable. The main advantage of the Log-ARMA model lies in its greater flexibility, which originates from separation of the persistence in conditional mean from that in conditional variance. In particular, a time-varying conditional variance in the Log-ARMA model can take account of the dynamics of overdispersion in transaction data. Gouriéroux and Jasiak (2000, p. 462) argue that the overdispersion in real data is time-varying. While in the ACD model the conditional overdispersion is constant (recall that $V[d_t|I_{t-1}]/\pi_t^2 = V[\varepsilon_t]$), in the Log-ARMA it is path dependent. Moreover, the coefficients in the variance equation (7) provide an idea about clustering of observations with over- or underdispersion. At the same time, in spite of flexibility of this model, a practitioner may easily use standard statistical software to estimate the parameters in (5)–(7).

Another interesting feature of the Log-ARMA model is the type of implied conditional hazard function of the duration process. Grammig and Maurer (2000) come to the conclusion that allowing for non-monotonic hazard functions is an important issue in modeling durations. Even in a conditionally normal homoscedastic Log-ARMA model, the conditional distribution of durations is lognormal and thus the conditional hazard function is non-monotonic, increasing for small durations and decreasing for larger durations, which is consistent with estimated patterns (e.g., Grammig and Maurer, 2000, Figure 5). Moreover, in a heteroscedastic Log-ARMA model, the hazard function is time-varying, and it will be even more flexible after the assumption of conditional normality is relaxed.

Because the summary statistics and nonparametric estimates of unconditional log-duration distributions suggest the presence of platykurticity (slim-tailedness) and skewness, we in addition use the Skewed Generalized Error (SGE) distribution introduced in Bali and Theodossiou (2003) instead of the standard normal distribution to fit to the standardized innovations η_t . This distribution is an extension of the GED¹³ that captures

¹²Indirect evidence of the convenience in using log-durations as a variable to be modeled is contained in Meitz and Teräsvirta (2003), who set lagged values of $\ln d_t$ to be transition variables in their STACD specification.

¹³GED is an acronym for “Generalized Error distribution.” It was used, for example, by Nelson (1991) in an EGARCH model of stock returns.

leptokurtic, platykurtic, and mesokurtic distributions, and in addition allows nonzero skewness. The density of the SGE distribution is

$$f_{SGE}(\eta; \kappa, \varphi) = \frac{\kappa}{2\Lambda\Gamma(1/\kappa)} \exp\left(-\left|\frac{\eta - \Delta}{\Lambda(1 + \operatorname{sgn}(\eta - \Delta)\varphi)}\right|^\kappa\right),$$

where

$$\Lambda = \frac{1}{B} \sqrt{\frac{\Gamma(1/\kappa)}{\Gamma(3/\kappa)}}, \quad \Delta = -\frac{2\varphi A}{B}, \quad A = \frac{\Gamma(2/\kappa)}{\sqrt{\Gamma(1/\kappa)\Gamma(3/\kappa)}}, \quad B = \sqrt{1 + (3 - 4A^2)\varphi^2},$$

$\kappa > 0$ is a tail-thickness parameter ($\kappa > 2$ implying platykurticity), and $-1 < \varphi < 1$ is a skewness parameter (negative φ implying skewness to the left).

3.2 Estimation results

The estimation of all models is performed using Gauss v. 6.0.8 with the maximum likelihood library `maxlik` v. 5.0.3. For numerical optimization, the BFGS and BHHH algorithms are utilized; standard errors are computed using numerical derivatives. All autoregressive processes are reinitialized at the beginning of each day, as in, for example, Engle and Russell (1998) and Zhang, Russell and Tsay (2001); hence, the duration between the last “normal regime” trade of one day and the first “normal regime” trade of the following day never contributes to the likelihood.

To judge the quality of fit, we use several measures. Ljung–Box tests for (standardized) residuals and their squares, and values of sample mean loglikelihood together with Akaike and Schwarz information criteria (AIC and BIC), are standard in empirical implementation of HF data models. It is sometimes reported that parameters in HF data models are highly unstable; the Nyblom test (Nyblom, 1989) is a natural portmanteau check for structural stability when the estimation framework is the method of maximum likelihood. We report the Nyblom statistic for a whole model and the maximal value of Nyblom statistics for individual parameters. Also, taking into account huge sample sizes, we treat diagnostic testing results liberally and display reasonable tolerance to rejections by portmanteau statistics.

Table 4 presents the results of exponential QML estimation of ACD(2,2), LogACD₁(2,2), LogACD₂(2,2), and Log–ARMA(1,2)–GARCH(1,1) models for Lukoil as a representative Russian stock (the results for other stocks are similar), while Table 5 contains analogous output for Alcatel. The loglikelihood values are quite close across the three ACD-type models, but there is an appreciable gain in loglikelihood values and reduction in information criteria for the Log–ARMA model in comparison with ACD-type models. The

LogACD₂ model tends to reduce the autocorrelation in raw and squared residuals slightly better than the ACD model and much better than the LogACD₁ model. Note also that in the Log-ARMA model the variance persistence indexed by $\lambda + \nu$ far exceeds the mean persistence indexed by $(\beta + \gamma_1 + \gamma_2) / (1 + \gamma_1 + \gamma_2)$. This indicates the need to separate the mean and variance dynamics, which is facilitated by the use of Log-ARMA models. This model does seem ideal for the Russian stocks as any autocorrelation in residuals and their squares is removed, and the model exhibits high stability (unlike in case of Alcatel) despite huge sample sizes. The latter fact may be explained in the following way. At the MICEx the non-informed traders know that informed traders are constantly learning and exploiting news, and still trade mimicking their behavior, trying to benefit from the overall growing market. As a result, an essential piece of news, while having an impact on the dynamics of trade intensity, does not influence the style of traders' behavior in the Russian stock market, and the trading process goes in the same regime.

