CORE DISCUSSION PAPER
2000/60

A COMPARISON OF FINANCIAL DURATION MODELS VIA DENSITY FORECASTS

Luc BAUWENS[1] Pierre GIOT[2] Joachim GRAMMIG[3] and David VEREDAS[4]

December 2000

Abstract

Using density forecasts, we compare the predictive performance of duration models that have been developed for modelling intra-day data on stock markets. Our model portfolio encompasses the autoregressive conditional duration (ACD) model, its logarithmic version (Log-ACD), the threshold ACD (TACD) model – in each case with alternative error distributions –, the stochastic conditional duration model (SCD), and the stochastic volatility duration model (SVD). The evaluation is done on transaction, price, and volume durations of four stocks listed at the NYSE. The results lead us to conclude that the ACD/log-ACD/TACD/SCD models capture the dynamic dependence in the data in a satisfactory way. They fit correctly the conditional distribution of volume durations, but fail to do so for trade durations. The evidence is mixed for price durations and ACD based models, poor for the SCD model. The SVD model in its original version performs worse than the (Log-)ACD models on the dynamics of trade durations, and offers no improvement with respect to the distributional aspect. The SVD is not suitable to model volume durations. Regarding price durations the performance of the SVD is comparable to those of (Log-)ACD specifications that provide the best results.

Keywords: duration, high frequency data, density forecast.

JEL Classification: C41, C52, C53, G14.î

[1]CORE and Department of Economics, Université Catholique de Louvain, Belgium. E-mail: bauwens@core.ucl.ac.be

[2]CORE, Université Catholique de Louvain, Belgium and Department of Quantitative Economics, Maastricht University. E-mail: giot@core.ucl.ac.be

[3]Department of Economics, Johann Wolfgang Goethe-University Frankfurt.

[4]CORE, Université Catholique de Louvain, Belgium

Corresponding author: Luc BAUWENS, CORE, 34 Voie du Roman Pays, B-1348 Louvain-la-Neuve, Belgium. Tel: +32-10-474336, fax: +32-10-474301
We thank S. Frey, R. Straub, and S. Vuletic for excellent research assistance. Financial support from the Landeszentralbank Hessen is greatfully acknowledged.
This paper presents research results of the Belgian Program on Interuniversity Poles of Attraction initiated by the Belgian State, Prime Minister's Office, Science Policy Programming. The scientific responsibility is assumed by the authors.

1 Introduction

Several contributions to the market microstructure literature[1] emphasize that the waiting times between intra-day market events like trades, quote updates, price changes, and order arrivals play a key role for understanding the processing of private and public information in financial markets. Hence, the accessibility of financial data at a micro level, which, in an ideal case, includes real time recordings of trades, order arrivals and quote updates, as well as the corresponding prices, volumes and time stamps, opened new perspectives for the empirical analysis of market microstructure. By appropriately thinning these intra-day data, it is possible to define almost any event of interest, and the corresponding durations. An econometric framework for the modeling of intertemporally correlated event arrival times is provided by Engle and Russell (1998), who introduced the autoregressive conditional duration (ACD) model. The ACD approach combines elements from transition analysis (see e.g. Lancaster 1990) and Engle's (1982) autoregressive conditional heteroskedasticity (ARCH) model. Upon closer inspection, the motivation behind the ACD and the ARCH model appears similar: financial market events, like trades and quote arrivals, occur in clusters.

Following the contribution of Engle and Russell (1998), several duration models have been put forward. Bauwens and Giot (2000a) introduced a logarithmic version of the ACD model, called the Log-ACD model, which implies a nonlinear relation between the duration and its lags, and is more convenient than the ACD model when conditioning variables are included in the model in order to test microstructure effects. As an alternative to the Weibull distribution used by Engle and Russell (1998) in the ACD model, Grammig and Maurer (1999) introduced an ACD model based on the Burr distribution (which includes the Weibull as a particular case). Zhang, Russell and Tsay (1999) proposed the threshold ACD model, in the spirit of threshold autoregressive models, to capture a possible nonlinear relation between the duration and the past information variables. Ghysels, Gouriéroux, and Jasiak (1997) proposed the stochastic volatility duration model, which accounts both for mean and variance dynamics in duration processes. Bauwens and Veredas (1999) put forward the stochastic conditional duration model, which uses a stochastic volatility-type model instead of a GARCH-type model to model the durations.

Given the availability of several competing models, it becomes relevant to examine if some models perform better than others on some criteria, and which type of model is particularly suited, or inadequate, for specific duration processes. Our aim is to contribute to this issue using a criterion of predictive ability. Since Diebold, Gunther, and Tay (1998) have proposed tools to evaluate density forecasts produced by a given model, we can apply these tools to several models on the same data and get therefore an indication of whether some models fail to produce correct forecast densities. This procedure has the advantage of being easy to implement and suitable for comparing models which are not necessarily nested. Moreover, even if a model nests all the others, this model still needs to be evaluated and may be a bad performer in some respects.

[1] See O'Hara's book (1995) for a survey and references, and Goodhart and O'Hara (1997) for a more recent survey.

The paper is organized as follows. In Section 2, we explain the use of density forecasts for model evaluation. In Section 3, we review the duration models that are compared in this paper. The results of the application to real data are reported and discussed in Section 4. The last section contains our conclusions.

2 Density forecasts

In order to provide a comparative assessment of the financial duration models reviewed in the next section we employ the methods for evaluating density forecasts revived and formalized by Diebold, Gunther and Tay (1998), henceforth referred to as DGT.[2] DGT's approach makes it possible to compare non-nested models and evaluate their forecasting performance regardless of the users' loss functions.

A density forecast is a density defined for the next observation of the variable that is modelled, a duration in this paper. The forecast density is usually implied by the model that is used for estimation. In a dynamic context, the forecast density is taken as conditional on the information available at the previous occurence (all durations observed before the one-step-ahead prediction).

The basic ideas behind DGT's method are easily understood.[3] We deal with an ordered sequence of durations denoted x_i, for $i = 1, \ldots, m$. Let us denote by $\{f_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ a sequence of one-step-ahead density forecasts produced by any duration model and by $\{p_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ the sequence of densities defining the data generating process (DGP) governing the duration series x_i.[4] DGT show that the correct density is weakly superior to all other forecasts, i.e. will be preferred, in terms of expected loss, by all forecast users regardless of their loss functions. This suggests that forecasts should be evaluated by assessing whether the forecasting densities are correct, i.e. whether

$$\{f_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m = \{p_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m . \tag{1}$$

At first sight, testing whether (1) is true appears difficult because $p_i(x_i \mid \mathcal{H}_i)$ is never observed. However, the distributional properties of the probability integral transform

$$z_i = \int_{-\infty}^{x_i} f_i(u)du \tag{2}$$

provide the solution to this problem. Rosenblatt (1952) derived that the distribution of the probability integral transform under the null hypothesis (1) is uniform. DGT extend this result and show that under the null hypothesis, the distribution of the sequence of probability transforms $\{z_i\}_{i=1}^m$ of $\{x_i\}_{i=1}^m$ with respect to $\{f_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ is IID uniform. Hence, the empirical sequence of probabilty integral transforms produced by the forecasts of some model can be used for testing. DGT propose graphical tools that complement goodness-of-fit and independence

[2] The method relies on a result that dates back at least to Rosenblatt (1952). Applications of the method are included in Shephard (1994) and Kim *et al.* (1998).

[3] In Appendix 2, we have collected the propositions on which the DGT method is based.

[4] We denote by $\mathcal{H}_i$ the conditioning information set generated by the durations preceding x_i. For notational convenience, we shall sometimes use an abbreviated notation and suppress the information set $\mathcal{H}_i$ and, in one-period contexts, the time subscripts i. However, the dependence on $\mathcal{H}_i$ and the time dependence of the forecasts should always be kept in mind.

tests for IID uniformity. By plotting a histogram based on an empirical z sequence, departures from uniformity can easily be detected. A visual inspection will often provide obvious hints to the reasons for model failure. For instance, a humped shape of the z-histogram would indicate that the issued forecasts are too narrow and that the tails of the true density are not accounted for. On the other hand, a U-shape of the histogram would suggest that the model issues forecasts that either under- or overestimate too frequently. Confidence intervals for the z-histogram bins as well as a χ^2 goodness-of-fit test can easily be computed by exploiting familiar statistical properties of the histogram under the null hypothesis of uniformity. Autocorrelation functions (ACF) of $(z_i - \bar{z}), (z_i - \bar{z})^2 \ldots$ reveal potential deficiencies of a model to account for the dynamics of the duration process.

Although the basic ideas behind DHT's approach are straightforward, the implications for testing and comparing non-nested models are far reaching. Note that the method does not require that the data generating process remains the same over time. Instead, the true data generating process $\{p_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ may exhibit all kinds of structural change, as indicated by its time subscript. The important point is that a model density forecasts $\{f_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ are able to account for these features. In the next section, we define the density forecasts that are associated to each of the duration models that we review. In Section 4, we use the graphical tools (histogram and ACF), as well as goodness-of-fit and autocorrelation tests, for an empirical comparison of the duration models.

3 Model Review

In this section, we review several duration models that have been proposed in the literature. We distinguish five classes of models: autoregressive conditional duration (ACD), logarithmic ACD (Log-ACD), threshold ACD (TACD), stochastic conditional duration (SCD), and stochastic volatility duration (SVD) models. For each class, we describe the model, the estimation method that we use, and we define the corresponding density forecast.

3.1 ACD Models

The basic reference for ACD models is Engle and Russell (1998). ACD models specify the duration x_i as

$$x_i = \Psi_i \epsilon_i, \tag{3}$$

where ϵ_i (for $i = 1, \ldots, n$) is an IID sequence with $\mathrm{E}(\epsilon_i) = \mu$, such that $\mathrm{E}(x_i|\mathcal{H}_i) = \Psi_i \mu$. Ψ_i is specified as a linear function of past durations and conditional durations and is called the conditional duration.[5] For simplicity we restrict our attention to the ACD(1,1) case,

$$\Psi_i = \omega + \alpha x_{i-1} + \beta \Psi_{i-1}, \tag{4}$$

[5]This is a slight abuse of language since Ψ_i is the conditional expectation of x_i/μ, not of x_i. Equations (3)-(4) imply that x_i is equal to its conditional expectation times a random variable with expectation 1, and that the conditional expectation is a linear function of the previous duration and conditional expectation.

where $\omega > 0$, $\alpha \geq 0$, and $\beta \geq 0$ (but $\beta = 0$ if $\alpha = 0$). Although not necessary, these sign restrictions are convenient to ensure the positivity of Ψ_i in estimation.

A parametric model is obtained when the distribution of ϵ_i is specified up to a finite number of parameters. Engle and Russell (1998) proposed the standard exponential distribution (i.e. with parameter equal to 1), and as an extension the Weibull distribution with shape parameter equal to γ and scale parameter equal to 1. The exponential specification provides the quasi-likelihood function for parametric models, such that if $\mathrm{E}(x_i|\mathcal{H}_i)$ is correctly specified, the QML estimator of (ω, α, β) is consistent and asymptotically normal (under suitable regularity conditions). The Weibull distribution is more flexible in that it nests the exponential one (for $\gamma = 1$), and allows a non-flat hazard function. However, the Weibull hazard function is necessarily monotone: increasing if $\gamma > 1$, decreasing if $\gamma < 1$. As documented independently by Bauwens and Veredas (1999), and by Grammig and Maurer (1999), the hazard function of several types of financial durations may be increasing for small durations and decreasing for long durations (see Section 4). To account for this stylized fact, Grammig and Maurer proposed to use the Burr distribution which can have a hump-shaped hazard and nests the Weibull distribution as a particular case. The Burr distribution has two shape parameters, so that there is no one-to-one correspondance (contrary to the case of the Weibull distribution) between the properties of overdispersion (underdispersion) and of decreasing (respectively, increasing) hazard.[6] Another distribution which has also two shape parameters and breaks this one-to-one correspondance is the generalized gamma. It has been used by Lunde (1999) with the ACD model. It nests the Weibull and the gamma distributions.[7] For convenience, we have gathered in Appendix 1 the definitions and main properties of the densities that we refer to in this paper.[8]

Estimation of the parameters of ACD models can be done by maximizing the likelihood function, which is the product of the n densities $f(x_i\Psi_i^{-1})\Psi_i^{-1}$, where Ψ_i is defined by (4), and $f(.)$ is the appropriate density (see Appendix 1).

The density forecast of ACD models is the conditional density of x_i given the past information, thus simply $f(x_i\Psi_i^{-1})\Psi_i^{-1}$, e.g. a Burr density in the case of the Burr ACD model.

3.2 Log-ACD Models

When additional explanatory variables are added linearly to the right-hand side of eqn (4) and have negative coefficients, Ψ_i may become negative which is not admissible. Imposing non-negativity restrictions on these coefficients is tantamount

[6] Overdisperion means that the standard deviation is larger than the mean.

[7] The gamma density resembles the Weibull when their shape parameters take roughly the same value in a range from 0.5 to 2 (in particular they both correspond to the exponential density when that parameter is equal to 1). The gamma density is less convenient for evaluating density forecasts than the Weibull since its cdf is not known analytically. This is also the case for the generalized gamma. Other families could be considered, such as the lognormal distribution, or any distribution which has positive support. The lognormal density is less flexible than the Burr and generalized gamma, since it has necessarily a hump-shaped hazard, and less convenient since its cdf is also not available.

[8] More details on the properties of these distributions can be found in Bauwens and Giot (2001, Appendix to Ch. 4).

to delete the corresponding variables, which is self-destructive. This motivated Bauwens and Giot (2000a) to introduce logarithmic versions of the ACD models. Bauwens and Giot (2000b) provide the moments of first-order Log-ACD models. In Log-ACD models, eqn (3) is written as

$$x_i = \exp(\psi_i)\,\epsilon_i, \tag{5}$$

such that ψ_i is the logarithm of the conditional duration $\Psi_i = \exp(\psi_i)$. The difference with ACD models is that the autoregressive equation bears on the logarithm of the conditional duration rather than on the conditional duration itself. Two possible specifications of this equation are

$$\psi_i = \omega + \alpha \log x_{i-1} + \beta\psi_{i-1} \tag{6}$$

and

$$\psi_i = \omega + \alpha\epsilon_{i-1} + \beta\psi_{i-1}. \tag{7}$$

Notice that no sign restrictions are needed on the parameters to ensure the positivity of the conditional duration.

The hypotheses about ϵ_i (for $i = 1, \ldots, n$) are the same as in ACD models, as well as the possible probability distributions. Thus the most general models in this class are the Burr and generalized gamma Log-ACD models. Estimation and density forecasts are done similarly to the ACD case, taking into account that Ψ_i has been redefined.

3.3 Threshold ACD Models

In the standard ACD framework, the conditional mean dynamics are determined by the simple linear specifications (4), (6) or (7). Zhang, Russell, and Tsay (1999) henceforth referred to as ZRT, argue that financial duration processes require a more flexible specification. ZRT's threshold ACD model (TACD) is a three regime model where the regimes are allowed to have different duration persistence and error distributions. In a logarithmic version[9] of the TACD, ψ_i evolves as

$$\psi_i = \begin{cases} \omega_1 + \alpha_1 \log x_{i-1} + \beta_1\psi_{i-1} & \text{if} \quad 0 < x_{i-1} \le r_1 \\ \omega_2 + \alpha_2 \log x_{i-1} + \beta_2\psi_{i-1} & \text{if} \quad r_1 < x_{i-1} \le r_2 \\ \omega_3 + \alpha_3 \log x_{i-1} + \beta_3\psi_{i-1} & \text{if} \quad r_2 < x_{i-1} < \infty. \end{cases} \tag{8}$$

The threshold parameters r_1 and r_2 determine the regime boundaries. The idea behind (8) is that there exist regimes where the trading or quoting process is 'slow', 'normal' and 'fast', and that a linear specification would lead to an overshooting of the expected duration after a very short duration or a long duration. ZRT assume the generalized gamma distribution for ϵ_i (see Appendix 1) and allow the shape parameter ν to vary across regimes, thereby rendering the duration distribution to be regime-specific.

[9] We advocate a logarithmic version of the TACD model in order to avoid the non-negativity restrictions discussed in the previous sub-section.

Estimation is performed by a grid search over r_1 and r_2 and by maximizing the likelihood function conditional on given threshold values. In our application r_1 and r_2 are chosen to be the duration deciles. The grid search uses all possible permutations of the deciles under the restriction $r_1 < r_2$. We focus on a TACD version that employs the generalized gamma distribution and the logarithmic specification (8). Alternative TACD specifications would make use of the alternative distributions outlined in Sub-section 3.1 or employ a regime-adapted version of (4) or (7). Density forecasts are obtained like in the previous cases. Note that (8) implies that the density forecast depends on the regime (i.e. x_{i-1}) from which it is issued.

