Dynamics of Trade and Price Durations of Currency Futures

Serguei Zernov*
Department of Economics, McGill University

January 12, 2003

Abstract

In this paper we study dynamic properties of trade and price durations of currency futures (Japanese yen/US dollar) traded on the Chicago Mercantile Exchange using a model consisting of two parts: a non-stochastic seasonal component and a stochastic component that follows the Stochastic Conditional Duration (SCD) model. We assume that the multivariate seasonal component has a multiplicative (additive in logarithms) form and estimate it using a backfitting algorithm and kernel regression. The SCD model is estimated using the QML approach with the Kalman filter for computing the quasi-likelihood. We investigate implications of the information structure of the SCD model for the problem of estimation and identification of the latent process. We find evidence that trade and price durations may have long-memory properties. In the appendix, we propose a Fractionally-Integrated SCD (FISCD) model that can account for long-memory properties of the data. We estimate the FISCD model for trade and price durations of currency futures using the QML approach in spectral domain.

****

First version: Sep.01, 2002.

Preliminary and incomplete. Please, do not distribute, do not cite without an explicit permission of the author.

1 Introduction

Recent technological progress has made the computing facilities necessary to process high-frequency financial data easily available, as companies providing financial services and research have eased their grip on high-frequency trading records. There has been an academic response to these new opportunities

*The author thanks David Veridas for very inspirational discussions and John Galbraith for his attention and important advices regarding this work. The author also acknowledges the financial support of the CIRANO.

including a number of publications that make use of high-frequency financial data[1].

Trading data (trade records, prices, trade volumes) can be viewed as observable manifestations of the economic process that involves interaction of the market participants given the institutional structure and the flow of economic and financial information. These data are the empirical basis of modeling the economic process and its components. Financial models where asset prices are described by a stochastic process have become a de facto standard in the science of finance. Models featuring stochastic volatility, jumps and other refinements reflect how we see the reaction of the market participants to the heterogeneous flow of economic information to which they are exposed.

Price in economics is realized only through transactions. Thus, if we consider a dynamic asset pricing model (continuous-time or discrete time, or mixed), price variables of such a model are not usually directly observable, in contrast with the textbook wisdom, but manifest themselves through transactions. Whatever inference we desire to obtain about the constituent parts of financial markets, it will be based on the transactions data plus some exogenous data series.

Fitting a dynamic model to transactions data may serve two purposes: first, the dynamic model may give insights into operation of financial markets, and second, it may be used as a black-box for the purposes of forecasting. Transactions data are usually modeled within the framework of so-called *point processes.* A (temporal) point process is a sequence of events (points) on the time line, with each point representing a random arrival time, and a set of associated random variables called *marks* (in our case it may be, for example, the price or the volume of a transaction). A point process with marks can serve as a generator for associated thinned processes, the arrival times of which are a sub-sequence of the arrivals of the original process. The criteria to select these sub-sequences are functions of the arrival times and of the marks of the generator.

Several models have been proposed recently to describe dynamics of arrival times (these models are often termed in the literature *models of durations).* Engle and Russell (1998) were among the pioneers in this area of research and proposed the autoregressive conditional durations model (ACD) which has been amended in several ways by other researchers. Parametric assumptions about the distribution of innovations different from those of Engle and Russell (1998) were considered (Grammig and Maurer 1999); Jasiak (1998) suggested the FI-ACD model that allowed for long-memory properties in the dynamics of durations; Bauwens and Giot (2000) suggested a logarithmic ACD model (the logarithmic specification avoids positivity restrictions on the parameters necessary in the original ACD model). Ghysels, Gourieroux, and Jasiak (1997) noticed that modeling only the conditional mean of the durations might be not sufficient to capture the properties of empirical series. These authors intro-

[1]The financial industry has been studying high-frequency data long before the recent surge of interest in the academia. The industry had the resources and the data. Results of this research were usually available only in-house, and also the research in the industry has somewhat different objectives and methods from the academic research.

duced the stochastic volatility duration (SVD) model in which the conditional second moment of durations also followed a stochastic process. Bauwens and Veredas (1999) suggested a somewhat different approach to modelling dynamics of durations. In their model called the stochastic conditional durations (SCD) model the conditional expectations of durations are presented as e^{ψ_i} where $\{\psi_i\}$ follow an $AR(1)$ process with normal i.i.d. innovations. The relation between the ACD class of models and the SCD models is similar to that between the GARCH models and the SV models.

Approaches have also been suggested that look jointly at the dynamics of trades and of the marks of the process; Ghysels and Jasiak (1998) is one example These later models are important because they are a step towards connecting the econometrics of financial point processes and the theories of dynamic asset pricing.