Because KS statistics indicate the inadequacy of the normal conditional distribution for standardized log-durations in the Log-ARMA model (although in a much lesser degree than of the exponential conditional distribution in ACD-type models), we estimate the Log-ARMA(1,2)-GARCH(1,1) assuming the SGE distribution for standardized innovations. The results are presented in Table 6. One can see that the SGE-Log-ARMA model fits the Russian duration data very well, removing all autocorrelation in standardized residuals and their squares, and exhibiting a high degree of stability (except the least liquid Sberbank stock). In addition, KS test statistics are much lower than for other combinations of models and conditional distributions¹⁴, which means that the SGE distribution captures well the skewness and platykurticity of conditional log-durations. The SGE-Log-ARMA also fits Alcatel durations quite well as far as serial dependence and conditional density are concerned, but some dynamic features (for example, regime switches) are left unaccounted for, which is indicated by extremely high values of both stability statistics.

3.3 Density forecasting results

In addition to in-sample fitting, we also verify an ability of the Log-ARMA model to provide reliable out-of-sample forecasts and compare them with those provided by ACD

¹⁴This follows from non-reported results of adapting the three ACD-type models to the Weibull, Generalized Gamma, and Burr distributions.

models. We follow the algorithm of evaluating density forecasts proposed by Diebold, Gunther, and Tay (1998) and adapted by Bauwens, Giot, Grammig, and Veredas (2004) to duration data. In a nutshell, to test the specification of conditional distribution, one may generate a sequence of probability integral transforms

$$z_t = \int_{-\infty}^{d_t} f(u) du,$$

where $f(u)$ denotes the conditional one-step-ahead density forecast, and verify if the sequence z_t is serially independent and distributed uniformly on $[0, 1]$. Appropriate testing tools include Pearson's goodness-of-fit statistic for uniformity and Ljung–Box statistics for serial uncorrelatedness of the z -sequence and of its squares; in addition, visual inspection of histograms of z is able to provide informal evidence about possible causes of model failure (for more details, see Bauwens, Giot, Grammig, and Veredas, 2004).

We implement the density forecasting exercise by re-estimating each model of interest (including making intraday seasonal adjustment) on the data from the first 35 days (28 days in case of Alcatel), and collecting forecast statistics during the remaining 15 days. Table 7 contains Pearson and Ljung–Box test statistics for selected Russian stocks, Lukoil among them, and Alcatel; the models of interest are ACD, LogACD₁ and LogACD₂ with Exponential, Weibull, Generalized Gamma and Burr innovations, and Log–ARMA with normal and SGE innovations. Figure 6 shows histograms accompanied by 95% confidence bands (computed using multinomiality of histogram heights under the null) of probability integral transforms for Lukoil and Alcatel; the models are represented by best performing ACD specifications for each distribution. Figure 7 displays analogous diagrams when the working model is Log–ARMA with normal and SGE innovations.

Although relative out-of-sample performance of different models varies from stock to stock, it is clear that Log–ARMA models fare no worse than some preferable ACD specifications. This is especially true as far as Pearson's test for uniformity is concerned; the situation with autocorrelation tests is somewhat less clear.¹⁵ The benchmark normal Log–ARMA model decidedly beats any benchmark exponential ACD model and even any Weibull ACD model most of the time. At the same time, the SGE Log–ARMA specification and Generalized Gamma logarithmic ACD specifications (especially LogACD₁) fare

¹⁵Note that for Alcatel the Ljung–Box statistics display too much serial dependence in z . One possible explanation of this phenomenon is that no account is made of pronounced long memory in Alcatel durations (Jasiak, 1999); another is that trading intensity for Alcatel is subject to medium-run regime switches (see Figure 1).

comparably well¹⁶. Visually, the bottom left panels of Figures 6 and 7 showing predictive z -distributions for the Generalized Gamma LogACD₁ and SGE Log-ARMA models fitted to Lukoil seem equally ideal, while only the bottom right panel of Figure 7 representing the SGE Log-ARMA specification is indefectible for Alcatel.

4 Factors influencing trading intensity

4.1 Previous empirical studies

We use the experience of previous empirical studies to choose variables to include in our dynamic model. The regressors that various researchers have used can be roughly divided into three categories: volume, price, and spread variables. As our database does not contain quote data, we focus on the first two.

The category “volume variables” encompasses those based on trading volumes in recent transactions. Hafner (2005) includes lagged volume as an explanatory variable in the ACD equation. This yields a non-significant coefficient which justifies his assumption of non-causality from volume to durations. Russell (1999) too finds that large transacted log-volume does not appear to be followed by higher intensity of transactions. However, more researchers find significant effects of volume variables. Bauwens and Giot (2000) add lagged average volume per trade into the LogACD₂ model, which turns out to be significant both when entered separately and when accompanied by trading intensity and spread variables. Dufour and Engle (2000) include five lags of past log-volume in the ACD model, which seems to have great explanatory power. Spierdijk (2004) finds jointly significant several volume variables involving up to the five most recent transactions. It appears that a conclusion about significant volume effects highly depends on the definition of corresponding variables. As additional confirmation, Engle and Lunde (2003) introduce jointly into their LogACD₂ model the square root of the size of the previous trade, which always turns out significant, the mean of the square roots of the sizes of the 10 previous trades, which only sometimes is significant, and the absolute value of accumulated signed size of the same trades, which never turns out significant.

¹⁶Note that p -values for Pearson’s test in Bauwens, Giot, Grammig, and Veredas (2004) are straight zeros for all models for all stocks as far as trade durations are concerned. For Russian duration data, such p -values are often larger than 1% both for the SGE Log-ARMA model and for Generalized Gamma ACD models; for Alcatel durations, the p -value is larger than 1% only for the SGE Log-ARMA model.

Related variables take into account imperfections of the data such as zero price changes and discreteness of time stamps. Bauwens and Giot (2000) include the number of trades recorded during the price duration (i.e., the time interval between successive price changes) divided by the length of that duration, and find that a higher value of this variable tends to shorten the next expected duration. Zhang, Russell and Tsay (2001) instead add as a regressor an indicator that at last trade, multiple trades have occurred, and find the corresponding coefficient significant and negative. An additional variable constructed similarly for the second to last trade turns out to be insignificant.