3.4 Stochastic Conditional Duration Models

Bauwens and Veredas (1999), hereafter BV, assume that durations are driven by a dynamic latent factor that is stochastic. This latent variable can be interpreted as inversely related to the information that arrives dynamically and randomly on the market and that triggers bursts of activity. BV assume for simplicity that the latent variable follows the first order autoregressive process

$$\psi_i = \omega + \beta\psi_{i-1} + u_i, \tag{9}$$

where u_i is independently normally distributed with zero mean and variance σ^2. Combining (9) with (5) yields a state space model when (5) is transformed in logarithm. Given this structure, the latent factor $\Psi_i = \exp(\psi_i)$ is lognormally distributed, both conditionally on the past durations and unconditionally. Notice that even if ϵ_i is assumed to follow a standard distribution for durations (like those used for ACD models), the conditional distribution of x_i is the mixture of the assumed distribution of ϵ_i by the conditional lognormal of Ψ_i. This mixture is the one-step-ahead forecast density that is implied by the SCD model. Likewise the unconditional distribution of x_i is the mixture of the distribution of ϵ_i by the unconditional lognormal of Ψ_i. BV compute these mixture distributions by unidimensional numerical integration[10] and illustrate the wide range of shapes of the hazard function that can be obtained depending on the distribution specified for ϵ_i. For example, if ϵ_i is distributed as Weibull, the shape of the hazard function vary from that of the Weibull distribution for very small values of σ^2 to humped shapes for large values of σ^2.

Estimation of the parameters of the SCD model can be done in several ways. Since the model is a nonlinear space state model, due to the non-normality of the distribution of $\ln \epsilon_i$ (unless ϵ_i is itself assumed lognormal), simulation-based methods are required for ML estimation. The approach followed by BV and in this research consists in approximating the distribution of $\ln \epsilon_i$ by a normal distribution with mean and variance that depend on the parameters of the assumed true distribution. A consistent (albeit inefficient) estimator of all the parameters can then be obtained by quasi-maximum likelihood using the Kalman filter and the prediction error decomposition. BV reported results for Weibull and gamma

[10] Thus, computing the probability integral transform z of a duration in the SCD model requires a bidimensional numerical integration: one for computing the forecast density, the other for computing z with this density.

assumptions about ϵ_i, and found them to be very similar. In the present paper, we study the forecast performance of the SCD model using the Weibull distribution.

3.5 Stochastic Volatility Duration Models

The ACD and Log-ACD models reviewed in the previous sub-sections share the assumption that the dynamics of higher moments of the duration process are governed by the dynamics of the conditional mean. Ghysels, Gouriéroux, and Jasiak (1997), henceforth referred to as GGJ, have argued that this feature restricts the flexibility that is needed when modeling financial duration processes. As a more versatile tool, GGJ introduce a nonlinear two factor model that disentangles the movements of the mean and of the variance of durations. Since the second factor is responsible for the variance heterogeneity, the model is referred to as the stochastic volatility duration (SVD) model. The departure point for the SVD is a standard static duration model where it is assumed that the durations are independently and exponentially distributed with a gamma heterogeneity,

$$x_i = \frac{U_i}{aV_i}, \tag{10}$$

where U_i and V_i are two independent variables with distributions gamma(1,1) (i.e. exponential) and gamma(b, b), respectively (see Appendix 1). Equation (10) can be rewritten to include Gaussian factors:

$$x_i = \frac{H(1, F_{1i})}{aH(b, F_{2i})} \tag{11}$$

where F_{1i} and F_{2i} are independent standard normal variables, and $H(b, F) = G(b, \varphi(F))$, where $G(b, .)$ is the quantile function of the gamma(b, b) distribution and φ the cdf of the standard normal. GGJ extend this static setup and propose to model duration dynamics through the two underlying Gaussian factors. We focus on a bivariate VAR representation for the process $F_i = (F_{1i}\, F_{2i})'$. The marginal distribution of F_i is constrained to be $N_2(0, I)$. This ensures that the marginal distribution of x_i belongs to the class of exponential distributions with gamma heterogeneity. Restricting our attention to a VAR of order one, we have

$$F_i = \Lambda F_{i-1} + u_i, \tag{12}$$

where u_i is a Gaussian white noise random vector with covariance matrix $\Sigma(\Lambda)$ such that $\mathrm{Var}(F_i) = I$. For a VAR of order one this is achieved by setting $\Sigma = I - \Lambda\Lambda'$.

Due to the nonlinear dynamic latent factor structure, the likelihood function of the SVD model is difficult to evaluate. On the other hand, the SVD model is easy to simulate. This suggests parameter estimation via simulation based techniques (see Gouriéroux and Monfort 1996). The SVD model belongs to the class of nonlinear dynamic factor models for which Gouriéroux and Jasiak (1999) propose a set of convenient, simulation based estimators. Estimation is performed in two steps. In the first step, we exploit that the marginal distribution of x_i is a

Table 1: List of Models

Short Name	Equations	Long Name
BACD	(3)-(4)-(22)	Burr ACD
GGACD	(3)-(4)-(15)	generalized gamma ACD
WACD	(3)-(4)-(26)	Weibull ACD
EACD	(3)-(4)-(29)	Exponential ACD
BLACD_1	(5)-(6)-(22)	Burr Log-ACD (type 1)
GGLACD_1	(5)-(6)-(15)	generalized gamma Log-ACD (type 1)
WLACD_1	(5)-(6)-(26)	Weibull Log-ACD (type 1)
BLACD_2	(5)-(7)-(22)	Burr Log-ACD (type 2)
GGLACD_2	(5)-(7)-(15)	generalized gamma Log-ACD (type 2)
WLACD_2	(5)-(7)-(26)	Weibull Log-ACD (type 2)
TACD	(5)-(8)-(15)	generalized gamma threshold Log-ACD (type 1)
SVD	(11)-(12)	stochastic volatility duration
WSCD	(5)-(9)-(26)	Weibull SCD

Pareto distribution of the second kind (see Appendix 1) that depends only on the parameters a and b, with density given by

$$f_{P_2}(x_i) = \frac{ab^{b+1}}{(ax_i + b)^{b+1}}. \tag{13}$$

Hence, it is straightforward to estimate a and b by quasi-maximum likelihood. In the second step, the parameter matrix Λ is estimated by the method of simulated moments, adopting the procedure outlined in Gouriéroux and Jasiak (1999, section 5.2). For the empirical application, autocorrelations of x_i and x_i^2 up to lag order 7 are selected for estimating Λ which is assumed to be a diagonal matrix.[11]

Because of analytical intractability, the SVD one-step-ahead density forecast has to be computed by simulation. The idea to generate a density forecast rests on the possibility to draw the actual state of the latent factors from the observable duration series. The next step is to simulate the movement of the latent factors in the next period. One can then exploit that in the SVD model the latent factors are directly connected to the actual duration. Given a sufficiently large number of simulations, the value of the probability integral transform can be computed based on the simulated duration forecasts. For a long sequence of forecasts this procedure can be time consuming, since it has to be carried out for each one-step-ahead forecast.

Table 1 summarizes the definitions of the models for ease of reference.

[11] In their empirical application, Ghysels, Gouriéroux, and Jasiak (1997) restricted all elements of Λ except λ_{22} to be zero.

4 Application

4.1 Data

In this section, we deal with four stocks traded on the New York Stock Exchange (NYSE): Boeing, Coca-Cola, Disney, and Exxon.[12] These stocks are very actively traded and belong to the Dow Jones index. Trading at the NYSE is based on a hybrid mechanism that combines a market maker system and an order book system. Assigned to each stock is a specialist, who makes the market for this stock, i.e. he manages the trading and quoting processes and provides liquidity when necessary by taking the other side of the trades. Apart from an opening auction, trading is continuous from 9h30 to 16h.

For each stock we define three types of durations: trade durations, price durations, and volume durations. Trade durations provided the first application of ACD models (Engle and Russell 1998). They are defined as the time spells between successive trades. When an ACD model is fitted to this set of durations, its conditional hazard provides a measure of the instantaneous trading intensity.

Price durations are defined by thinning the quote process with respect to a minimum change in the mid-price of the quotes. More precisely, a price duration is defined as the time interval needed to observe a cumulated change in the mid-price not less than a threshold that we set at \$0.125 (see Giot 1999, for a more thorough discussion of price durations). As pointed out by Engle and Russell (1998), ACD models applied to price durations provide a measure of the instantaneous volatility of the quoted mid-price process.

Finally, thinning the quote process such that the selected durations are characterized by a total traded volume equal to at least 25,000 shares defines a set of so-called volume durations.[13] Volume durations have an immediate appeal for characterizing the liquidity of a stock as they indicate the time needed to trade a given amount of shares.

For all sets of durations, the so-called interdaily durations (i.e. the durations between the first event of a day and the last recorded event of the previous day) were removed from the data. Moreover, durations corresponding to events recorded outside the regular opening hours of the NYSE were removed.

As documented in previous empirical work, the three sets of durations we use feature a strong time-of-day effect (intra-day seasonality, called diurnality), which stems from predetermined market characteristics such as opening/closing of trading or lunch time for traders. To take into account this quasi-deterministic diurnality, the time-of-day standardized durations are computed as

$$X_i = x_i \, \phi(t_i) \tag{14}$$

where X_i is the raw (trade, price or volume) duration, $\phi(t_i)$ is the time-of-day effect at time t_i, and x_i denotes the time-of-day standardized (trade, price or volume)

[12] The data were extracted from the Trade and Quote (TAQ) database. This database consists of two parts: the first reports all trades, while the second lists the bid and ask prices posted by the specialists (on the NYSE). We use data for the months of September, October, and November 1996.

[13] Gouriéroux, Jasiak and Le Fol (1999) first introduced volume durations for the trade process.

Table 2: Descriptive Statistics for the Data

	Trade durations	Price durations	Volume durations
		BOEING	
Number	23930	2620	1576
Overdispersion	1.218	1.338	0.804
$Q(10)$	1357.15	322.30	477.69
		COCA-COLA	
Number	39622	1609	3022
Overdispersion	1.181	1.171	0.876
$Q(10)$	2012.14	69.71	304.34
		DISNEY	
Number	32821	2160	1778
Overdispersion	1.226	1.209	0.720
$Q(10)$	1629.18	137.31	1136.12
		EXXON	
Number	28371	2717	2045
Overdispersion	1.194	1.196	0.658
$Q(10)$	797.21	68.18	237.85

Trade durations are the durations between consecutive trades. Price durations are the durations between bid-ask quotes such that a cumulative change in the mid-price of at least \$0.125 is observed. Volume durations are the durations between bid-ask quotes such that a cumulative traded volume of at least 25,000 shares is observed. Overdispersion is defined by the ratio standard deviation/mean (all means are equal to 1 to the third decimal). $Q(10)$ denotes the Ljung-Box Q-statistic of order 10 on the durations.

duration. The deterministic time-of-day effect for each set of durations is defined as the expected (trade, price or volume) duration conditioned on time-of-day and on the day of the week (so that, for example, the time-of-day effect of Monday can be different from the time-of-day effect of Tuesday). The expectation is computed by averaging the (trade, price or volume) durations over thirty minutes intervals for each day of the week. Cubic splines are then used on the thirty minutes intervals to smooth the time-of-day function. Hereafter, time-of-day standardized durations are called durations.[14]

As shown in Table 2, for all stocks the three types of durations feature a large Ljung-Box Q-statistic of order 10, indicating that the durations are strongly autocorrelated.[15] This feature is the strongest for trade durations. Both trade and price durations are characterized by an overdispersion index (defined as standard deviation/mean) greater than 1, while volume durations are underdispersed. In

[14]Engle and Russell (1998, p 1137) explain that the time-of-day function can be estimated jointly with the parameters of an ACD model. They also mention that proceeding in two steps (seasonal adjustment followed by estimation of the ACD on the adjusted data) gives very similar results in large samples. We stick to the more simple practice of proceeding in two steps. This is also what is done by Ghysels *et al* (1997) for the SVD, for which a joint estimation would be unfeasible.

[15]Correspondingly, autocorrelation functions for the durations (not reported) usually start at a rather small value (around 0.25) and slowly decrease to zero.

Figures 1-3, we plot the densities for the three types of durations. Volume durations are characterized by clearly hump-shaped densities, with most of the mass around 0.5-1. This is in sharp contrast with trade and price durations which exhibit essentially decreasing density functions, but with a small hump at very small durations. These salient characteristics of the data are similar to what has been documented in the literature (see Engle and Russell 1998, Giot 1999, and Bauwens and Giot 2000a). The hump close to the origin is not an artefact of the kernel density estimation of the density of a positive variable. To take this into account, we used the gammal kernel proposed by Chen (1998). The bandwith was set at $(0.9\, s\, n^{-0.2})^2$, where s is the standard deviation of the data and n the number of observations.

4.2 Results

4.2.1 Description

Density forecasting with the different models for the three types of durations gives the empirical results listed in Tables 3, 4, and 5. More specifically, for each stock and for each type of durations, we estimate the ACD, Log-ACD, TACD, SCD and SVD models summarized in Table 1 and we compute the z-series (the probability integral transforms) of each model.[16] The exercise is done 'in-sample' and 'out-of-sample'. In the former case, the estimation is done on the last third of the data, and then the forecast densities and z are computed on the same sample as used for estimation. In the latter case, the estimation is performed on the first two-thirds of the sample, and then the forecast densities and z are computed on the last third of the sample using the estimates obtained on the first part.

SVD distribution parameters could not be estimated for volume durations. The problem is caused by the first estimation step. This step serves to estimate the parameters of the Pareto distribution, which cannot be hump-shaped. However, given the humped shape of the unconditional volume duration distribution (see Figure 3), assuming a Pareto distribution for this kind of data is definitely inappropriate. What happens is that one parameter of the Pareto distribution tends to infinity during the estimation. Since one needs an estimate of this parameter to compute the z, the SVD model, in its present specification, is not appropriate for volume durations.

The results reported in Tables 3 (trade durations), 4 (price durations), and 5 (volume durations) focus on two statistics for the z-series: the goodness-of-fit (GOF) statistic for testing the null hypothesis that the distribution of the z is uniform, and the number of autocorrelations for z that are significant at the 5% level. Given the large number of models, we do not report all the results. In each table, we report systematically the results for the ACD and the two types of Log-ACD models using the Burr, the generalized gamma, and the Weibull distributions, and for the single version of the SVD, Weibull SCD and TACD models that we have presented in Section 3. We report also the results for the exponential ACD model, but not for their logarithmic counterparts since they are

[16] We stick to models with no more than one lag mainly because these are the models found in most of the papers.

quite similar.

To complement the results given in the tables, we show the histograms of z, along with the 90% confidence bands, for a subset of the models. These histograms are given in Figures 5 to 8 for trade durations, 10 to 13 for price durations, and 15 to 18 for volume durations. Since the SVD is not estimable on volume duration data, the bottom right panel of the histograms in Figures 15-18 is reserved for a specific illustration. For Boeing, it depicts the histogram for the exponential ACD model, illustrating how poorly it performs on this criterion (for the other stocks, the exponential ACD histograms look similar). For Coke, it relates to the Weibull Log-ACD model of type 2, showing that it barely differs from the Burr ACD histogram. For Exxon we choose the histogram produced by Weibull Log-ACD model of type 1. For Disney, the last panel contains the 'in-sample' histogram of the TACD, which looks even more uniform than the corresponding 'out-of-sample' histogram on its left.

Finally, we display in Figures 9, 14, and 19 the ACF of z for each stock. In each figure, the left panels correspond to the Burr ACD model - which is generally representative for the ACD/Log-ACD/TACD/SCD models - and the right panels to the SVD model (for trade and price durations). In the case of volume durations (Figure 19) the right panels are chosen to show an ACF that is relatively different from those of the Burr ACD models. For example, consider the bottom panels which are for the Exxon stock: the ACF for the Poisson model shows that the first 12 autocorrelations are outside of the 95% confidence interval, while this is not the case for the Burr ACD model. The Poisson model assumes that durations are independently and exponentially distributed. It results as a particular case of the ACD and Log-ACD models with an exponential distribution, when the parameters α and β are equal to 0.