The main subject of interest of this paper is the empirical study of trading dynamics (the dynamics of trade and price durations) of currency futures. We use ten years of transactions data on Japanese yen (JPY)/US dollar (USD) futures traded at the Chicago Mercantile Exchange (CME). Following the approach adopted by other researchers we represent the durations process as consisting of two components: the non-stochastic seasonal component and the stochastic component[2]. The assets that has been most often analyzed in empirical studies of durations are stocks. The seasonal component of futures, compared to that of the stocks, has a dimension related to the life-cycle of the contract; we model and describe the seasonal behavior of JPY/USD futures with respect to their life cycle. We model the stochastic part of dynamics of currency futures using the SCD model of Bauwens and Veredas (1999), which we estimate using the QML approach and Kalman filter.

Standard specification diagnostics of the model with SCD dynamics estimated on the futures data (we test how well the model fits the dependency properties of the data and we test parametric specifications about the distributions of the innovations) show that the SCD model with an $AR(1)$ latent process does not capture very well the dependency properties of the data. The analysis of the joint information structure of the model and the estimation algorithm suggests that increasing the order of the of the latent process is not a feasible alternative, at least with the QML approach to estimation. We suggest in the appendix an extension to the SCD model that allows for long memory in the latent process. We call this extended model Fractionally Integrated SCD, $FISCD(p, x, q)$[3] (the FISCD(p,d,q) model is mathematically equivalent to the LMSV model of Breidt, Crato, and de Lima (1998)). Coming back to basics, SCD FISCD and LMSV are all essentially models of a discrete signal measure with white, possibly non-Gaussian, noise. Using the spectral QML approach we estimate the $FISCD(1, x, 0)$ model for trade and price durations of currency futures. Breidt, Crato, and de Lima (1998) have proved strong consistency of such estimates but other properties of these estimates are not known and will

[2]In models developed in Engle and Russell (1998) or in Veredas, Rodriguez-Poo, and Espasa (2002), all the components of the dynamics of durations are estimated simultaneously.

[3]The latent process of the FISCD follows ARFIMA(p,x,q).

be a subject of future research.

The data set we analyze in this paper spans a much longer period of time than the data used in previous publications on a similar subject. Thanks to this fact we are able to look at the evolution of estimated model parameters over a period of years. We are also able to investigate how the parameters of the SCD model differ over the periods of the life-cycle of the futures contracts. The analysis of such behavior allows to draw conclusions about how well models like SCD or FISCD capture invariant dynamic properties of the trading process.

We also investigate the asymptotic properties of the QML estimates of the SCD model in the case when the durations process is not considered to be seamless but is initialized in the beginning of every trading day.

Our exposition proceeds as follows. Part two describes the data and the transformations that have been applied to the data. Part three formulates the econometric model. Part four describes the estimation and specification diagnostic methods. Part five presents estimation results, their interpretation and discussion. Part six concludes. Some technical details as well as notes on specification and estimation of the FISCD model are brought to the Appendix.

2 Description of the data

We examine the dynamics of currency futures traded on the CME. Specifically, we study JPY/USD futures. CME currency futures contracts follow the usual March-June-September-December cycle. New contracts are listed the day after the front month expires; the contracts expire on the second business day before the third Wednesday (in our data set, it is always Monday). For example, the first trading date of the March 2003 futures contract (the tick symbol - JYH3) was the 18th of September, 2001 and the last trading date is the 17th of March, 2001. Trading opens at 7:20 and ends at 14:00 Central time.

The data set we analyze spans the period from January 2, 1991 to August 31, 2001 and consists of almost 4 million records. The records of trades of 47 contracts are present in the set, from the contract expiring in March of 1991 to that expiring in January of 2002. For the purposes of our analysis the data had to be filtered. The records in the data set represent either transactions, bid quotes or ask quotes. We removed records that are marked as ask or bid quotes and do not represent actual transactions. As well, we do not consider contracts that have fewer than 130 trading days within the time span of the data set and we excluded contracts that expire after August 31, 2001. This leaves 37 contracts in the set , from the contract expiring in June of 1991 to the June, 2001 contract. The data after filtering consist of 2,743,740 records.

Trade durations are defined as time intervals between consecutive trades; the last duration of a day precedes the first duration of the next day in the duration series. One of the problems that we have had to resolve was the treatment of multiple trades that happened within one second (one second is the precision of the time stamp) i.e. when recorded trade durations were equal to zero. There are several possible ways to deal with this problem, some more sophisticated than others. Here we have taken a simple solution to this problem of censored measurements. We assign the value $\frac{1}{k-1}$ to every zero duration, where k is the number of trades happened within the current second, and we increase the subsequent duration by $\frac{1}{k-1}$ so that the sum of all durations is not changed. Another naive approach would be, following Bauwens and Veredas (1999), deleting all null durations. Bauwens and Veredas (1999) motivated the latter approach by arguing that the trades that happened within a very short period of time were likely to be from the same trader who split a large block of shares.