The category “price variables” encompasses those based on returns from recent transactions and price volatility over them. There is more agreement among different studies concerning such variables, although McCulloch and Tsay (2001) find price changes insignificant in a very simple Log-ARMA model, and Russell (1999) finds that volatility of price changes has little impact on transaction rates. However, Russell (1999) finds that a price change between subsequent transactions has a positive and significant coefficient; Spierdijk (2004) includes two more ancient lags of absolute returns, and both turn out significant for most stocks in her sample. Dufour and Engle (2000) find mixed evidence on significance of lagged absolute returns for a sample of several NYSE stocks, with a negative effect on durations prevailing. Russell and Engle (2005) find that expected transaction intensity tends to be higher when conditional price volatility per transaction is higher. Likewise, Grammig and Wellner (2002) obtain positive and highly jointly significant effects of the latent expected price volatility and price volatility shock.

Market microstructure theories such as those of Admati and Pfleiderer (1988) and Easley and O’Hara (1992) presume the coexistence of informed and uninformed (liquidity) traders in the market. Increased trading volumes and price volatility serve as indicators of informed trading, which may either enhance trading activity if informed traders are further attracted by opportunities to exploit their private information, or slow it down if liquidity traders are intimidated by the possibility of informed trading (see, e.g., Spierdijk, 2004).

4.2 Estimation results

We perform estimation on the basis of the Log-ARMA(1,2)-GARCH(1,1) model that worked well before, with conditionally normal innovations for simplicity, calculating standard errors in the Bollerslev-Wooldridge (1992) form. We use the following *ad hoc* mod-

eling procedure: starting from this specification with the indexed by t mean equation supplemented by factor variables indexed by $t - 1, t - 2, \dots, t - 10$, we remove one by one the least significant variables until all p-values are lower than 1%. The following factor variables are chosen:

VOL: log-volume for each trade;

ST: number of simultaneous transactions occurring at each trade;

RV: absolute value of return times log-volume for each trade;

PR: absolute value of latest non-zero price change for each trade.

The choice of these variables is driven primarily by tendencies in the literature reviewed in the previous subsection. In addition, some experimentation for the Russian stocks revealed the following trends. The variable ST works better than an indicator of more than one transaction, in contrast to the finding of Zhang, Russell and Tsay (2001). The variable RV is a proxy for the degree of “market overheating.” Finally, previous returns or non-zero returns are insignificant if used without taking absolute values. Further, to investigate possible asymmetric effects of positive and negative returns, we took the final estimated mean equations, and in addition included $\text{PR}(-1)\mathbb{I}[\text{R}_{\neq 0}(-1) < 0]$, where $\mathbb{I}[\text{R}_{\neq 0}(-1) < 0]$ is an indicator of whether the latest non-zero price change is negative. Table 8 contains the estimation results for all Russian stocks and Alcatel. The table reports only point estimates of coefficients on the factors and AR and MA terms; all are significant at the 1% confidence level. The values of the Ljung–Box statistics for standardized log-duration innovations shown at the bottom once again indicate good performance of the Log–ARMA framework.

One can immediately see that the introduced factors do have strong effects, and influence the dynamics of intertrade durations for all stocks pretty much in the same way. The lagged volume has positive impact on trading intensity as in Engle and Lunde (2003), Bauwens and Giot (2000), Dufour and Engle (2000), and Spierdijk (2004), attesting that informed traders are encouraged by observing increasing trading volumes. The effect of the presence of simultaneous transactions during the last trade on the current duration is stably negative, as in Zhang, Russell and Tsay (2001). This effect is purely a result of the aggregation of trades that happened within one second, and hardly has structural interpretation. Market overheating naturally leads to a higher trading intensity, but the effect nearly fades away during later trades (evidently, there is no overheating at the Paris Bourse, and the corresponding coefficient is probably spurious). Both long-run and short-run influence of the last non-zero price changes on durations is significantly positive

for all stocks. The positive sign is in agreement with the majority of stocks in the sample of Spierdijk (2004), but at variance with the majority of NYSE stocks in the sample of Dufour and Engle (2000). It is likely that in the Russian stock market large transaction price changes signal to liquidity traders that new information unavailable to them has entered the market, making them avoid risky trading.

Inclusion of asymmetric effects of positive and negative returns yields that coefficients belonging to $PR(-1)$ are insignificantly different (with p-values ranging from 8% to 83%), depending on whether the latest non-zero price change has been positive or negative. Thus at the MICEx, the traders react evenly to upward and downward recent price trends, i.e. both directions of the market are perceived as equally “normal,” at least in the short run.

5 Concluding remarks

This paper is fully focused on durations. Future research may be devoted to a careful analysis of price (and possibly volume) data, including modeling price movements in a framework that acknowledges discreteness of their changes, modeling joint dynamics of durations and mark processes, and analysis of effects of market-wide, industry-wide, and stock-specific news.

References

- Admati, A.R. and P. Pfleiderer (1988) *A theory of intraday patterns: volume and price variability*. Review of Financial Studies 1, 3–40.
- Anatolyev, S. and S. Korepanov (2003) *The term structure of Russian interest rates*. Applied Economics Letters 10, 867–870.
- Aradhyula, S.V. and A.T. Ergün (2004) *Trading collar, intraday periodicity and stock market volatility*. Applied Financial Economics 14, 909–913.
- Bali, T.G. and P. Theodossiou (2003) *Risk measurement performance of alternative distribution functions*. Annals of Operations Research, forthcoming.
- Bollerslev, T. and J.M. Wooldridge (1992) *Quasi maximum likelihood estimation and inference in dynamic models with time varying covariances*. Econometric Reviews 11, 143–172.
- 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.
- Bauwens, L., P. Giot, J. Grammig, and D. Veredas (2004) *A comparison of financial duration models via density forecasts*. International Journal of Forecasting 20, 589–609.
- Chen, S.X. (2000) *Beta kernel smoothers for regression curves*. Statistica Sinica 10, 73–91.
- Diebold, F.X., T.A. Gunther, and A.S. Tay (1998) *Evaluating density forecasts with applications to financial risk management*. International Economic Review 39, 863–883.
- Drost, F.C. and B.J.M. Werker (2004) *Semiparametric duration models*. Journal of Business & Economic Statistics 22, 40–50.
- Dufour, A. and R.F. Engle (2000) *Time and the price impact of a trade*. Journal of Finance 55, 2467–2498.
- Easley, D. and M. O’Hara (1992) *Time and the process of security price adjustment*. Journal of Finance 47, 577–606.
- Engle, R.F. (2000) *The econometrics of ultra-high frequency data*. Econometrica 68, 1–22.
- Engle, R.F. and G. Gonzalez-Rivera (1991) *Semiparametric ARCH Models*. Journal of Business & Economic Statistics 9, 345–359.
- Engle, R.F. and A. Lunde (2003) *Trades and quotes: a bivariate point process*. Journal of Financial Econometrics 1, 159–188.