4.2.2 General comments

1) The ACD/Log-ACD/TACD/SCD models are quite successful in capturing the dynamic structure of the durations. If that is the case, as explained in Section 2, the z sequence should be independent. As can be seen in Tables 4 and 5 (price and volume durations), the number of significant autocorrelations for z is small (often less than 5, out of 50). Regarding trade durations, this number is slightly higher, which is not very surprising as trade durations are much more autocorrelated than price or volume durations (see Table 2). As can be seen in Figures 9, 14, and 19, a general pattern for the autocorrelations of z does not emerge and most of the few significant autocorrelations are very close to the 95% confidence bands. Indeed, it is quite remarkable that the ACD, Log-ACD and TACD models capture the dynamic structure in most cases in a satisfactory way, keeping in mind that all models have a ARMA(1,1)-type lag structure. The SCD is also quite successful with only a simple AR(1) structure on ψ_i. Moreover, we stress that the models using the exponential distribution usually perform in this respect as well as the models based on more flexible distributions. Regarding in- and out-of-sample performance, the results are usually quite similar, except for the price and volume durations of the Disney stock. Thus, the models estimated on the first part of the data seem to be adequate for forecasting. The poor performance

of the SVD model, which occurs especially for the trade durations, may be due to a too restrictive dynamic specification. Choosing a richer lag structure, such as a VARMA(1,1), is most likely needed. In every case, we also checked the autocorrelations of z^2 and we found that the number of significant autocorrelations was close to the number for the z-autocorrelations.

2) The goodness-of-fit statistics deliver mixed results. For trade durations, all models are strongly rejected (the p-values are very small), but it is worth noting that the confidence bands are very narrow given the large number of observations. Furthermore, for Boeing and Exxon, some models do not always perform as badly as suggested by the very small p-values, since a reasonable number of frequencies are within the confidence bands (see Figures 5-8). For price durations, the best results are achieved by one of the ACD or Log-ACD models using the generalized gamma distribution. For each stock, there is always at least one model with a p-value larger than 5% (except for Disney, out-of-sample), and a Log-ACD model that performs better than the ACD models. The models based on the Weibull or exponential distributions are rejected, except WACD and WLACD_2 for Coke (out-of-sample). The SCD has a rather poor performance as it features a p-value of 0 for all 4 stocks.[17] The SVD model does not perform as well as the generalized gamma ACD/Log-ACD models, but it performs better than Weibull ACD/Log-ACD models. Lastly, the TACD model is far from outperforming the other models on the data we have used. For volume durations, generalized gamma and Burr models are often accepted (except for Disney out-of-sample), Weibull models perform often as well, while the exponential ACD model is clearly rejected (see the bottom right panel in Figure 15). The SCD performs better than it did for price durations but it does not give stellar results either as its p-value is usually the lowest of the group of models (EACD excluded). The corresponding histograms of z tell the same story.

3) Several models cannot account sufficiently for the very small trade, and to a lesser extent, price durations. More precisely, very small durations (i.e. durations such that the corresponding z is less than 0.05) are under-represented, compared with what is expected from the forecast densities. For example, in Figures 5 and 10, the frequency for $0 < z < 0.05$ is clearly below the lower limit of the 90% confidence interval. This defect is less serious in general for models employing the generalized gamma distribution. For trade durations, a look at the graphs reveals that the TACD and SVD models perform better than the Burr-based ACD/Log-ACD models for very small trade durations in the case of Boeing and Exxon. For Coke and Disney, no model obviously outperforms the others. Of all distributions, the Weibull distribution is the most affected. This can be understood by looking at Figure 4, where we show a few forecast densities corresponding to the estimated parameters for the Exxon trade durations (out-of-sample, $\mathrm{Log\text{-}ACD}_1$

[17] Although the use of the Weibull distribution seems to be restrictive, estimating the model with a more flexible distribution does not improve the results. For example, we tried a SCD model with the Burr distribution and the additional parameter of the Burr distribution always tended to 0 (i.e. Weibull case). This is probably due to the fact that the variance of the second error term tends to dominate any other additional parameter.

models). Because the estimate of the parameter γ of the Weibull is smaller than 1 (actually being equal to 0.93, this being required to fit the overdispersion of the data), the density tends to infinity as x tends to zero. As a consequence, there are not enough very small durations with respect to the Weibull distribution, and correspondingly there are too few very small z-values (at least in the 0 to 0.05 bin). Because the other distributions start at a finite value, they are less prone to this defect. On the other hand, small (but not very small) durations are over-represented in the data, as the frequencies associated with the second, and sometimes third and fourth histogram bins of z are usually above the confidence interval. We conclude that there is a need for a distribution that puts a lot of mass on small durations, but not too much on very small durations.

4) ACD and Log-ACD models perform globally in the same way. For each stock and type of duration, one can find at least one 'preferred' model (i.e. a model that has the highest p-value for the GOF statistic and the lowest number of significant autocorrelations). Doing this exercise for the price durations, we find that (at least) one of the Log-ACD models with the generalized gamma distribution is preferred for all stocks 'in-sample' and for 2 stocks (Coke and Exxon) 'out-of-sample', while the TACD model is the best one for Exxon (in-sample) and for Boeing and Disney (out-of-sample). For volume durations, there is more variety in the preferred models, while for trade durations the results are so poor on the GOF statistic that a ranking is not worth considering. The performance of the Weibull SCD model does not match the performances of the ACD and Log-ACD models as this model performs well only for volume durations. Surprisingly, since it is in principle more flexible and uses more parameters, the TACD model does not outperform the other ACD/Log-ACD models; in some cases, it even underperforms on the goodness-of-fit criterion. The performance of the SVD is also disappointing. In particular it is a surprise that autocorrelagrams of the z^2 statistics (not reported) are not improved by this model since its main objective is to account for a stochastic volatility effect in the durations. However, there is room for improving the specification, for example by employing a richer dynamic structure than VAR(1). Given that TACD, SCD and SVD models are much more costly to estimate and to evaluate than ACD and Log-ACD models, we tend to conclude that they should not be considered as a starting point in an empirical research.

5 Conclusion

Density forecasts are a powerful tool to evaluate financial duration models models in- and out-of-sample. We have compared ACD, Log-ACD, TACD, SCD, and SVD models based on a variety of distributions for the innovations, which have been proposed in the recent literature. We have found that these models (except the SVD) are generally able to account well for the dynamic structure of the data even if the distribution of the innovations is not correctly specified, and that diagnostics based on the probability integral transform can reveal a problem such as a structural break (as seems to be the case for the Disney stock). Concerning the (conditional) distribution of the durations, we have found that for most data sets the ACD and Log-ACD models based on the Weibull (and exponential) distributions are misspecified for all types of durations (trade, price, volume). The Burr and generalized gamma-based models are correctly specified for volume durations, and for price durations in several cases. We found that the main problem of financial duration models is to account for the peculiar shape of the empirical distribution at small durations. Log-ACD models perform as well as their ACD counterparts, which is a useful result for empirical financial microstructure modeling, since the Log-ACD specifications do not require positivity restrictions on the parameters of the autoregressive equation.

Concerning the trade durations, further work is needed to determine the reasons for the rejection of the distributional hypothesis. A more flexible distribution than the Burr and the generalized gamma is needed, unless one wants to rely on a semiparametric approach.[18] Other factors could be in play. In particular, given the large sample sizes, changes of regimes may have occured. In that respect, the TACD approach to regime switching does not yield generally an improved performance compared to simpler models. Another factor to be considered is the procedure used for eliminating the intra-day seasonality. Indeed, it should be kept in mind that the durations used for evaluating the models are the outcome of such a procedure. Thus it is possible that some of the problems related to the distribution of the durations could be alleviated by using a more refined procedure.

Several extensions are on our research agenda. Firstly, we shall use Monte Carlo experiments to examine the properties of the (GOF and autocorrelation) diagnostic statistics we have used. Secondly, such experiments could also be used to analyze the properties of the estimators of the parameters (and of the diagnostic statistics) of a misspecified model. In particular, it remains to be seen if one can discriminate between closely related specifications (e.g. ACD versus Log-ACD, or Burr versus generalized gamma). Thirdly, we could compare the models using other econometric tools (such as information criteria and test statistics for non-nested models) than the diagnostics we have used. Fourthly, we shall check the sensitivity of the results to different procedures for adjusting the data for the intra-day seasonality.

Appendix 1: Distributions

Generalized gamma distribution

We define the generalized gamma density function for $\epsilon > 0$ as

$$f_{GG}(\epsilon) = \frac{c^{-\gamma\nu}\gamma\epsilon^{\gamma\nu-1}}{\Gamma(\nu)}\exp(-(\epsilon/c)^{\gamma}), \text{ for } \nu > 0, \gamma > 0, c > 0, \tag{15}$$

where $\Gamma(.)$ denotes the usual gamma function. The parameter c is a scale parameter that we normalize to 1 when using this distribution and its particular cases with ACD and Log-ACD models. The uncentered moment of order p is given by

$$\mu_p = c^p \frac{\Gamma(\nu + \frac{p}{\gamma}),}{\Gamma(\nu)} \tag{16}$$

[18]See Drost and Werker (2000).

and the cdf by

$$F_{GG}(\epsilon) = \frac{\Gamma\left(\nu, (\epsilon/c)^{\gamma}\right)}{\Gamma(\nu)}, \tag{17}$$

where

$$\Gamma(\nu, x) = \int_{0}^{x} u^{\nu-1} e^{-u} du, \tag{18}$$

which requires a numerical integration.

Gamma distribution

For $\gamma = 1$, the generalized gamma density defined in (15) reduces to the gamma density with shape parameter ν and scale parameter c:

$$f_G(\epsilon) = C_G^{-1}(\nu, c)\epsilon^{\nu-1}\exp(-\epsilon/c), \tag{19}$$

i.e. $\epsilon \sim \text{gamma}(\nu, c)$, where

$$C_G(\nu, c) = \Gamma(\nu)c^{\nu}. \tag{20}$$

Clearly, if $\epsilon \sim \text{gamma}(\nu, c)$, then $\epsilon/c \sim \text{gamma}(\nu, 1)$. The uncentered moment of order p of ϵ is given by

$$\mu_p = \frac{C_G(\nu + p, c)}{C_G(\nu, c)} = c^p \frac{\Gamma(\nu + p)}{\Gamma(\nu)} \quad \text{for } \nu + p > 0. \tag{21}$$

and the cdf by (17) with $\gamma = 1$.

Burr distribution

We define the Burr density function for $\epsilon > 0$ as

$$f_B(\epsilon) = \frac{\gamma}{c}\left(\frac{\epsilon}{c}\right)^{\gamma-1}\left[1 + \sigma^2\left(\frac{\epsilon}{c}\right)^{\gamma}\right]^{-(1+\sigma^{-2})}, \text{ for } \gamma > \sigma^2 > 0, \tag{22}$$

where

$$c = \mu_1 \frac{(\sigma^2)^{1+\gamma^{-1}}\Gamma(1 + \sigma^{-2})}{\Gamma(1 + \gamma^{-1})\Gamma(\sigma^{-2} - \gamma^{-1})} \text{ for } \mu_1 > 0. \tag{23}$$

The uncentered moment of order p is given by

$$\mu_p = c^p \frac{\Gamma(1 + p/\gamma)\Gamma(\sigma^{-2} - p/\gamma)}{(\sigma^2)^{1+p/\gamma}\Gamma(1 + \sigma^{-2})}, \text{ for } \gamma > p, \tag{24}$$

so that the parameter μ_1 appearing in (23) is indeed the mean (if $\gamma > 1$). The Burr cdf is

$$F_B(\epsilon) = 1 - (1 + \sigma^2 c^{-\gamma}\epsilon^{\gamma})^{-1/\sigma^2}. \tag{25}$$

See Lancaster (1990: p68) for a derivation of the Burr distribution as a gamma mixture of Weibull distributions.

Weibull distribution

For $\nu = 1$, the generalized gamma density defined in (15) reduces to the Weibull density with shape parameter γ and scale parameter c:

$$f_W(\epsilon) = \frac{\gamma}{c}\left(\frac{\epsilon}{c}\right)^{\gamma-1} \exp(-(\epsilon/c)^{\gamma}). \tag{26}$$

It is also obtained from the Burr density defined in (22) by letting σ^2 tend to 0. The uncentered moment of order p is given by

$$\mu_p = c^p\, \Gamma(1 + \frac{p}{\gamma}), \tag{27}$$

and the cdf by

$$F_W(\epsilon) = 1 - \exp\left(-(\epsilon/c)^{\gamma}\right). \tag{28}$$

Exponential distribution

For $\gamma = 1$, the Weibull density reduces to the exponential density with parameter c:

$$f_E(\epsilon) = \frac{1}{c}\exp\left(-\frac{\epsilon}{c}\right). \tag{29}$$

It is also a gamma(1,c) distribution. The uncentered moment of order p is given by

$$\mu_p = c^p\, \Gamma(1 + p), \tag{30}$$

and the cdf by

$$F_E(\epsilon) = 1 - \exp(-\epsilon/c). \tag{31}$$

Pareto distribution

The Pareto density function for the random variable $\epsilon > k$, with parameters $\theta > 0$ and $k > 0$, is given by

$$f_{P_1}(\epsilon) = \frac{\theta k^{\theta}}{\epsilon^{\theta+1}}, \tag{32}$$

Its uncentered moments of order p are given by

$$\mu_p = \frac{\theta k^p}{\theta - p} \text{ if } \theta > p, \tag{33}$$

and its cdf by

$$F_{P_1}(\epsilon) = 1 - (k/\epsilon)^{\theta}. \tag{34}$$

The above distribution is known as the Pareto of the first kind (see Johnson *et al.* 1994, pp 574-575). The Pareto distribution of the second kind results from the transformation $x_i = \epsilon - k$, so that its density function is

$$f_{P_2}(x_i) = \frac{\theta k^{\theta}}{(x_i + k)^{\theta+1}}. \tag{35}$$

As reminded in Section 3.4, this distribution can be obtained as the mixture of an exponential distribution by a gamma distribution. Formula (35) corresponds to formula (13) for $k = b/a$ and $\theta = b$.

Appendix 2: Propositions on density forecasts

In the following we collect the propositions on which Diebold, Gunther and Tay's (1998) method to evaluate density forecasts is based. Let $L(a, x)$ denote a forecast user's loss function where a refers to an action choice and x the realisation. If the user believes that the density forecast $f(x)$ is correct, then the action a_* is chosen such that[19]

$$a_*(p(x)) = \underset{a \in A}{\operatorname{argmin}} \int L(a, x) f(x) dx. \tag{36}$$

Proposition 1 states that there is no way to rank to incorrect density forecasts such that all users will agree with the ranking.

Proposition 1: Let $p(x)$ be the density of x, let $a_*^{(j)}$ be the optimal action based on forecast $f^{(j)}$ and let $a_*^{(k)}$ be the optimal action based on forecast $f^{(k)}$. Then there does not exist a ranking r of arbitrary density forecasts $f^{(j)}$ and $f^{(k)}$, both distinct from f such that for all loss functions $L(a, x)$,

$$r_j \geq r_k \Leftrightarrow \int L\left(a_*^{(j)}, x\right) p(x) dx \geq \int L\left(a_*^{(k)}, x\right) p(x) dx. \tag{37}$$

This result seems to rule out the possibility to evaluate density forecasts issued by competing models. However, Diebold, Gunther and Tay (1998) point out that if a forecast coincides with the true data generating process, then it will be preferred by all forecast users, regardless of loss function. This is formalized in the following proposition that is also contained in Granger and Pesaran (1996):

Proposition 2: Suppose that $f^{(j)}(x) = p(x)$, so that $a_*^{(j)}$ minimizes the expected loss with respect to the true distribution. Then

$$\int L\left(a_*^{(j)}, x\right) p(x) dx \leq \int L\left(a_*^{(k)}, x\right) p(x) dx, \forall k \tag{38}$$

The result follows from the assumption that $a_*^{(j)}$ minimizes expected loss over all possible actions, including those that might be chosen under alternative density forecasts.

In other words, regardless of loss function, the correct density is weakly superior to all other forecasts. This result provides the motivation for evaluating density forecasts by assessing whether $f_i(x_i \mid \mathcal{H}_i) = p_i(x_i \mid \mathcal{H}_i)$.[20]

The follwing lemma, dating back to Rosenblatt (1952) is used to circumvent the problem that $p_i(x_i \mid \mathcal{H}_i)$ is never observed:

Lemma 1: Let $p(x)$ be the true density of x, let $f(x)$ be the density forecast of x, and let z be the probability integral transform of x with respect to $f(x)$. Then assuming that $\frac{\partial F^{-1}(z)}{\partial z}$ is continuous and nonzero over the support of x, z has support on the unit interval with density

$$q(z) = \left| \frac{\partial F^{-1}(z)}{\partial z} \right| p(F^{-1}(z)) = \frac{p(F^{-1}(z))}{f(F^{-1}(z))}. \tag{39}$$

[19] To be able to assume a unique minimizer it is sufficient that A be compact and that L is strictly convex in A.

[20] If this is not the case then some forecast users could potentially be better served by a different density forecast (depending on their loss function).