A price duration is defined as the lapse of time that is required to observe a price change not less that a given amount. It is in some sense natural to measure the change of the price in percentage points (or to measure the logarithm of the change of the price): dynamic models in finance are most often formulated with respect to the logarithm of the price. The matter is complicated by the fact that the transaction price is quoted with a given number of significant digits, i.e. that the observed prices take their values on a discrete set. This is yet another illustration that transactions may be seen as a manifestation of the latent economic process; finite accuracy of the reported prices is a property of the "transmission mechanism" - the market. Russel and Engle (1998) develop a model of price durations where the prices are explicitly discrete-valued. We have chosen to use a change in the logarithm of price as the criterion for thinning. The empirical results presented here are for the case when the change in the logarithm of the price is equal or larger than 0.05%[4].

3 Modelling trade and price durations of currency futures

Let $\{D_i\}$ denote the recorded durations (trade durations or price durations). In what follows, when there is no ambiguity, we shall use small letters to denote logarithms of the values denoted by the corresponding capital letters i.e. $d_i \equiv \ln(D_i), .\psi_i \equiv \ln(\Psi_i)$ etc. The model that we are estimating is formally specified as follows:

$$D_i = \Phi(\boldsymbol{\kappa}_i)\,\Psi_i\varepsilon_i \qquad (1)$$

We assume that $\varepsilon_i|I_{i-1} \sim iid\ \mathrm{D}(\boldsymbol{\eta})$ where I_{i-1} denotes the information set at the beginning of the spell of the duration d_i and $\mathrm{D}(\boldsymbol{\eta})$ is a distribution with a positive support with a parameter $\boldsymbol{\eta}$. The fourth moment of $\mathrm{D}(\boldsymbol{\eta})$ exists and is finite. Usual choices of the parametric form of the distribution $\mathrm{D}(\boldsymbol{\eta})$ in the context of duration studies are the Weibull distribution and the standard Gamma distribution. The process $\psi_i = \ln\Psi_i$ follows, in the general case, a stationary $ARMA(p,q)$ with Gaussian innovations.

[4]Stocks prices are usually recorded with lower accuracy that the currency futures, which is why accounting for discretization error is more important in the former case.

The function $\Phi(\boldsymbol{\kappa})$ is assumed to be non-stochastic and strictly positive for all admissable values of $\boldsymbol{\theta}$. Taking the logarithm of the equation above,

$$d_i = \phi(\boldsymbol{\kappa}_i) + \mu(\boldsymbol{\eta}) + \psi_i + \xi_i \tag{2}$$

where $(\xi_i + \mu(\boldsymbol{\eta}))$ is distributed as logarithm of ε_i and $\mathbf{E}[\xi_i|I_{i-1}] = 0$. Under the specifications above log-durations are sums of the non-stochastic part $\phi(\boldsymbol{\kappa}_\iota) + \mu(\boldsymbol{\eta})$ and the stochastic part $\psi_i + \xi$. If we define $\hat{d}_i = d_i - \phi(\boldsymbol{\kappa}_i)$ and assume that ψ_i follows $AR(1)$, the model in terms of $\hat{d}_i$ will be the SCD model as it has been formulated and studied in (Bauwens and Veredas 1999).

We shall argue later in this study that it is not practical to consider SCD models with the latent process of order higher than $AR(1)$. FISCD model introduced in Appendix (7.3) is a more flexible alternative than SCD. FISCD is a complex econometric object; many properties of the parameter estimates under the FISCD are unknown and are a subject of future research.

3.1 Seasonality in the dynamics of durations

The literature presents strong empirical evidence for seasonality of trade and price durations (for example, in Engle and Russell (1998), Gouriéroux, Jasiak, and Fol (1999) and in Bauwens and Veredas (1999)), which is why the seasonal component $\Phi(\boldsymbol{\kappa})$ is present in Equation 1. Unlike stocks that may be thought of as having an infinite time horizon, assets like derivative contracts and bonds have a life cycle, from contract inception to its expiration. This life cycle is reflected in the "seasonal" behavior of time series describing dynamics of such contracts. This form of seasonal behavior of trade and price durations of futures, due to their life cycle, has been give less attention in the empirical literature that the diurnal or weekly seasonality[5].

Under the model adopted in this study, duration series have two components: the deterministic seasonal component and the stochastic component that follows the SCD dynamics. We can approach the estimation of the seasonal component parametrically, semi-parametrically or non-parametrically, and there exist several possibilities in each of these classes of estimation techniques. An attractive feature of non-parametric modelling is its flexibility. We shall model the seasonal dynamics of the durations in the non-parametric spirit (strictly speaking, the approach that we use is semi-parametric, as the reader will see from the exposition below), similar to the approach adopted in .Veredas, Rodriguez-Poo, and Espasa (2002) and in Bauwens and Veredas (1999) with the difference that the the futures considered in this paper have a more complex seasonal structure.