- Engle, R.F. and J.R. Russell (1998) *Autoregressive conditional duration: a new model for irregularly spaced transaction data*. *Econometrica* 66, 1127–1162.
- Fedorov, P. and S. Sarkissian (2000) *Cross-sectional variations in the degree of global integration: the case of Russian equities*. *Journal of International Financial Markets, Institutions and Money* 10, 131–150.
- Ghysels, E., C. Gouriéroux, and J. Jasiak (2004) *Stochastic volatility duration models*. *Journal of Econometrics* 119, 413–433.
- Gouriéroux, C. and J. Jasiak (2000) *Durations*. In: “A Companion to Theoretical Econometrics”, B. Baltagi (ed.), Blackwell Publishers.
- Grammig, J. and K.-O. Maurer (2000) *Non-monotonic hazard functions and the autoregressive conditional duration model*. *Econometrics Journal* 3, 16–38.
- Grammig, J. and M. Wellner (2002) *Modeling the interdependence of volatility and inter-transaction duration processes*. *Journal of Econometrics* 106, 369–400.
- Hafner, C.M. (2005) *Durations, volume and the prediction of financial returns in transaction time*. *Quantitative Finance* 5, 145–152.
- Hanousek, J. and R. Podpiera (2003) *Informed trading and the bid-ask spread: evidence from an emerging market*. *Journal of Comparative Economics* 31, 275–296.
- Hayo, B., and A.M. Kutan (2005) *The impact of news, oil prices, and global market developments on Russian financial markets*. *Economics of Transition* 13, 373–393.
- Jasiak, J. (1999) *Persistence in intertrade durations*. *Finance* 19, 166–195.
- Lunde, A. (1999) *A generalized gamma autoregressive conditional duration model*. Working paper, Aalborg University.
- McCulloch, R.E. and R.S. Tsay (2001) *Nonlinearity in high frequency financial data and hierarchical models*. *Studies in Nonlinear Dynamics and Econometrics* 5(1), article 1.
- Medvedev, A.N. and G.V. Kolodyazhny (2001) *Financial crisis in Russia: the behavior of non-residents*. Working paper 2K/12E, EERC–Russia. Available at www.eerc.ru/publications/wp/WorkingPapers.htm.
- Medvedev, A.N. and G.V. Kolodyazhny (2003) *Russian stock market: participants and their strategies*. Working paper 01-060E, EERC–Russia. Available at www.eerc.ru/publications/wp/WorkingPapers.htm.
- Meitz, M. and T. Teräsvirta (2003) *Evaluating models of autoregressive conditional duration*. Manuscript, Stockholm School of Economics.
- Nelson, D (1991) *Conditional heteroskedasticity in asset returns: a new approach*. *Journal of Econometrics* 43, 227–251.

- Nyblom, J. (1989) *Testing the constancy of parameters over time*. Journal of the American Statistical Association 84, 223–230.
- Ostrovsky, A. (2003) *From chaos to capitalist triumph*. Financial Times (UK), Oct 9, pg. 4.
- Panas, E. (2005) *Generalized beta distributions for describing and analysing intraday stock market data: testing the U-shape pattern*. Applied Economics 37, 191–199.
- Russell, J.R. (1999) *Econometric modeling of multivariate irregularly-spaced high-frequency data*. Manuscript, GSB, University of Chicago.
- Russell, J.R. and R.F. Engle (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 23, 166–180.
- Spiersdijk, L. (2004) *An empirical analysis of the role of the trading intensity in information dissemination on the NYSE*. Journal of Empirical Finance 11, 163–184.
- Tyurin, K. (2002) *Modeling high-frequency dynamics of financial markets in continuous time: an event-based approach*. Manuscript, Indiana University–Bloomington.
- Zhang, M.Y., J.R. Russell, and R.S. Tsay (2001) *A nonlinear autoregressive conditional duration model with applications to financial transaction data*. Journal of Econometrics 104, 179–207.

A Appendix

Engle and Russell (1998) show that the $\text{ACD}(p, q)$ model can be reformulated as $\text{ARMA}(\max(p, q), q)$ for durations *in levels* with non-Gaussian innovations. These innovations exhibit conditional heteroskedasticity whose parameters are tied to the ACD coefficients. Consider now the $\text{LogACD}_1(p, q)$ model (3). From equations (3) and (1),

$$\begin{aligned} \ln d_t &= \omega + \sum_{i=1}^p \phi_i \ln d_{t-i} + \sum_{i=1}^q \psi_i \ln \pi_{t-i} + \ln \varepsilon_t \\ &= \omega + \left(1 - \sum_{i=1}^q \psi_i\right) m + \sum_{i=1}^{\max(p,q)} (\phi_i + \psi_i) \ln d_{t-i} + \zeta_t - \sum_{i=1}^q \psi_i \zeta_{t-i}, \end{aligned}$$

where $m = E[\ln \varepsilon_t]$ and $\zeta_t = \ln \varepsilon_t - m$, so that $\ln d_t$ follows an $\text{ARMA}(\max(p, q), q)$ process. If ε_t is distributed lognormally¹⁷ with parameters¹⁸ $(-\sigma^2/2, \sigma^2)$, then ζ_t is distributed normally with parameters $(0, \sigma^2)$. Thus, the LogACD_1 model is equivalent to a conditionally normal homoscedastic Log-ARMA model. Other distributional specifications for ε_t would imply conditionally non-normal homoscedastic Log-ARMA models. A conditionally heteroskedastic specification (7) relaxes serial independence in ζ_t and hence that in ε_t . If we start from the $\text{LogACD}_2(p, q)$ model (4), however, the past innovations ζ_{t-i} then enter the ARMA equation nonlinearly, but still are conditionally homoskedastic.