The important result is that if $f(x) = p(x)$ then $q(x)$ is simply the $U(0,1)$ density. Based on Lemma 1, Diebold, Gunther and Tay (1998) characterize both density and dependence structure of the entire z sequence when $\{f_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m = \{p_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ (adopting the extensive notation again).

Proposition 3: Suppose $\{x_i\}_{i=1}^m$ is generated from $\{p_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ where $\mathcal{H}_i = \{x_{i-1}, x_{i-2}, ...\}$. If a sequence of density forecasts $\{f_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ coincides with $\{p_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$, then under the usual condition of non-zero Jacobian with continuous partial derivatives, the sequence of probability integral transforms of $\{x_i\}_{i=1}^m$ with respect to $\{p_i(x_i \mid \mathcal{H}_i)\}_{i=1}^m$ is IID $U(0,1)$. That is,

$$\{z_i\}_{i=1}^m \sim \text{IID}\, U(0,1). \tag{40}$$

References

Bauwens, L. and Giot, P. (2000a). 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. and Giot, P. (2000b). The moments of first-order Log-ACD models. Mimeo. CORE, Louvain-La-Neuve.

Bauwens, L. and Giot, P. (2001). *Econometric Modelling of Stock Market Intraday Activity.* Forthcoming. Kluwer Academic Publishers.

Bauwens, L. and Veredas, D. (1999). The stochastic conditional duration model: a latent factor model for the analysis of financial durations. CORE Discussion Paper 9958. Université catholique de Louvain, Louvain-La-Neuve.

Chen, S.X. (1998). Probability density function estimation using gamma kernels. Manuscript. School of Statistical Science, La Trobe University.

Diebold, F.X., Gunther, T.A., and Tay, A.S. (1998). Evaluating density forecasts, with applications to financial risk management. *International Economic Review* 39, 863-883.

Drost,F.C. and Werker, B.J.M. (2000). Efficient estimation in semiparametric time series: the ACD model. Mimeo. CentER, Tilburg University.

Engle, R.F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50, 987-1007.

Engle, R.F. and Russell, J.R. (1998). Autoregressive conditional duration: a new approach for irregularly spaced transaction data. *Econometrica* 66, 1127-1162.

Giot, P. (1999). Time transformations, intra-day data and volatility models. Forthcoming in *Journal of Computational Finance.*

Ghysels, E., Gouriéroux, C., and Jasiak, J. (1997). Stochastic volatility duration models. Working Paper 9746. CREST, Paris.

Gouriéroux, C. and Jasiak, J.(1999). Dynamic Factor Models. CREST, Paris.

Gouriéroux, C. and Monfort, A. (1996). *Simulation-based Econometric Methods.* Oxford: Oxford University Press.

Gouriéroux, C., Jasiak, J. and Le Fol, G. (1999). Intra-day market activity. *Journal of Financial Markets* 2, 193-226.

Goodhart, C.A.E. and O'Hara, M. (1997). High frequency data in financial markets: issues and applications. *Journal of Empirical Finance* 4, 73-114.

Grammig, J. and Maurer, K.-O. (1999). Non-monotonic hazard functions and the autoregressive conditional duration model. Discussion Paper 50, SFB 373, Humboldt University Berlin. Forthcoming in *The Econometrics Journal.*

Granger, C.W.J. and M.H. Pesaran (1996). A decision theoretic approach to forecast evaluation. Working Paper 9618. Department of Applied Economics, University of Cambridge.

Johnson, N.L., Kotz, S., and Balakrishnan, N. (1994). *Distributions in Statistics: Continuous Univariate Distributions, Vol. 1 (2nd ed.).* New York: John Wiley and Sons.

Kim, S., N. Shephard, and S. Chib (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies* 65, 361-393.

Lancaster, T. (1990). *The Econometric Analysis of Transition Data.* Cambridge: Cambridge University Press.

Lunde, A. (1999). A generalized gamma autoregressive conditional duration model. Working Paper. Aalborg University.

O'Hara, M. (1995). *Market Microstructure Theory.* Oxford: Basil Blackwell.

Rosenblatt, M. (1952). Remarks on a multivariate transformation. *Annals of Mathematical Statistics* 23, 470-472.

Shephard, N. (1994). Partial non-gaussian state space. *Biometrika* 81, 115-131.

Zhang, M.Y., Russell, J.R., and Tsay, R.T. (1999). A nonlinear autoregressive conditional duration model with applications to financial transaction data. Mimeo. Graduate School of Business, University of Chicago.

Table 3: Results for Trade Durations

Stock	Model	In-Sample GOF	In-Sample AC(z)	Out-of-sample GOF	Out-of-sample AC(z)
	BACD	0.0	3	0.0	2
	GGACD	0.0	2	0.0	3
	WACD	0.0	3	0.0	2
	EACD	0.0	3	0.0	1
BA	$BLACD_1$	0.0	1	0.0	1
(7978)	$GGLACD_1$	0.0	1	0.0	2
	$BLACD_2$	0.0	2	0.0	2
	$GGLACD_2$	0.0	3	0.0	1
	TACD	0.0	1	0.0	3
	SVD	0.0	33	0.0	30
	WSCD	0.0	7	0.0	1
	BACD	0.0	5	0.0	5
	GGACD	0.0	6	0.0	5
	WACD	0.0	5	0.0	5
	EACD	0.0	6	0.0	6
KO	$BLACD_1$	0.0	5	0.0	4
(13205)	$GGLACD_1$	0.0	5	0.0	5
	$BLACD_2$	0.0	5	0.0	6
	$GGLACD_2$	0.0	6	0.0	4
	TACD	0.0	6	0.0	4
	SVD	0.0	23	0.0	18
	WSCD	0.0	9	0.0	5
	BACD	0.0	10	0.0	9
	GGACD	0.0	6	0.0	17
	WACD	0.0	9	0.0	8
	EACD	0.0	9	0.0	8
DIS	$BLACD_1$	0.0	9	0.0	9
(10941)	$GGLACD_1$	0.0	6	0.0	10
	$BLACD_2$	0.0	9	0.0	7
	$GGLACD_2$	0.0	8	0.0	9
	TACD	0.0	4	0.0	4
	SVD	0.0	46	0.0	47
	WSCD	0.0	7	0.0	15
	BACD	0.0	6	0.0	4
	GGACD	0.0	5	0.0	10
	WACD	0.0	6	0.0	4
	EACD	0.0	6	0.0	4
	$BLACD_1$	0.0	6	0.0	8
XON	$GGLACD_1$	0.0	6	0.0	7
(9458)	$WLACD_1$	0.0	5	0.0	6
	$BLACD_2$	0.0	6	0.0	4
	$GGLACD_2$	0.0	6	0.0	4
	$WLACD_2$	0.0	6	0.0	4
	TACD	0.0	5	0.0	8
	SVD	0.0	21	0.0	21
	WSCD	0.0	11	0.0	8

GOF: 100 × p-value of the $\chi^2(19)$-test statistic for a uniform distribution of z. AC(z): number of significant autocorrelations (out of 50) for z (5%-level). Models are defined in Table 1.

Table 4: Results for Price Durations

Stock	Model	In-Sample		Out-of-sample	
		GOF	AC(z)	GOF	AC(z)
BA (874)	BACD	1.5	4	0.01	4
	GGACD	1.0	4	0.04	5
	WACD	0.0	4	0.0	4
	EACD	0.0	4	0.0	4
	$BLACD_1$	2.5	4	0.01	4
	$GGLACD_1$	29	4	2.1	1
	$WLACD_1$	0.0	3	0.0	4
	$BLACD_2$	1.8	5	0.0	4
	$GGLACD_2$	72	5	6.6	4
	$WLACD_2$	0.0	4	0.0	4
	TACD	0.0	7	19	4
	SVD	5.3	6	2.7	5
	WSCD	0.0	8	0.0	5
KO (537)	BACD	6.4	0	6.1	1
	GGACD	10	0	4.6	1
	WACD	0.8	0	11	0
	EACD	0.01	0	1.8	0
	$BLACD_1$	6.8	0	9.2	0
	$GGALCD_1$	41	0	51	0
	$BLACD_2$	1.7	0	15	1
	$GGLACD_2$	11	0	71	1
	$WLACD_2$	0.8	0	16	1
	TACD	0.0	3	9.2	0
	SVD	4.7	0	4.5	1
	WSCD	0.0	3	0.7	0
DIS (721)	BACD	59	1	0.0	25
	GGACD	0.4	2	0.0	24
	WACD	0.01	2	0.0	21
	EACD	0.02	2	0.0	21
	$BLACD_1$	16	1	0.0	22
	$GGLACD_1$	78	1	0.0	27
	$WLACD_1$	0.0	1	0.0	21
	$BLACD_2$	19	2	0.0	21
	$GGLACD_2$	76	1	0.0	24
	$WLACD_2$	0.0	2	0.0	16
	TACD	23	1	6.3	31
	SVD	19	11	0.0	31
	WSCD	0.0	2	0.0	26
XON (907)	BACD	20	2	5.5	2
	GGACD	20	2	88	4
	WACD	0.0	2	0.0	2
	EACD	0.0	2	0.02	2
	$BLACD_1$	11	2	15	2
	$GGLACD_1$	53	2	50	2
	$BLACD_2$	18	2	15	2
	$GGLACD_2$	67	2	15	2
	TACD	15	1	10	5
	SVD	0.5	4	0.01	4
	WSCD	0.0	1	0.03	3

See Table 3 for explanations.

Table 5: Results for Volume Durations

Stock	Model	In-Sample GOF	In-Sample AC(z)	Out-of-sample GOF	Out-of-sample AC(z)
BA (526)	BACD	42	2	10	3
	GGACD	38	2	7.9	3
	WACD	52	2	6.5	3
	EACD	0.0	3	0.0	3
	BLACD$_1$	67	4	19	4
	GGLACD$_1$	65	4	15	4
	WLACD$_1$	60	4	25	4
	BLACD$_2$	24	2	12	2
	GGLACD$_2$	47	2	4.1	2
	TACD	21	1	0.0	3
	WSCD	4.3	2	18.1	3
KO (1008)	BACD	7.4	1	6.0	1
	GGACD	4.2	1	3.2	1
	WACD	4.5	1	3.2	1
	EACD	0.0	2	0.0	2
	BLACD$_1$	32	2	1.6	2
	GGLACD$_1$	47	2	1.1	2
	BLACD$_2$	8.1	1	1.2	1
	GGLACD$_2$	3.1	1	1.6	1
	WLACD$_2$	10	1	0.6	1
	TACD	0.0	1	0.2	2
	WSCD	11.7	4	43.3	2
DIS (594)	BACD	37	2	0.0	12
	GGACD	16	2	0.0	12
	WACD	41	2	0.0	14
	EACD	0.0	2	0.0	8
	BLACD$_1$	30	1	0.0	11
	GGLACD$_1$	38	1	0.0	11
	WLACD$_1$	74	1	0.0	12
	BLACD$_2$	4.4	1	0.0	5
	GGLACD$_2$	10	1	0.0	6
	WLACD$_2$	44	1	0.0	6
	TACD	19	2	31	5
	WSCD	0.03	3	0.0	6
XON (683)	BACD	99	2	99	1
	GGACD	94	2	96	1
	WACD	98	2	92	1
	EACD	0.0	1	0.0	1
	BLACD$_1$	35	2	89	1
	GGLACD$_1$	68	2	78	1
	WLACD$_1$	77	1	89	1
	BLACD$_2$	96	2	98	1
	GGLACD$_2$	95	2	97	1
	TACD	81	1	5.8	1
	WSCD	26	1	20	1

See Table 3 for explanations.