¹⁷This may be viewed as a limiting case of the generalized gamma distribution.

¹⁸Recall that if x is distributed lognormally with parameters (μ, σ^2) , then $\ln x$ is distributed normally with parameters (μ, σ^2) , and $E[x] = \exp(\mu + \sigma^2/2)$. Due to the constraint $E[\varepsilon_t] = 1$, μ and σ must be related by $\mu = -\sigma^2/2$.

Table 1. Trading regimes at MICEX

Trading regime	Timing	Description
Pre-trading period	10:00AM–10:30AM	The opening price is determined.
Trading session	10:30AM–6:45PM	The regular trading regime.
Post-trading period	6:45PM–7:00PM	Trades can be executed at a weighted average price determined during last 30 minutes of the trading session.
Negotiated deals regime	10:00AM–7:05PM	Designed for large trades not executed immediately, but whose execution may be delayed up to 30 days.
Special negotiated deals regime	7:05PM–10:00PM	The same as the negotiated deals regime, but only specific securities can be traded.
REPO regime	10:15AM–7:05PM	The repurchase agreement trading regime.
Incomplete lots regime	5:00PM–7:05PM	Nonstandard lots can be traded.

Table 2. Description of Russian securities

Company	Security code	Price step	Lot	Stock share	Observations in sample
Sberbank	RU0009029540	1	1	0.75%	13156
Yukos	RU0009054449	1	1	1.99%	20125
Lukoil	RU0009024277	1	1	15.71%	34316
SurgutNG	RU0008926258	0.1	100	5.27%	38275
Rostelecom	RU0008943394	1	1	2.75%	38246
MosEnergo	RU14MSNG3008	0.1	100	1.88%	32452

Note: price step is in kopecks (1/100 of the ruble), lots are in shares.

Table 3. Descriptive statistics of adjusted durations and log durations

Stock	Mean	Var	Skew	Kurt	Min	Med	Max	Acor
Durations d_t								
Sberbank	0.998	2.791	4.287	32.96	0.002	0.378	24.799	0.159
Yukos	1.004	2.286	3.818	26.62	0.005	0.442	23.344	0.145
Lukoil	1.001	2.434	4.042	28.51	0.010	0.446	24.538	0.167
SurgutNG	1.001	2.307	4.033	29.67	0.008	0.460	27.500	0.171
Rostelecom	1.004	2.410	4.361	38.37	0.009	0.458	36.084	0.151
MosEnerg	1.006	2.308	3.885	27.31	0.007	0.456	24.131	0.158
Alcatel	0.991	1.730	4.208	50.37	0.005	0.524	36.187	0.134
Log-durations $\ln d_t$								
Sberbank	-1.008	2.325	-0.157	2.579	-6.197	-0.974	3.211	0.266
Yukos	-0.864	2.025	-0.235	2.639	-5.270	-0.816	3.150	0.213
Lukoil	-0.848	1.892	-0.128	2.570	-4.604	-0.807	3.200	0.240
SurgutNG	-0.802	1.771	-0.120	2.606	-4.794	-0.776	3.314	0.234
Rostelecom	-0.816	1.802	-0.105	2.586	-4.717	-0.781	3.586	0.234
MosEnerg	-0.813	1.826	-0.144	2.629	-4.940	-0.786	3.183	0.227
Alcatel	-0.752	1.789	-0.345	2.674	-5.335	-0.646	3.589	0.213

Notes: Mean is sample mean, Var is sample variance, Skew is sample skewness coefficient, Kurt is kurtosis, Min is sample minimum, Max is sample maximum, Acor is sample first-order autocorrelation coefficient.

Table 4. Estimates of exponential ACD, LogACD and Log-ARMA-GARCH models for Lukoil

	ACD	LogACD ₁	LogACD ₂	Log-ARMA
ω	0.033 (0.004)	0.029 (0.004)	-0.046 (0.008)	
ψ_1	0.998 (0.041)	1.261 (0.060)	1.394 (0.075)	
ψ_2	-0.120 (0.035)	-0.317 (0.052)	-0.416 (0.071)	
ϕ_1	0.205 (0.009)	0.162 (0.006)	0.144 (0.005)	
ϕ_2	-0.113 (0.006)	-0.123 (0.007)	-0.100 (0.008)	
α				-0.091 (0.009)
β				0.896 (0.011)
γ_1				-0.708 (0.012)
γ_2				-0.057 (0.008)
μ				0.040 (0.008)
λ				0.019 (0.002)
ν				0.958 (0.006)
$Q(15)$	34.61	107.24	22.57	20.97
$Q^2(15)$	37.90	125.44	33.58	15.39
<i>ModelNyb</i>	0.90	0.83	0.74	1.24
<i>IndivNyb</i>	0.09	0.60	0.34	0.19
<i>Mean LL</i>	-0.9257	-0.9241	-0.9278	-0.8427
<i>AIC</i>	1.8518	1.8484	1.8559	1.6858
<i>BIC</i>	1.8521	1.8488	1.8563	1.6863
<i>KS</i>	20.66	21.33	20.69	3.44

Notes: Standard errors in parentheses are in the Bollerslev–Wooldridge form. “ $Q(15)$ ” and “ $Q^2(15)$ ” denote Ljung–Box statistics for standardized residuals and their squares, whose 1% critical value is 30.58. “*ModelNyb*” denotes the Nyblom statistic for the whole model, whose 1% critical value is approximately 1.87 for ACD and LogACD models, and 2.35 for Log-ARMA-GARCH model; “*IndivNyb*” denotes the maximal across the coefficients individual Nyblom statistic, whose 1% critical value is 0.74. “*MeanLL*” denotes the sample mean loglikelihood; “*AIC*” and “*BIC*” – Akaike and Schwarz information criteria, respectively. “*KS*” denotes the Kolmogorov–Smirnov statistic measuring divergence between the exponential and empirical distributions, whose 1% critical value is 1.63.