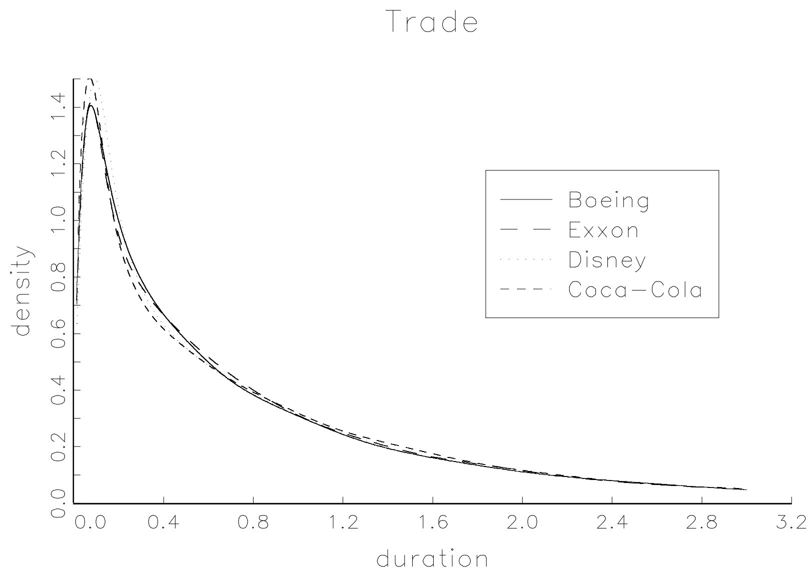

Figure 1: Kernel Densities of Trade Durations

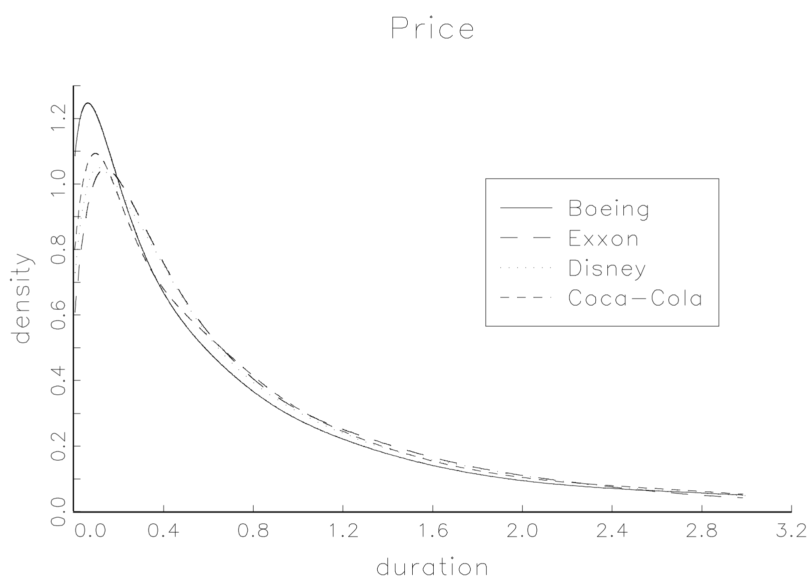

Figure 2: Kernel Densities of Price Durations

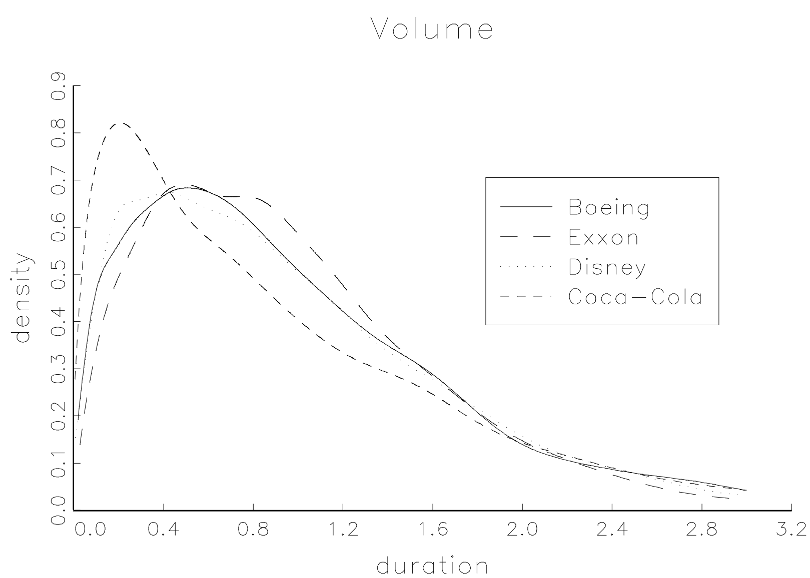

Figure 3: Kernel Densities of Volume Durations

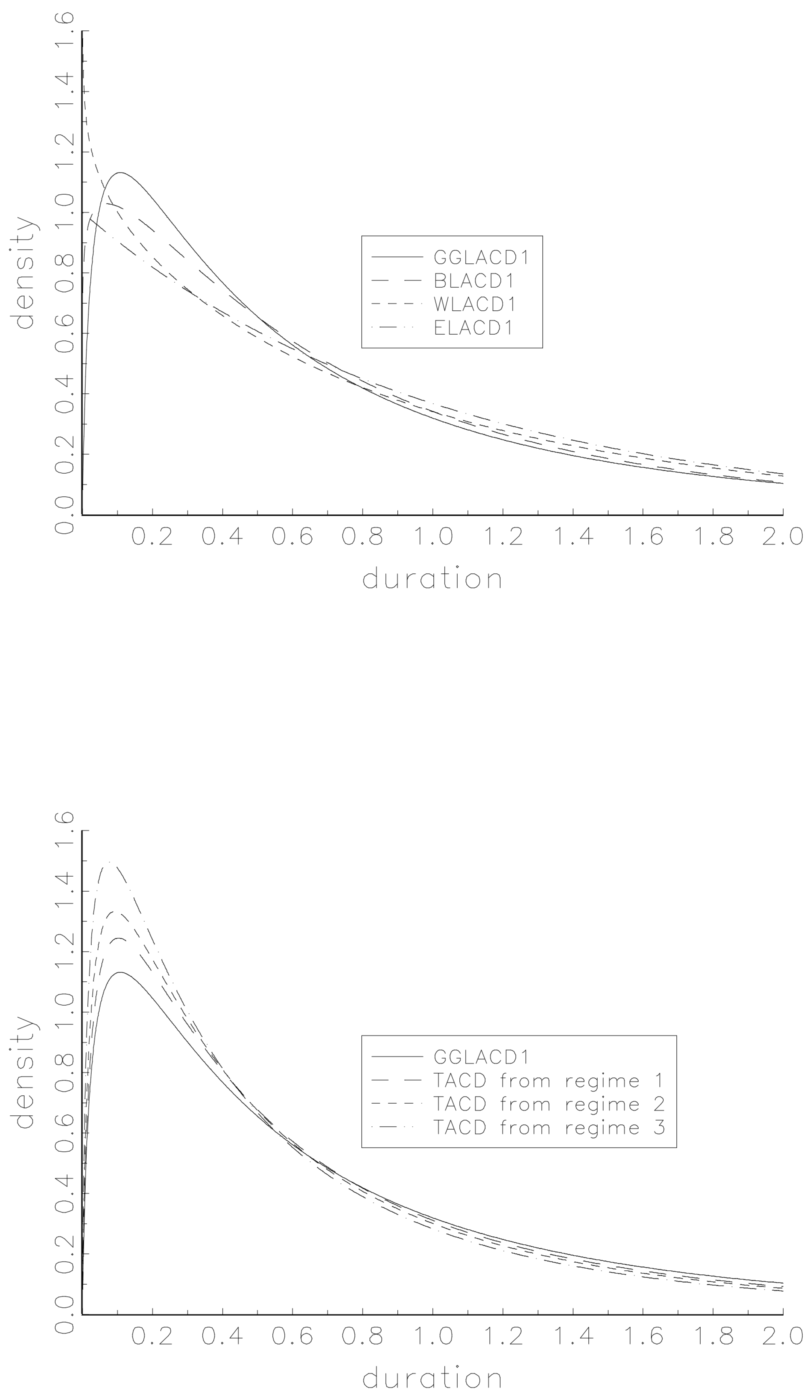

All densities have their mean equal to 1.

Figure 4: Forecast Densites of Log-ACD model of type 1 (Exxon Trade Durations)

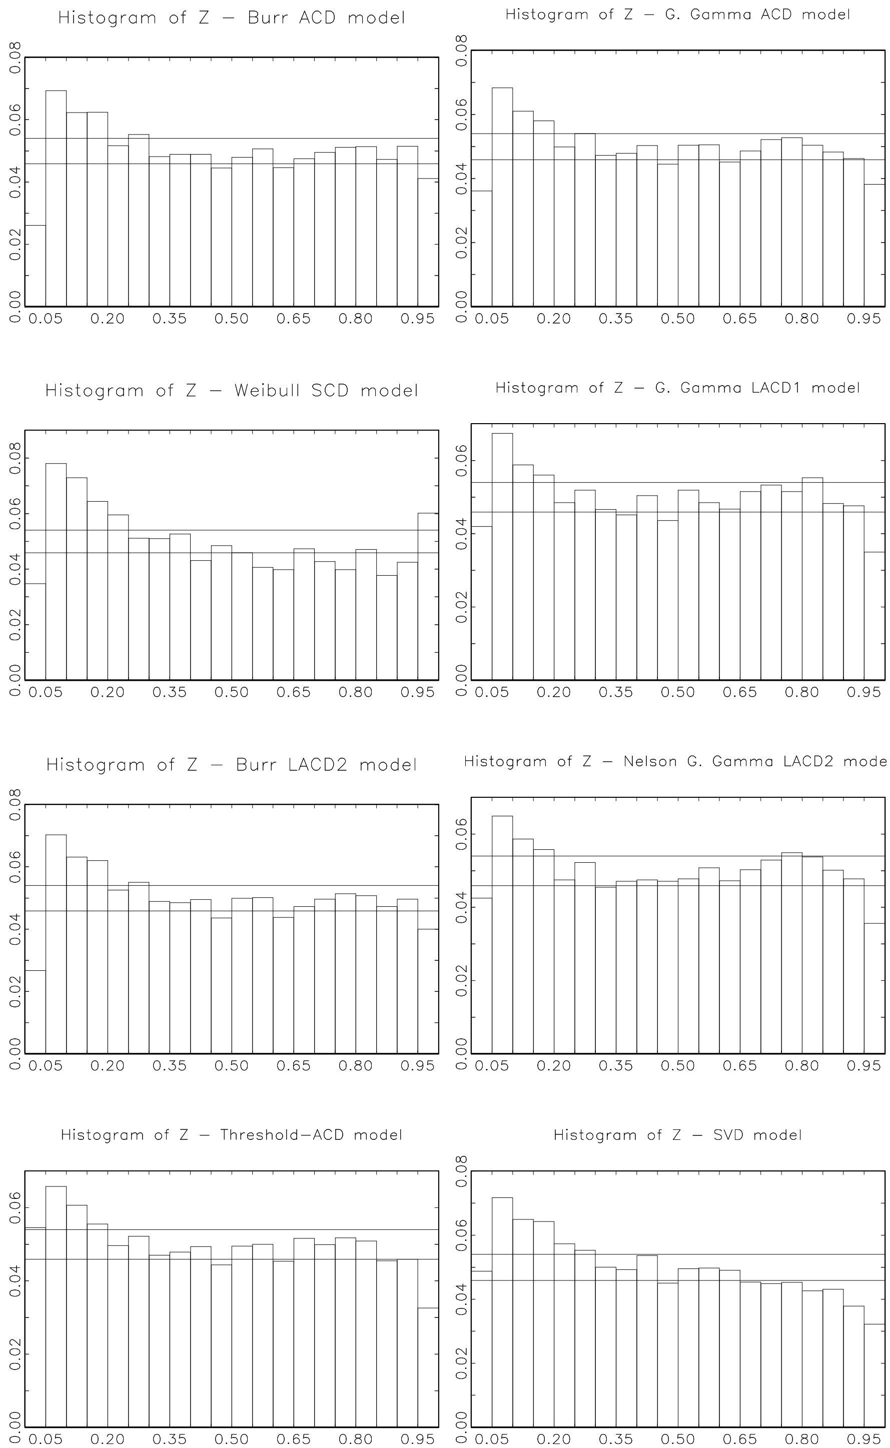

Figure 5: GOF for Boeing Trade Durations (out-of-sample)

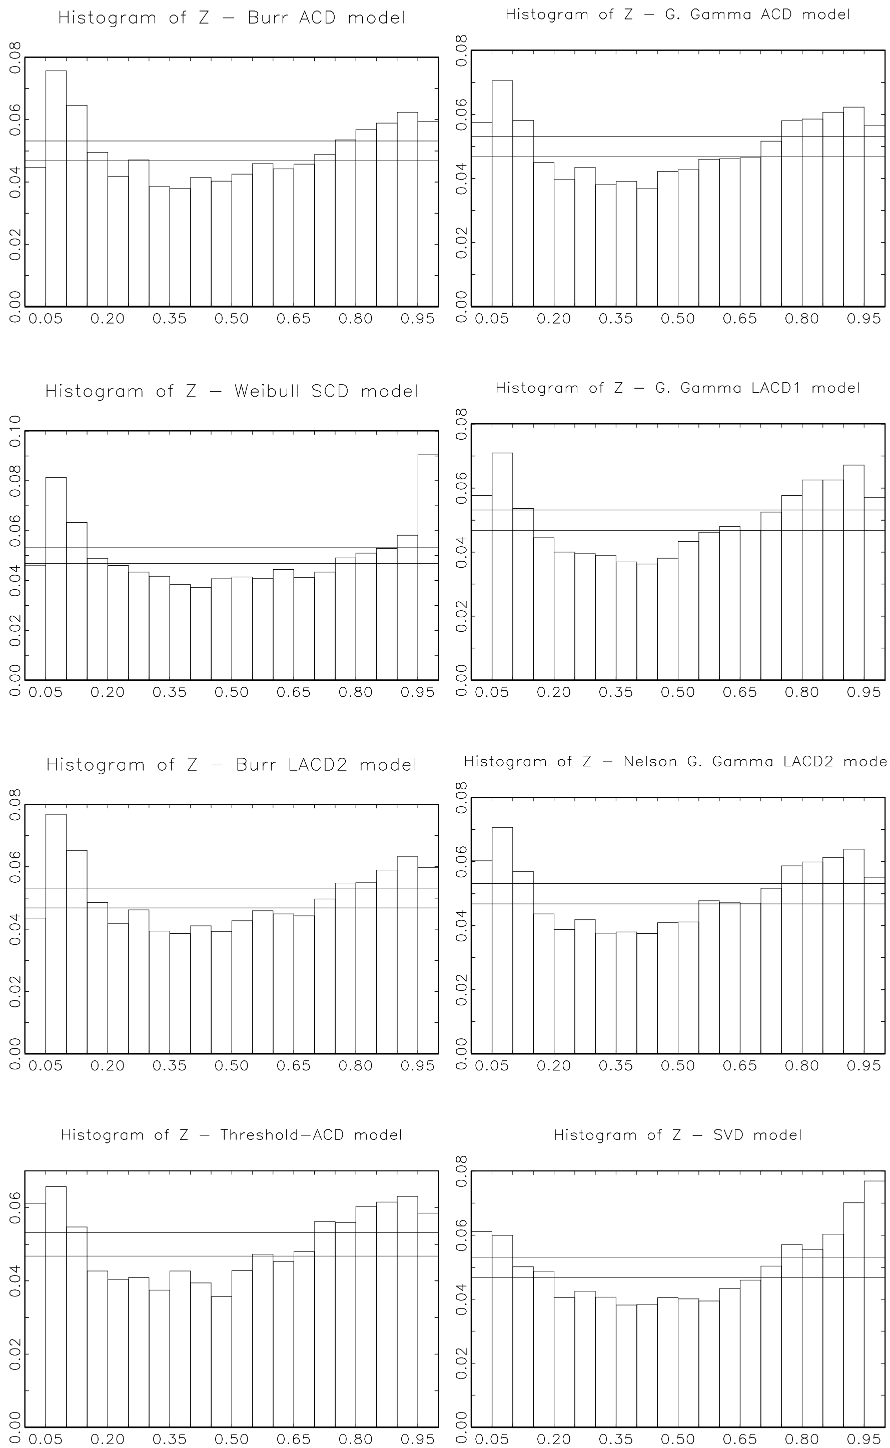

Figure 6: GOF for Coca-Cola Trade Durations (out-of-sample)

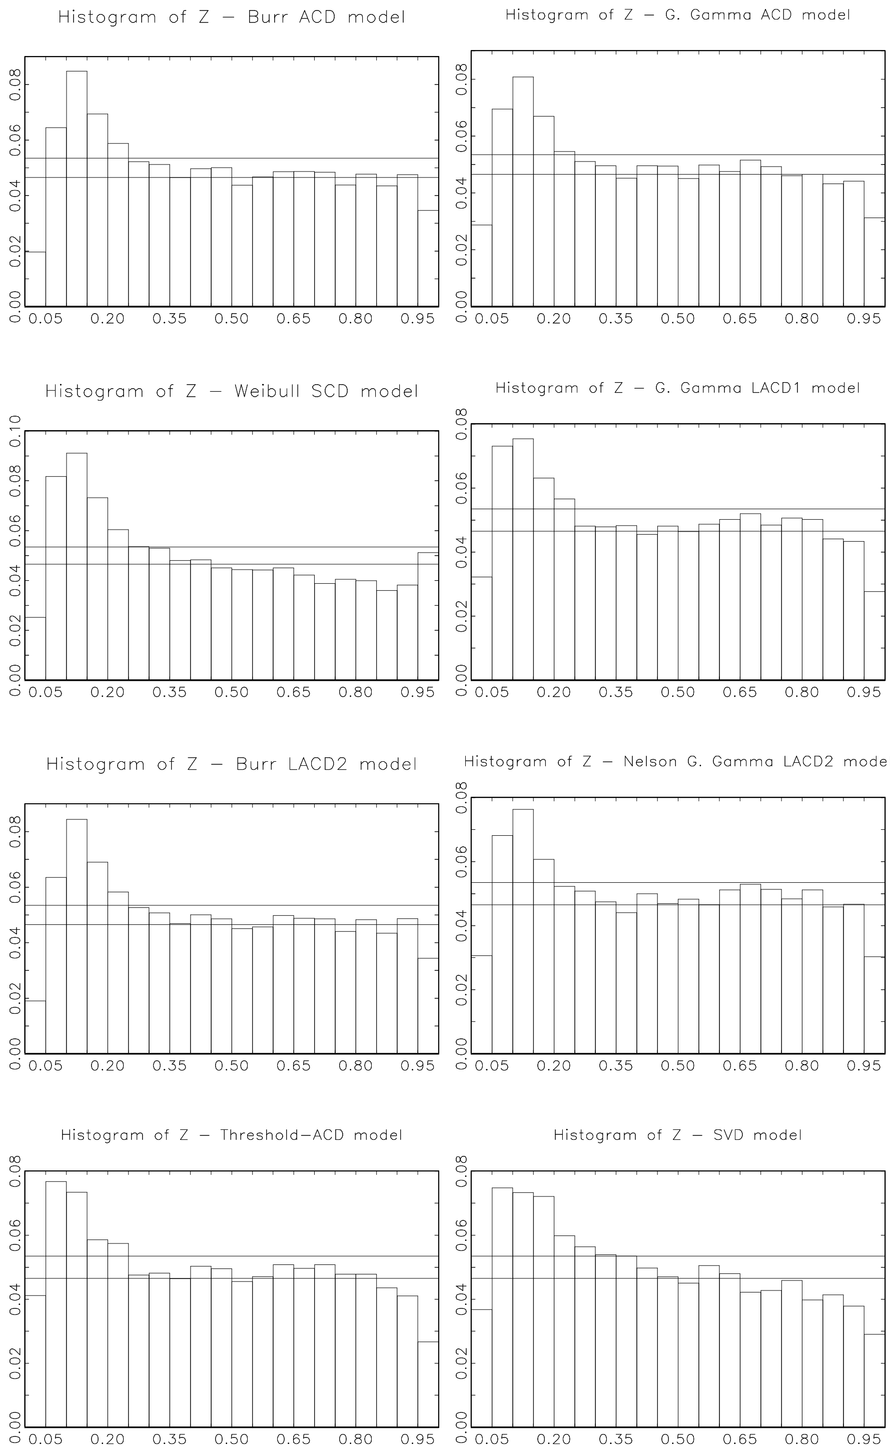

Figure 7: GOF for Disney Trade Durations (out-of-sample)

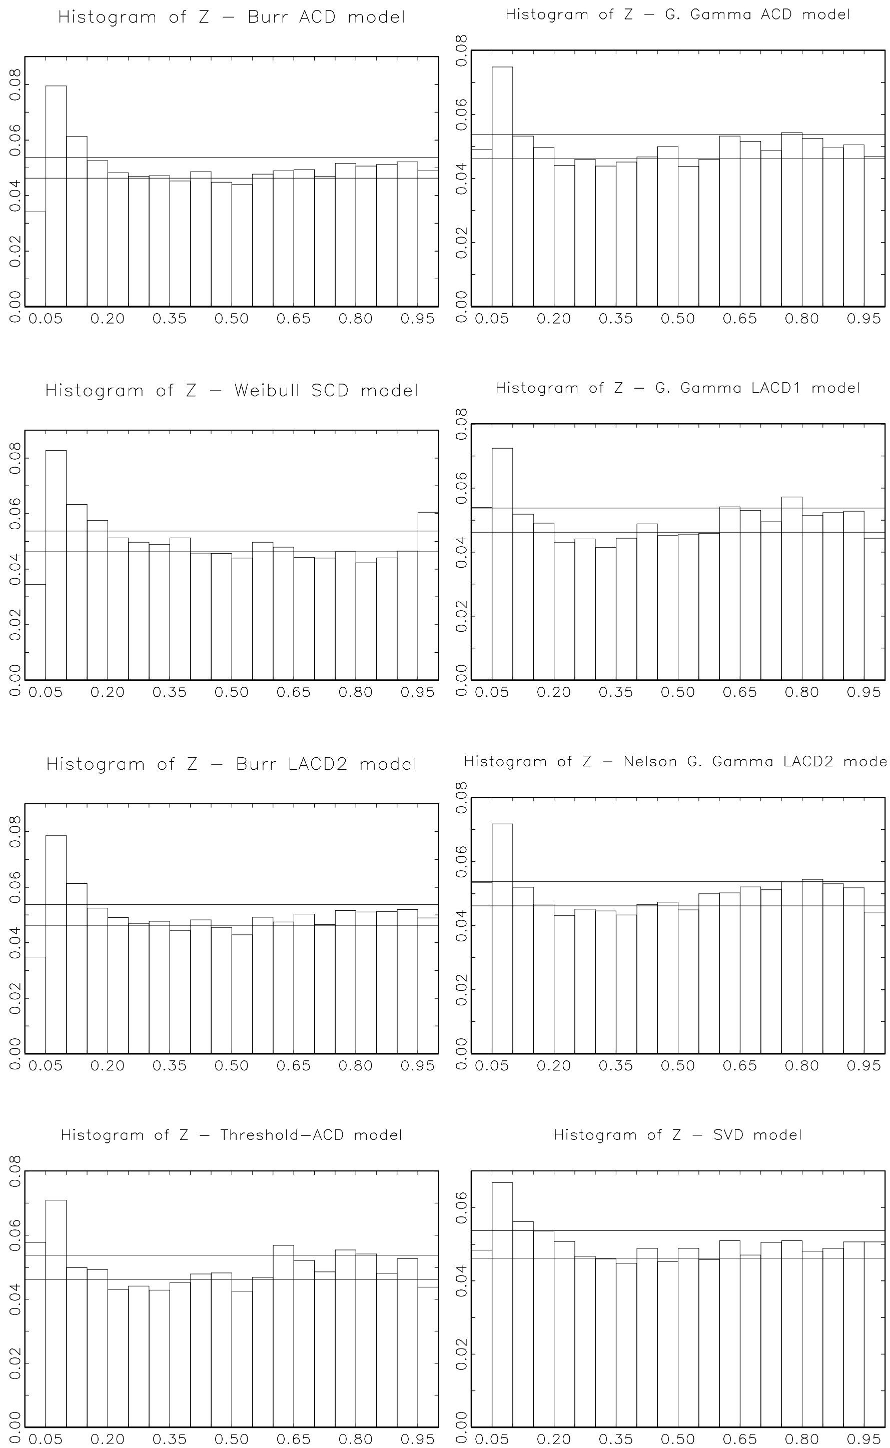

Figure 8: GOF for Exxon Trade Durations (out-of-sample)

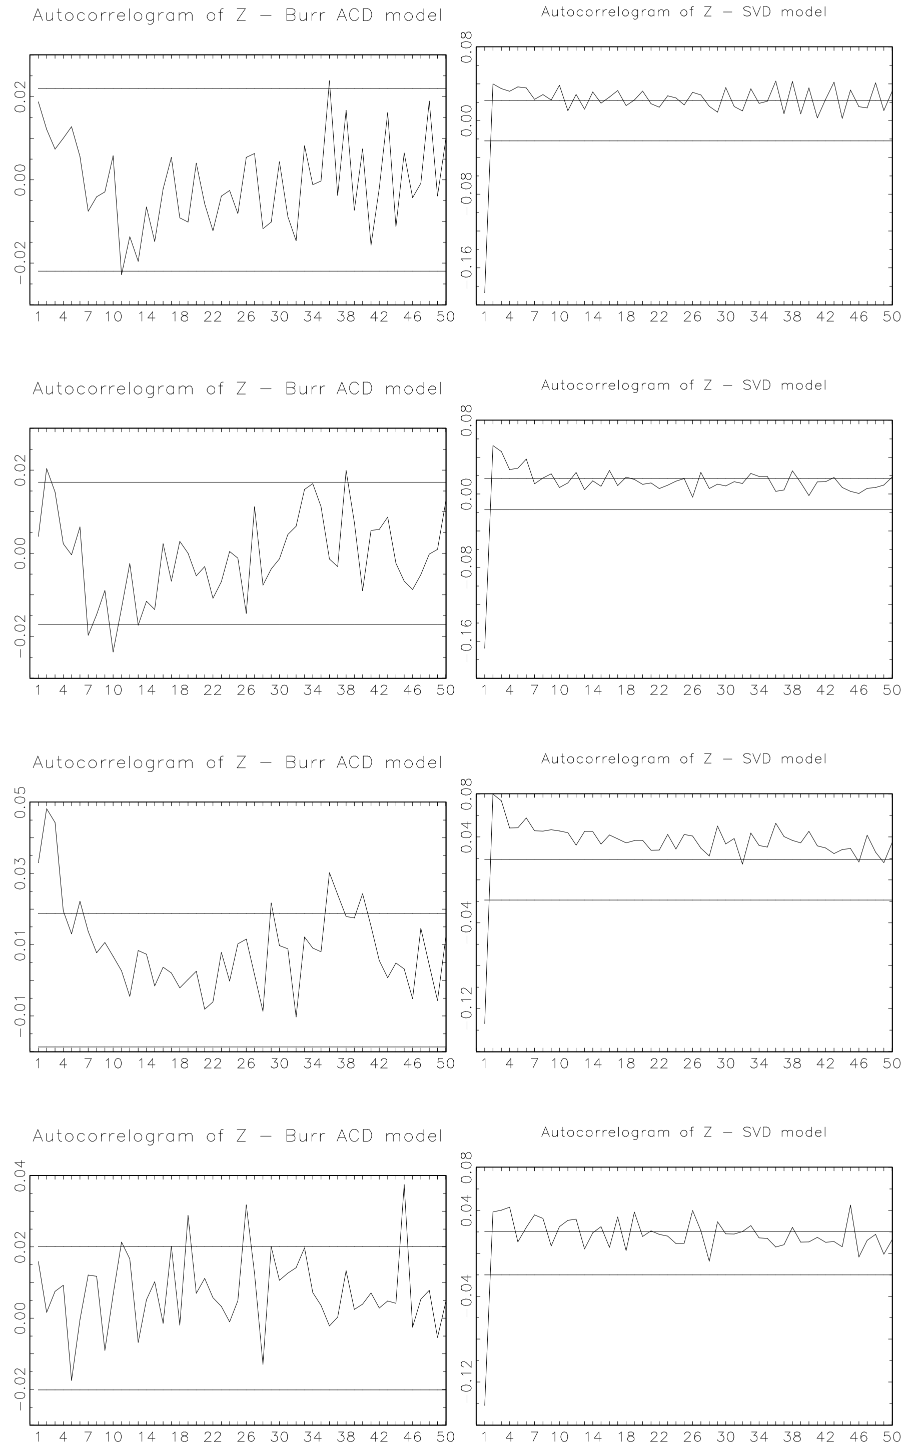

From up to down: Boeing, Coke, Disney, Exxon.

Figure 9: z-Correlograms for Trade Durations (out-of-sample)

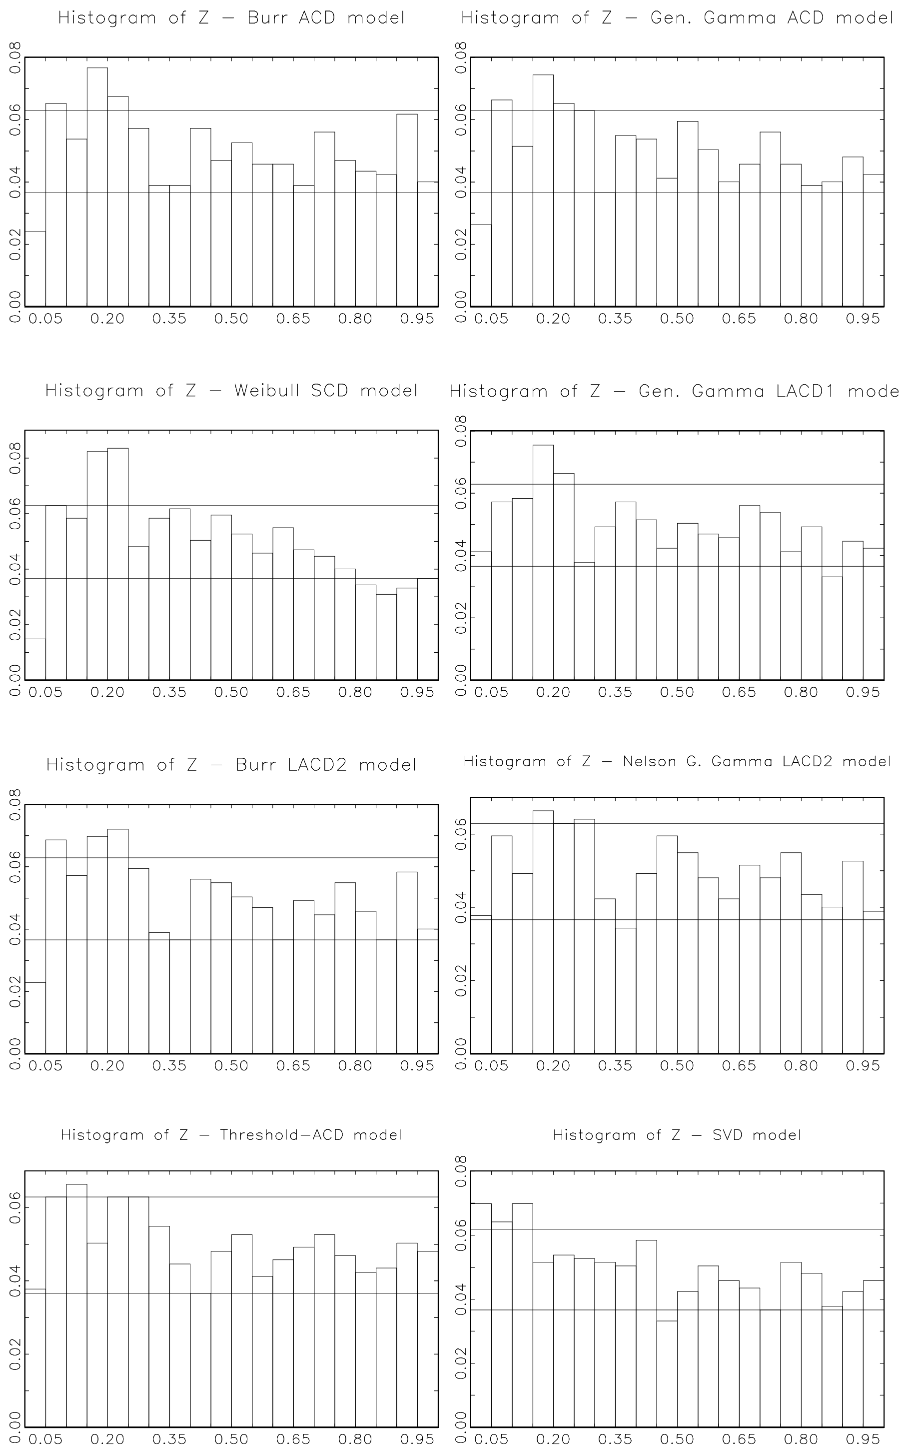

Figure 10: GOF for Boeing Price Durations (out-of-sample)

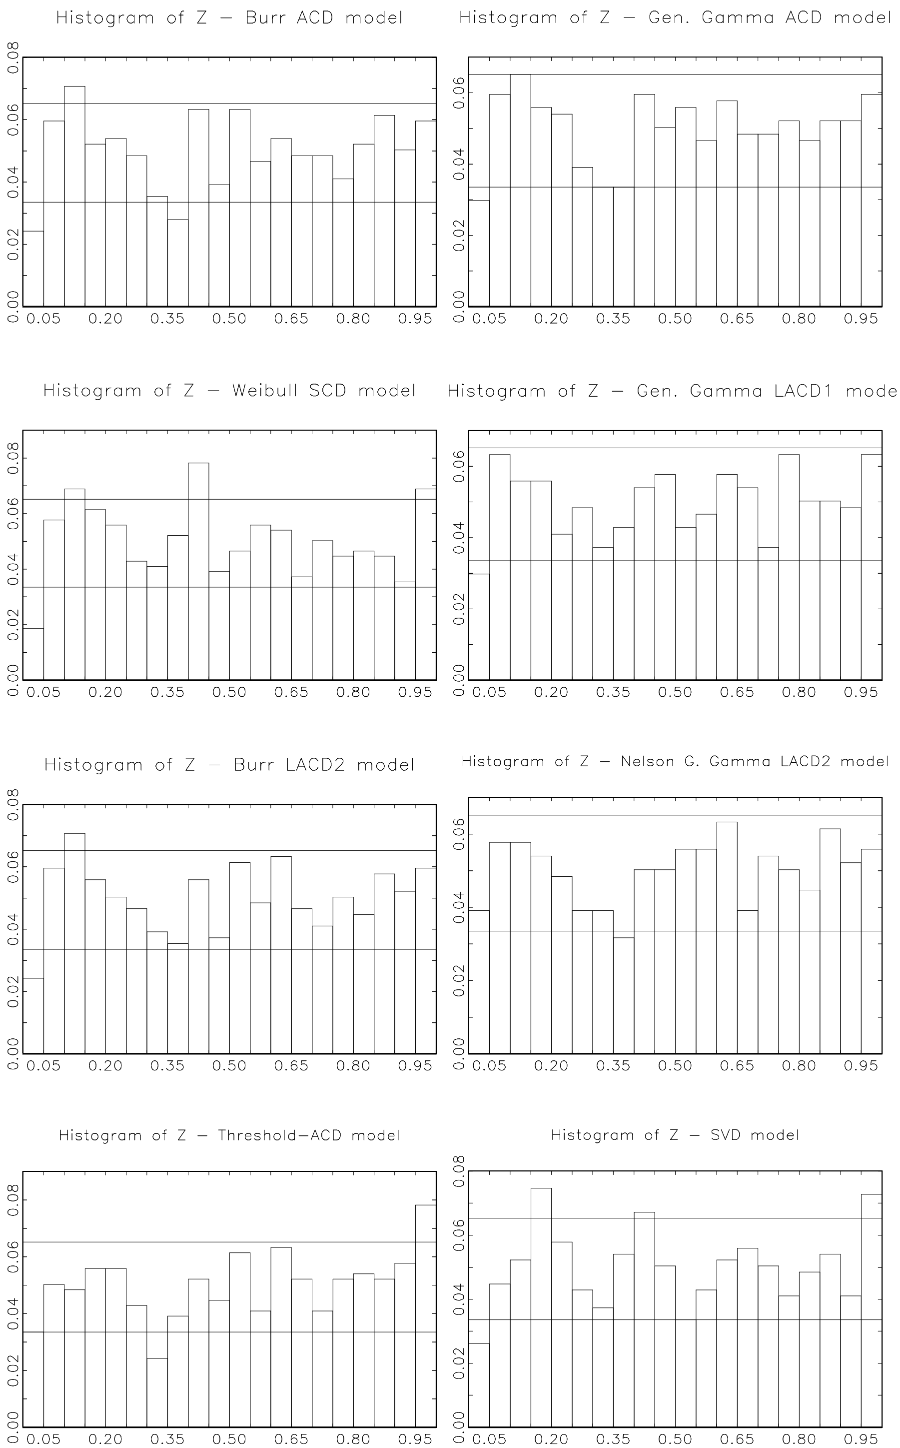

Figure 11: GOF for Coca-Cola Price Durations (out-of-sample)