Table 5. Estimates of exponential ACD, LogACD and Log-ARMA-GARCH models for Alcatel

	ACD	LogACD ₁	LogACD ₂	Log-ARMA
ω	0.014 (0.004)	0.018 (0.004)	-0.027 (0.005)	
ψ_1	1.245 (0.046)	1.396 (0.068)	1.555 (0.051)	
ψ_2	-0.309 (0.043)	-0.442 (0.060)	-0.566 (0.050)	
ϕ_1	0.145 (0.010)	0.118 (0.007)	0.110 (0.006)	
ϕ_2	-0.094 (0.010)	-0.090 (0.007)	-0.084 (0.007)	
α				-0.062 (0.009)
β				0.924 (0.012)
γ_1				-0.765 (0.014)
γ_2				-0.032 (0.009)
μ				0.031 (0.002)
λ				0.011 (0.003)
ν				0.987 (0.002)
$Q(15)$	27.35	45.69	18.54	15.88
$Q^2(15)$	5.36	7.02	6.07	30.25
<i>ModelNyb</i>	7.70	4.05	6.60	11.26
<i>IndivNyb</i>	1.44	2.97	4.49	5.42
<i>Mean LL</i>	-0.9404	-0.9379	-0.9403	-0.9082
<i>AIC</i>	1.8813	1.8763	1.8811	1.8171
<i>BIC</i>	1.8818	1.8769	1.8817	1.8179
<i>KS</i>	10.93	11.31	10.91	4.65

Notes: Standard errors in parentheses are in the Bollerslev–Wooldridge form. “ $Q(15)$ ” and “ $Q^2(15)$ ” denote Ljung–Box statistics for standardized residuals and their squares, whose 1% critical value is 30.58. “*ModelNyb*” denotes the Nyblom statistic for the whole model, whose 1% critical value is approximately 1.87 for ACD and LogACD models, and 2.35 for Log-ARMA-GARCH model; “*IndivNyb*” denotes the maximal across the coefficients individual Nyblom statistic, whose 1% critical value is 0.74. “*MeanLL*” denotes the sample mean loglikelihood; “*AIC*” and “*BIC*” – Akaike and Schwarz information criteria, respectively. “*KS*” denotes the Kolmogorov–Smirnov statistic measuring divergence between the exponential and empirical distributions, whose 1% critical value is 1.63.

Table 6. Estimates of the SGE-Log-ARMA-GARCH model

	Sberbank	Yukos	Lukoil	SurgutNG	Rostelecom	MosEnergo	Alcatel
α	-0.191 (0.023)	-0.139 (0.016)	-0.084 (0.009)	-0.087 (0.007)	-0.064 (0.005)	-0.079 (0.007)	-0.055 (0.008)
β	0.810 (0.022)	0.841 (0.019)	0.904 (0.010)	0.895 (0.008)	0.924 (0.006)	0.905 (0.008)	0.934 (0.010)
γ_1	-0.614 (0.024)	-0.686 (0.020)	-0.725 (0.012)	-0.723 (0.010)	-0.760 (0.008)	-0.746 (0.010)	-0.792 (0.013)
γ_2	-0.023 (0.011)	-0.025 (0.009)	-0.056 (0.007)	-0.042 (0.006)	-0.051 (0.006)	-0.032 (0.007)	-0.034 (0.008)
μ	0.022 (0.008)	0.067 (0.029)	0.052 (0.011)	0.023 (0.004)	0.018 (0.003)	0.015 (0.003)	0.005 (0.001)
λ	0.026 (0.004)	0.017 (0.004)	0.023 (0.003)	0.014 (0.002)	0.015 (0.001)	0.015 (0.001)	0.014 (0.002)
ν	0.964 (0.006)	0.947 (0.019)	0.948 (0.008)	0.972 (0.004)	0.975 (0.003)	0.976 (0.002)	0.983 (0.001)
κ	2.538 (0.057)	2.420 (0.040)	2.374 (0.033)	2.357 (0.029)	2.425 (0.031)	2.328 (0.031)	2.283 (0.039)
φ	-0.126 (0.015)	-0.204 (0.011)	-0.143 (0.008)	-0.125 (0.008)	-0.122 (0.008)	-0.139 (0.008)	-0.285 (0.011)
$Q(15)$	36.36	28.91	27.63	24.38	28.46	19.01	35.17
$Q^2(15)$	16.92	15.34	14.42	28.62	20.58	13.33	26.03
$ModelNyb$	4.24	2.53	1.59	1.91	2.56	2.96	15.02
$IndivNyb$	1.25	0.60	0.18	0.24	0.70	0.56	5.62
$Mean LL$	-0.7700	-0.8617	-0.8367	-0.8496	-0.8420	-0.8518	-0.8918
KS	0.88	1.37	2.27	2.13	1.56	1.83	0.58

Notes: Standard errors in parentheses are in the Bollerslev-Wooldridge form. “ $ModelNyb$ ” denotes the Nyblom statistic for the whole model, whose 1% critical value is approximately 2.81; “ $IndivNyb$ ” denotes the maximal across the coefficients individual Nyblom statistic, whose 1% critical value is 0.74. For the meaning of other diagnostic statistics, see notes to Tables 4 and 5.