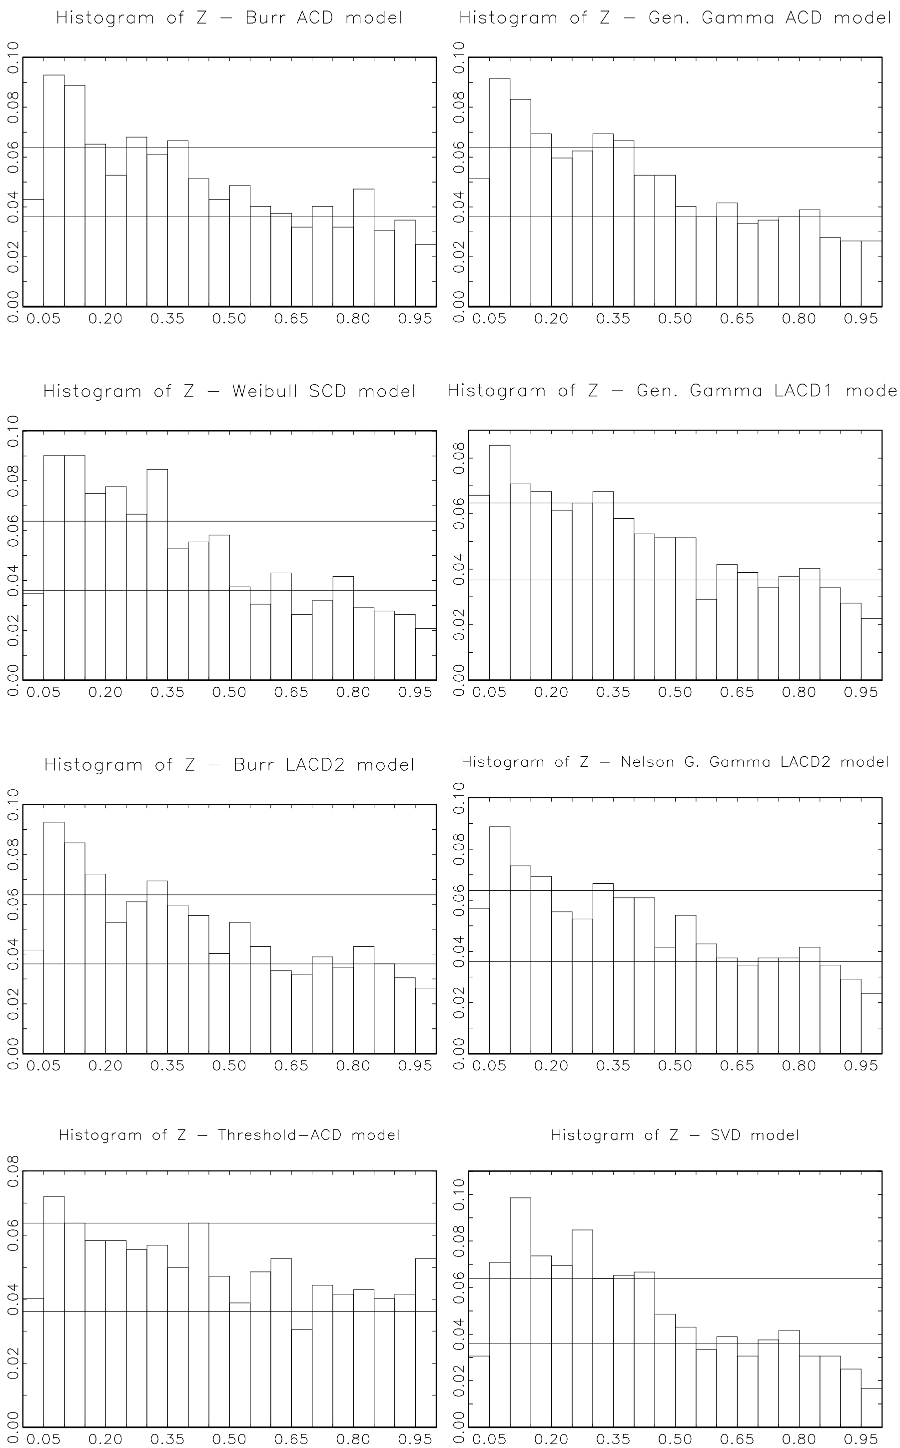

Figure 12: GOF for Disney Price Durations (out-of-sample)

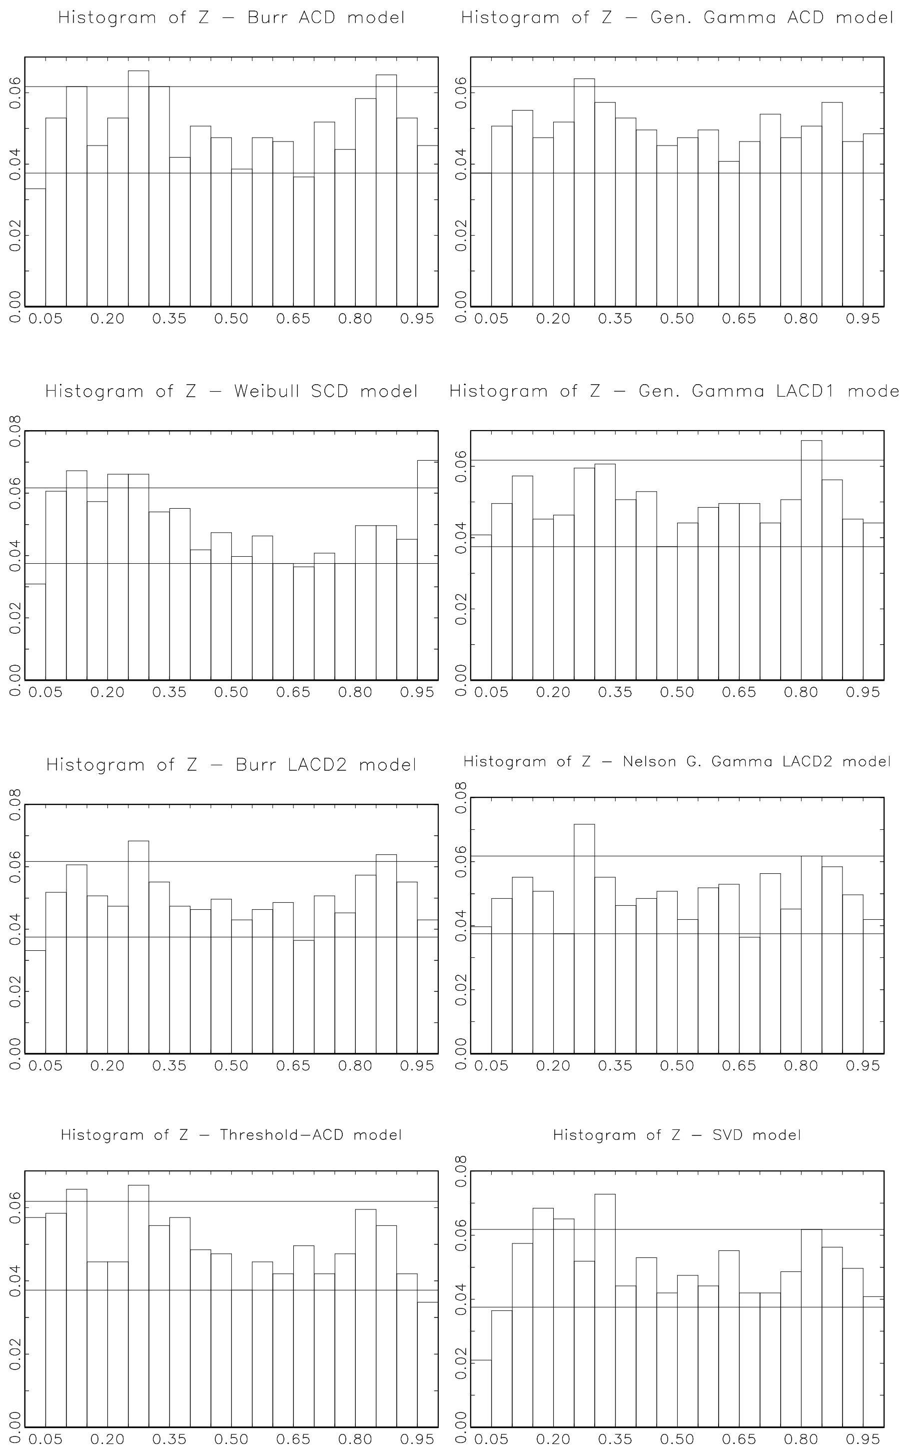

Figure 13: GOF for Exxon Price Durations (out-of-sample)

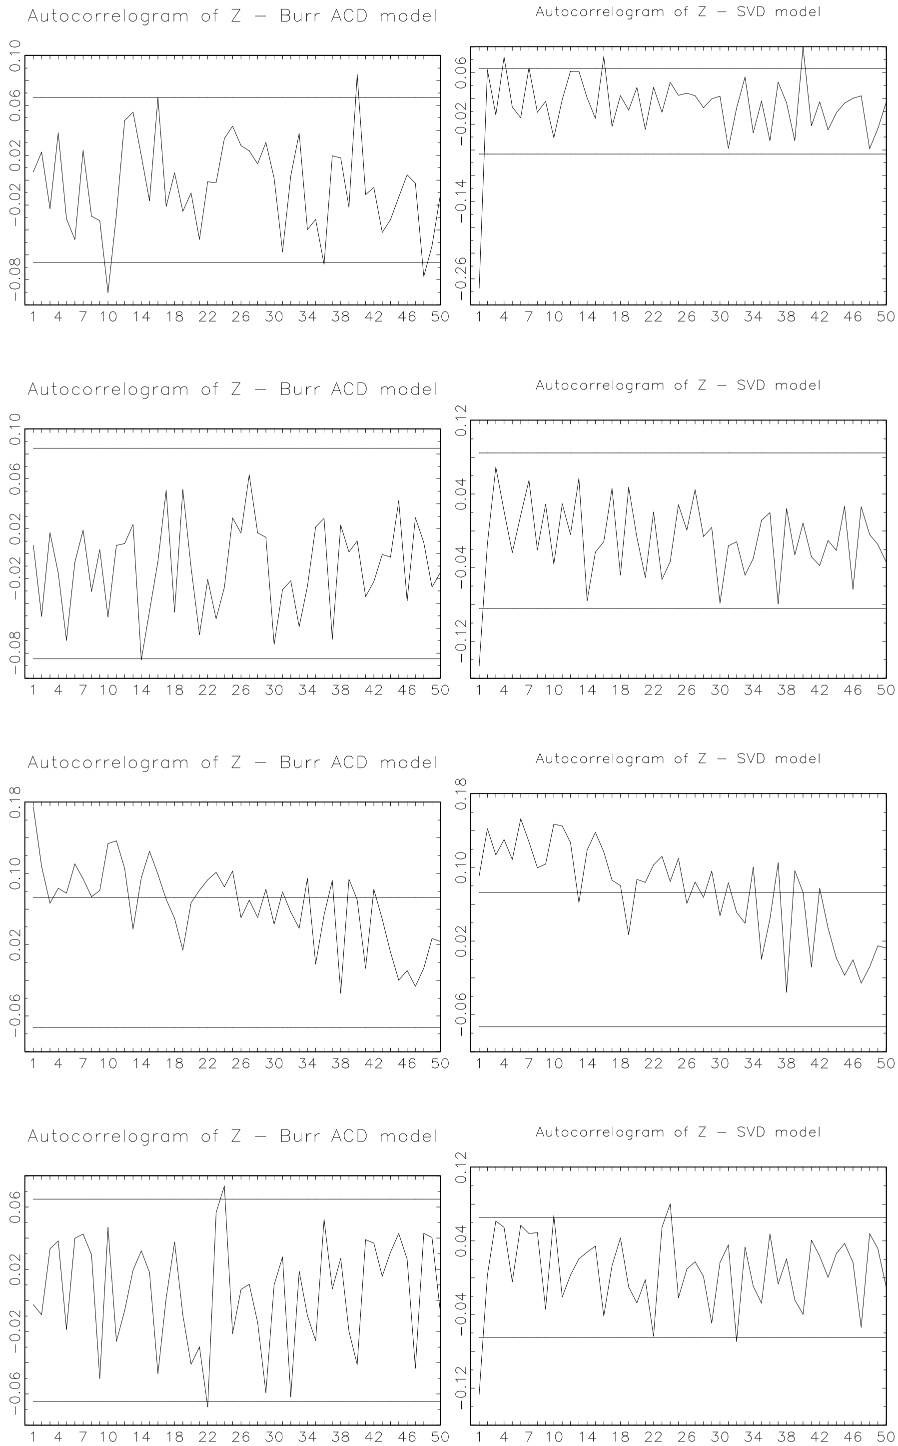

From up to down: Boeing, Coke, Disney, Exxon.

Figure 14: z-Correlograms for Price Durations (out-of-sample)

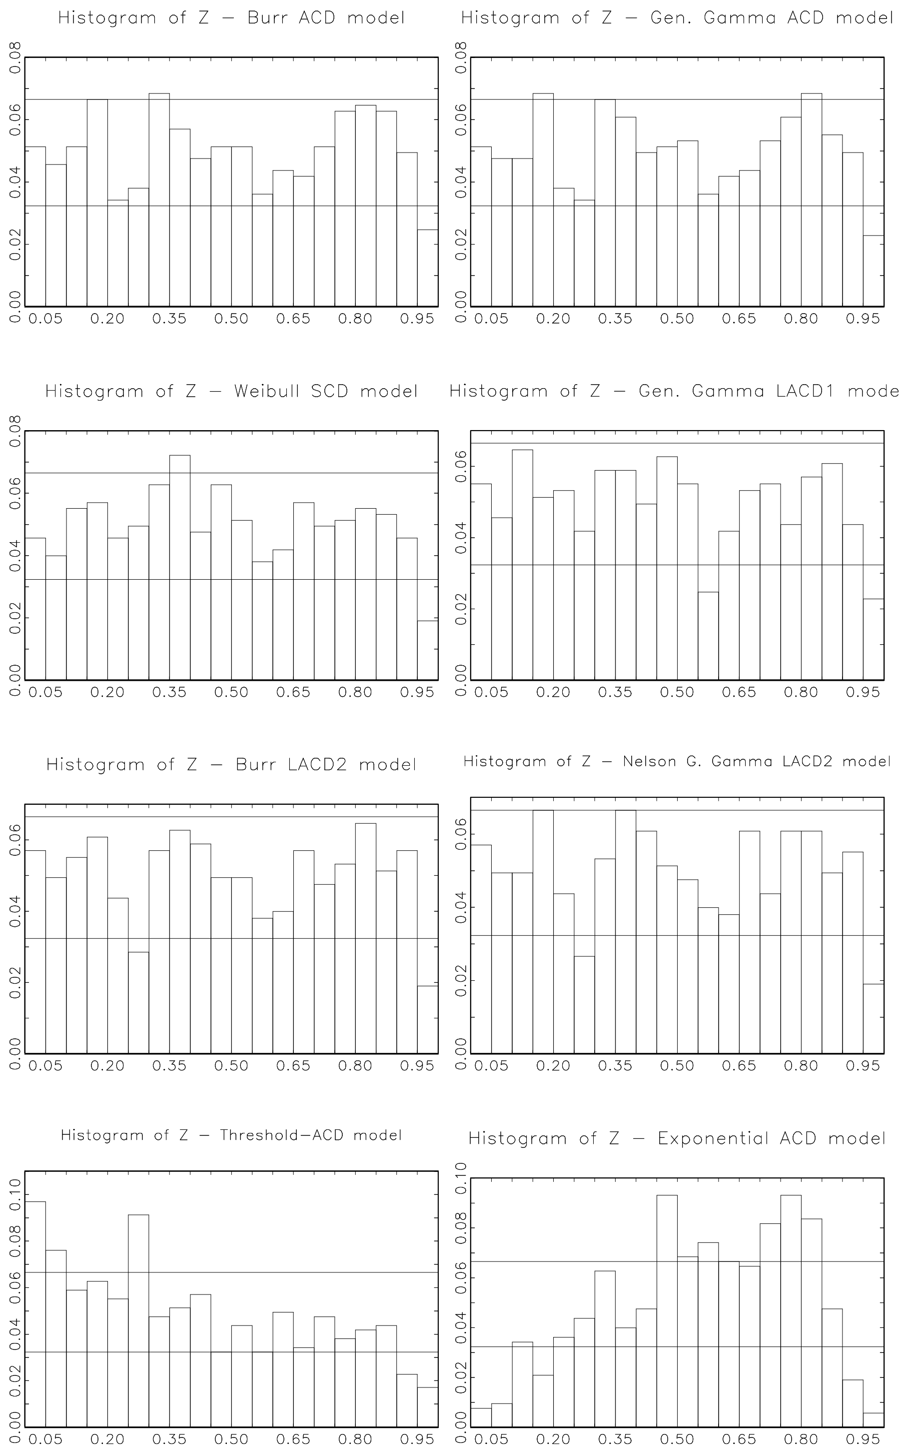

Figure 15: GOF for Boeing Volume Durations (out-of-sample)

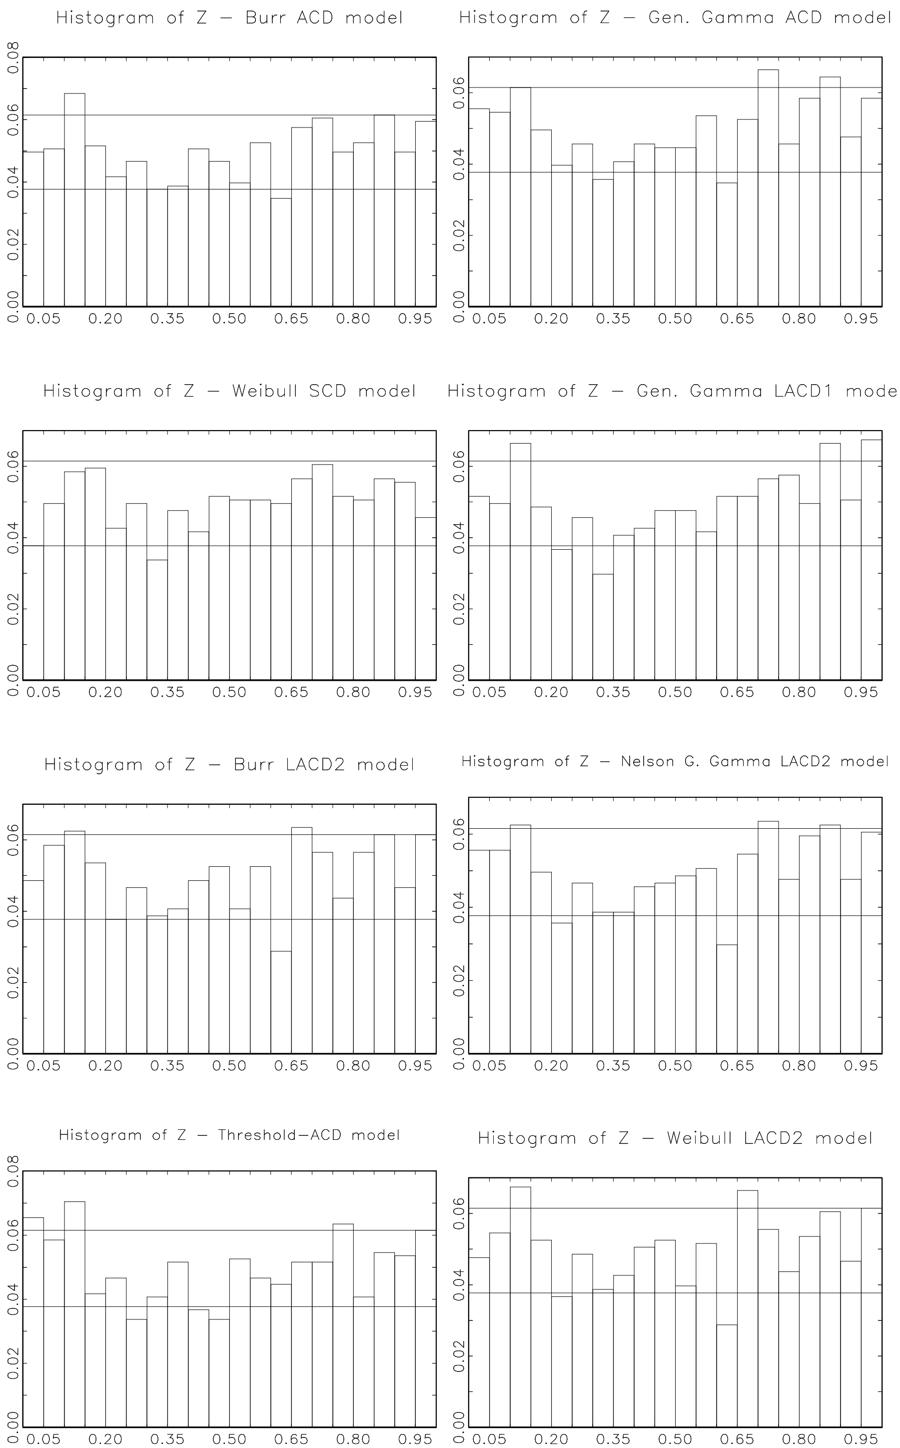

Figure 16: GOF for Coca-Cola Volume Durations (out-of-sample)

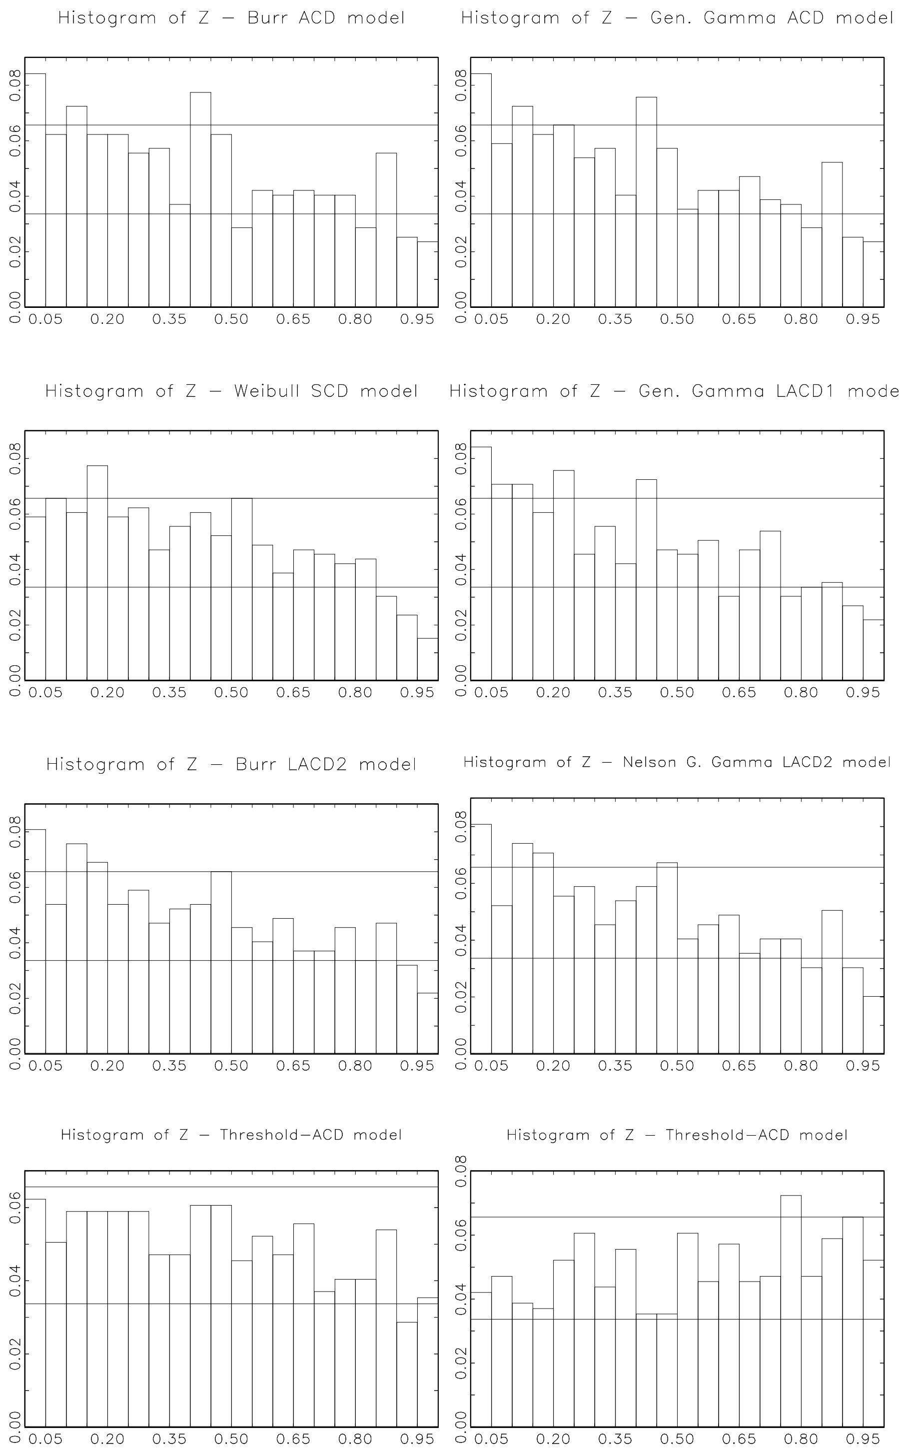

Figure 17: GOF for Disney Volume Durations (out-of-sample except bottom right panel)

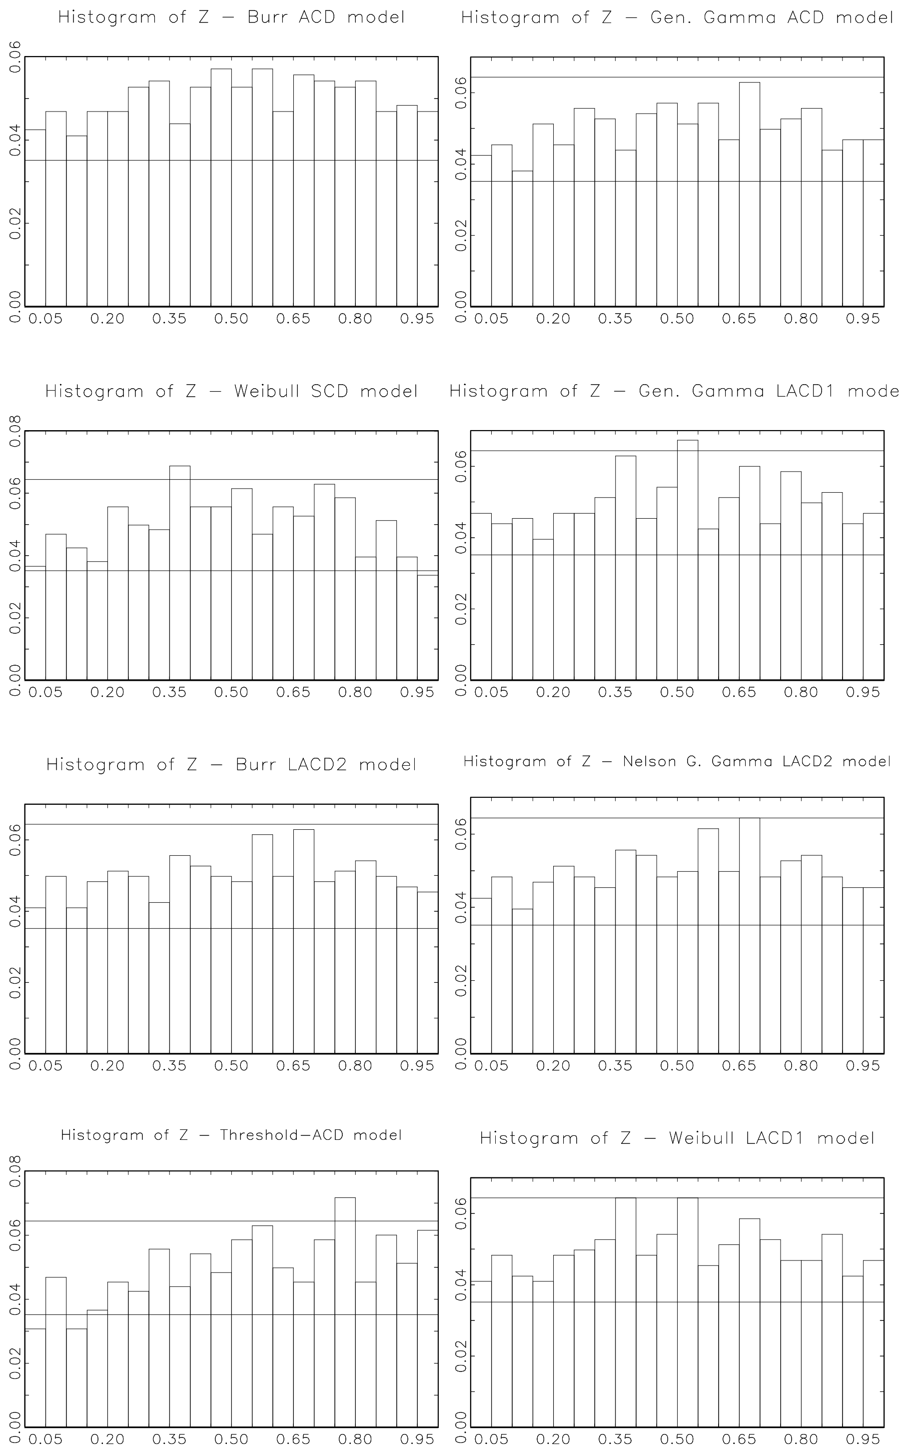

Figure 18: GOF for Exxon Volume Durations (out-of-sample)

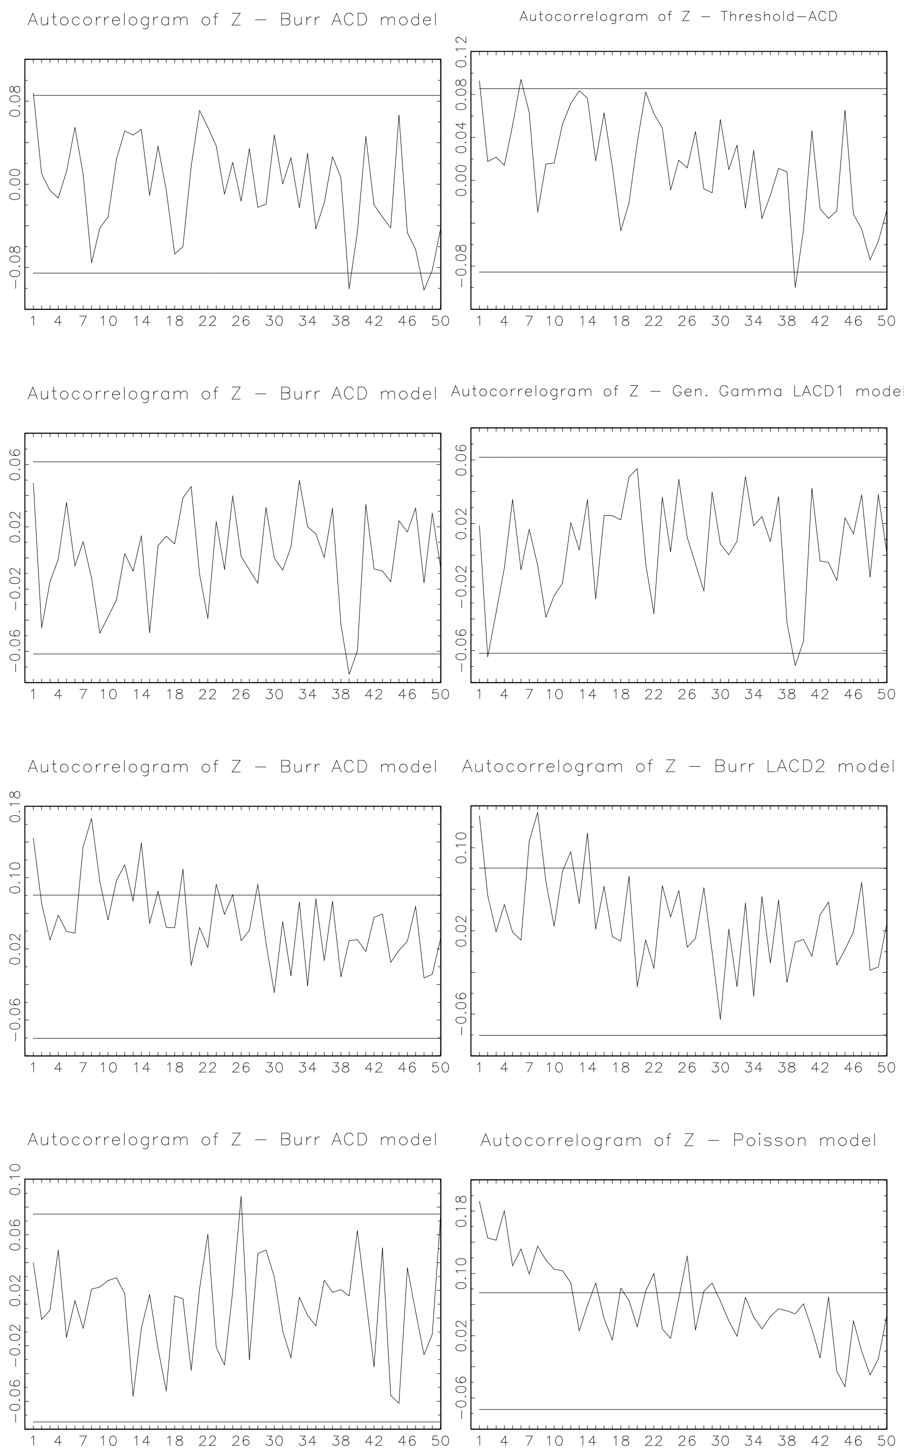

From up to down: Boeing, Coke, Disney, Exxon.

Figure 19: z-Correlograms for Volume Durations (out-of-sample)