Table 7. Results of out-of-sample testing

Model	Pearson	$Q(15)$	$Q^2(15)$	Pearson	$Q(15)$	$Q^2(15)$
	Sberbank			Yukos		
Exponential ACD	956.78	92.76	59.09	703.10	53.76	36.98
Weibull ACD	321.22	87.34	56.28	269.23	51.98	35.70
Gen. Gamma ACD	33.49	87.17	60.93	23.64	46.26	35.04
Burr ACD	91.94	66.83	45.44	70.86	39.51	30.93
Exponential LogACD ₁	1017.35	102.91	111.60	729.85	40.82	55.25
Weibull LogACD ₁	350.21	63.65	77.06	310.32	30.48	42.86
Gen. Gamma LogACD ₁	35.35	27.61	40.27	28.41	20.93	32.68
Burr LogACD ₁	89.24	20.39	31.31	72.69	20.03	31.60
Exponential LogACD ₂	985.06	169.50	111.78	707.61	88.78	62.05
Weibull LogACD ₂	292.87	184.98	127.36	268.13	91.68	64.67
Gen. Gamma LogACD ₂	39.61	217.01	160.71	28.19	94.30	69.20
Burr LogACD ₂	101.67	193.46	140.20	70.86	85.58	63.01
Normal Log-ARMA	59.62	39.72	52.21	108.61	30.85	43.99
SGE Log-ARMA	27.02	61.07	75.18	39.76	42.47	56.81
	Lukoil			Alcatel		
Exponential ACD	973.52	86.20	56.91	385.74	295.00	185.27
Weibull ACD	510.27	82.08	55.96	222.11	280.06	191.54
Gen. Gamma ACD	47.72	85.80	62.00	54.95	271.35	199.94
Burr ACD	114.46	66.26	47.20	112.77	227.35	164.74
Exponential LogACD ₁	1029.19	48.05	63.52	365.39	215.13	184.78
Weibull LogACD ₁	582.23	36.00	50.93	218.70	185.65	173.94
Gen. Gamma LogACD ₁	36.04	21.11	38.56	58.97	129.83	135.95
Burr LogACD ₁₂	95.64	16.25	29.29	89.21	114.25	119.41
Exponential LogACD ₂	922.93	145.16	92.92	355.68	300.31	194.54
Weibull LogACD ₂	496.10	151.20	101.78	183.69	298.85	207.35
GGam LogACD ₂	54.44	185.29	137.34	66.49	342.10	253.99
Burr LogACD ₂	103.51	162.35	117.23	121.39	293.01	212.12
Normal Log-ARMA	122.71	22.53	42.29	157.48	219.43	208.45
SGE Log-ARMA	35.53	28.14	47.85	31.04	280.87	245.61

Notes: “Pearson” denotes Pearson test statistic for uniformity of probability integral transforms z , whose 1% critical value is 36.19. “ $Q(15)$ ” and “ $Q^2(15)$ ” denote Ljung–Box statistics for z and z^2 , whose 1% critical value is 30.58.

Table 8. Factors influencing trading intensity, the normal Log-ARMA-GARCH model

	Sberbank	Yukos	Lukoil	SurgutNG	Rostelecom	MosEnerg	Alcatel
AR(1)	0.830	0.825	0.891	0.890	0.923	0.906	0.905
MA(1)	-0.630	-0.660	-0.717	-0.724	-0.761	-0.747	-0.755
MA(2)	-0.033	-0.026	-0.056	-0.044	-0.055	-0.037	-0.022
VOL(-1)	-0.023	-0.015	-0.008	-0.009	-0.013	-0.022	-0.020
VOL(-2)							-0.025
VOL(-3)					0.010		
VOL(-5)							0.010
ST(-1)	-0.047	-0.053	-0.051	-0.046	-0.027		
ST(-2)							-0.017
ST(-4)		0.016		0.008			
ST(-7)		-0.017		-0.007			
RV(-1)	-0.042	-0.021	-0.021	-0.018	-0.018	-0.021	
RV(-2)	0.040	0.024	0.018	0.020	0.016	0.016	
RV(-3)							0.008
PR(-1)	0.082	0.041	0.034	0.042	0.038	0.034	0.047
PR(-3)	-0.036		-0.016	-0.011			-0.020
PR(-4)					-0.012		
PR(-5)						-0.010	
PR(-9)						-0.010	
$Q(15)$	25.62	22.60	17.32	23.82	21.20	17.79	15.39
$Q^2(15)$	14.04	19.46	14.51	28.61	18.92	11.57	28.16

Notes: All coefficients are Bollerslev-Wooldridge significant at 1% level. “ $Q(15)$ ” and “ $Q^2(15)$ ” denote Ljung-Box statistics for standardized residuals and their squares, whose 1% critical value is 30.58.




Figure 1: Average durations for different days.




Figure 2: Estimated intraday patterns for Lukoil durations.




Figure 3: Intraday patterns (Monday and Wednesday) with 95% confidence bands for Lukoil.




Figure 4: Unconditional distributions of durations.




Figure 5: Unconditional distributions of log-durations.

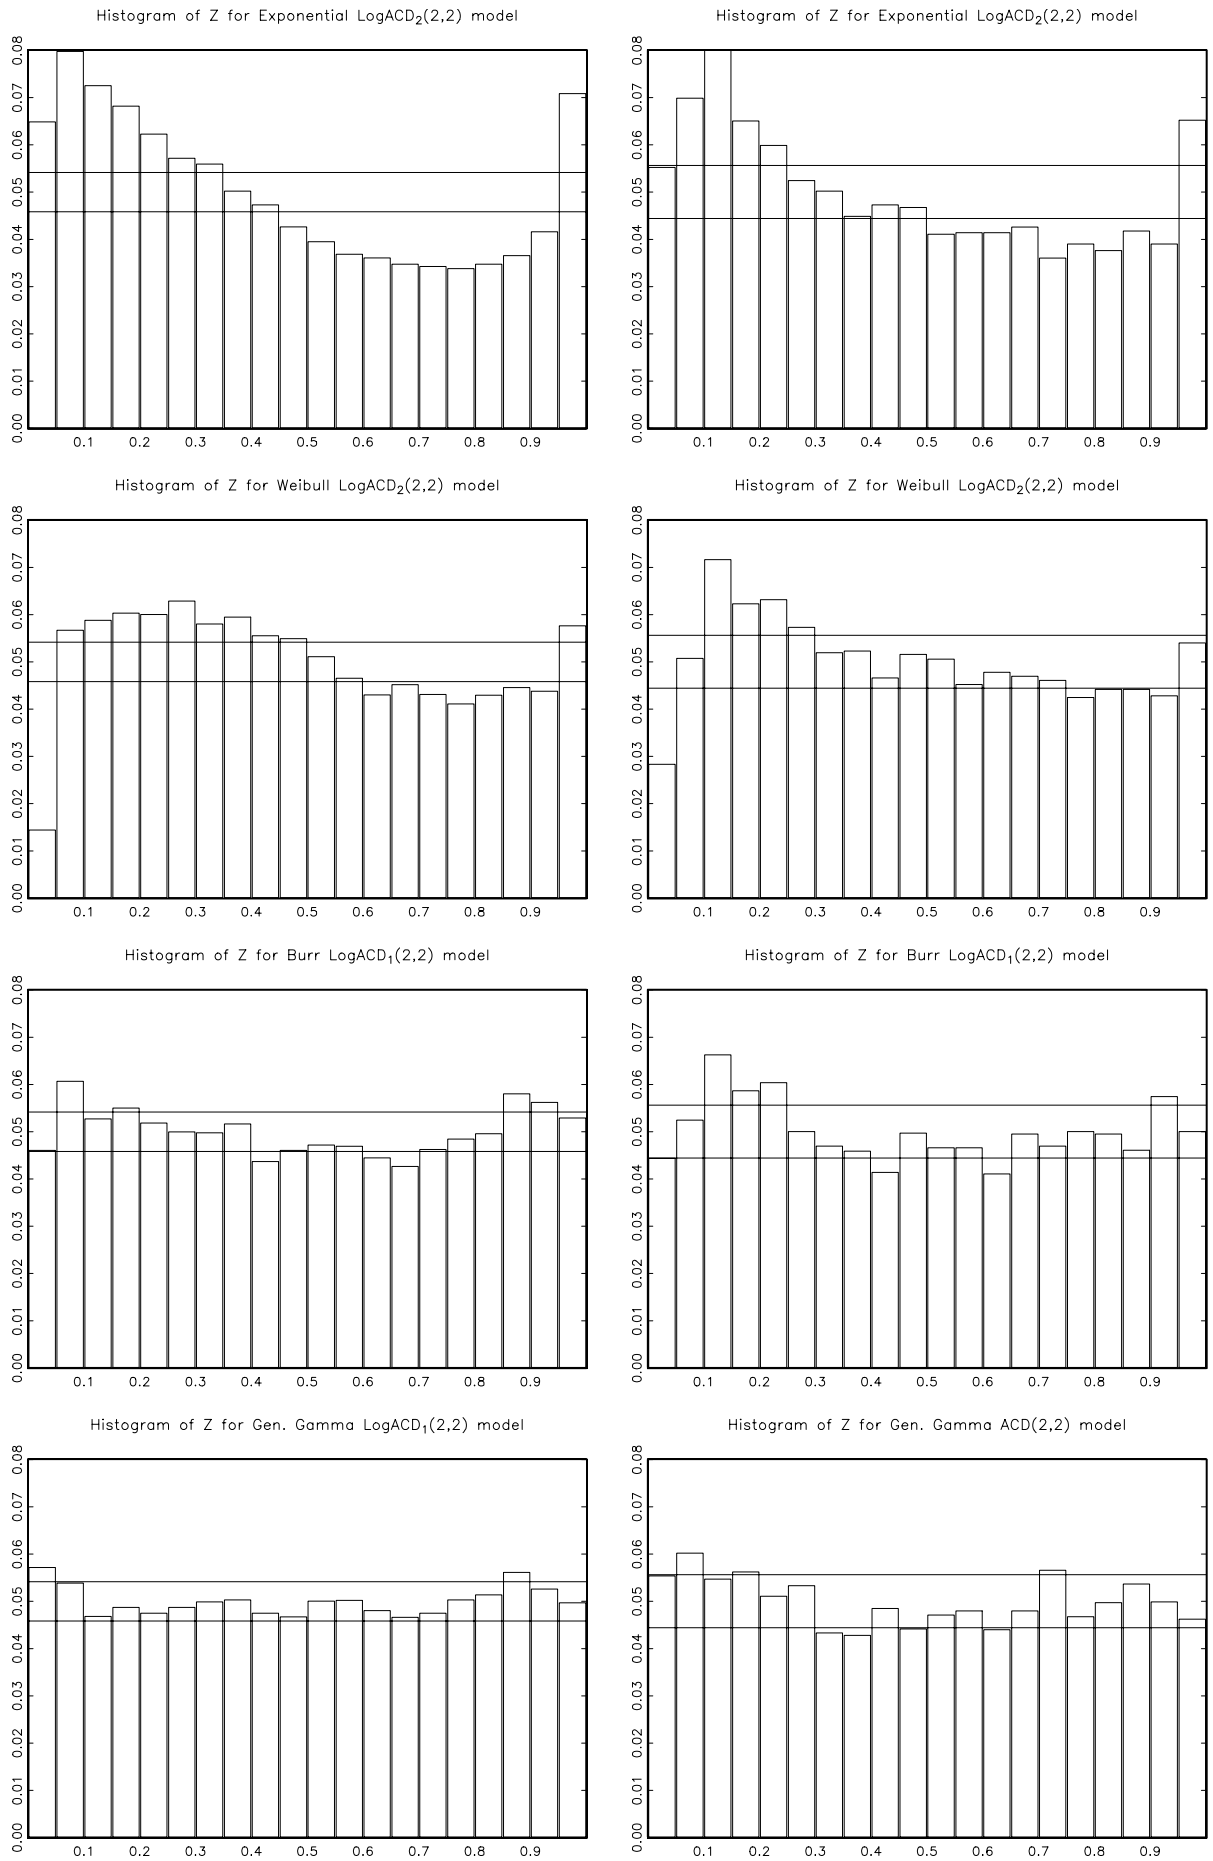


Figure 6: Histograms of probability integral transforms for Lukoil (left) and Alcatel (right).




Figure 7: Histograms of probability integral transforms for Lukoil (left) and Alcatel (right), continued.