The multiplicative presentation (1) and the additive presentation in logarithms (2) are equivalent at the stage of modelling. When it comes to estimation, whether one applies seasonal adjustment before taking the logs or after leads to different results. In Bauwens and Veredas (1999) the data are seasonally adjusted before taking the logarithms. The advantage of this approach is

[5] Gerhard and Hautsch (2002) desrcibe the seasonality over the maturity of intra-day volatility for BUND futures. They estimate intra-day volatility based on price durations.

that the results are easier to interpret and easier to apply to forecasting (we are interested ultimately in durations and not in the logarithms of durations). We have chosen however to apply the seasonal adjustment after taking the logarithms of the data, for two reasons. First, we shall assume later in this study that $\phi(\boldsymbol{\kappa})$ follows the additive model with the logarithm as the link function[6]. Properties and estimation of additive models are known better than properties of GAM, and we would like to build upon this knowledge. Second, the SCD is a model with dynamics linear in the logarithm of durations; estimation using the Kalman filter is based on this linearity property. Seasonal adjustment of the dynamic variable of the model seems to be more transparent than adjusting the non-linear transformation of this variable.

As we have mentioned just above, we impose additional structure on the seasonal component of the durations. Specifically, we assume that

$$\ln \Phi(\boldsymbol{\kappa}) \equiv \phi(\boldsymbol{\kappa}) = A_\delta + \chi(t) + \zeta(\tau), \tag{3}$$

where $\boldsymbol{\kappa} = \{\delta, t, \tau\}$, A_δ ($\delta \in \{Monday, ..., Friday\}$) describes the weekly seasonality, $\chi(t)$ - the seasonality due to contract life cycle, t is the time to expiration, and $\zeta(\tau)$ corresponds to the diurnal seasonal component, τ is the time elapsed from the beginning of the trading session. We test later in the paper the assumption that the weekly, the life-cycle and the diurnal components are orthogonal. The additive form of ϕ reduces the dimensionality of the non-parametric regression problem. The data set that we analyze is large. Preliminary analysis shows however that even with this size of the data set the curse of dimensionality cannot be escaped, especially when we estimate the seasonal component at longer horizons to expiration t where the trading is sparse. Under the assumption of the additive model one can achieve, under certain conditions, the same rates of convergence of non-parametric regression as in the univariate case (Linton and Nielsen 1995).

4 Estimation Methods

4.1 Estimation of the seasonal component

We model the seasonal deterministic component of log durations in the following manner:

$$\mathbf{E}[d_i | \delta_i, t_i, \tau_i] = \alpha + A_{\delta_i} + \chi(t_i) + \zeta(\tau_i)$$

[6] $\Phi(\boldsymbol{\theta})$ follows in this case the Generalized Additive Model (GAM). It is said that a non-parametric regression function follows the GAM if

$$f(m(\mathbf{X})) = \alpha + \sum_{j=1}^{d} \chi_j(X_j)$$

where $f(\cdot)$ is a known link function, $\chi_1, ..., \chi_j$ are unknown univariate functions and $\mathbf{X} = (X_1, ..., X_j)$.

One of the approaches to estimating additive models is the so-called *backfitting* algorithm Fan and Gijbels (1996), pp.266-267). In application to our problem, the algorithm works as follows:

1. Initialization: $\alpha = \frac{1}{N}\sum_{i=1}^{N} d_i$. We subtract the sample mean from $\{d_i\}$ and make initial guesses about all but one seasonal components (about A_δ and $\chi(t)$, for example) We force the sample expectation of each of these components to be equal to zero, so that the mean of the seasonally adjusted data is equal to zero.

2. $d_{\tau,t,v}^{(k+1)} = d^{(k)} - A_\delta^{(l)} - \chi^{(l)}$, where $A_\delta^{(l)}$, $\chi^{(l)}$ are the latest estimates of A and χ $d^{(k+1)}$ is used to obtain an updated estimate of ζ via a univariate non-parametric smoother.

3. We repeat step 2 for each of the seasonal components until convergence is achieved. (2 or 3 rounds are usually sufficient in practice).

The term *backfitting* referring to the action above was first used in Friedman and Stuetzle (1981). If the additive specification is not the true model, the algorithm is expected to give the estimates that are the best additive approximation to the regression surface (Breiman and Friedman 1985).

We have to specify a univariate smoother that will be used at the step 2 of the algorithm above. Our current choice is the kernel regression (other choices like smoothing splines, for example, will probably work equally well). The form of kernel regression is otherwise known as the Nadaraya-Watson estimator and it is defined as

$$f(x) = \frac{\sum_{i=1}^{N} K\left(\frac{x-x_i}{h}\right) d_i}{\sum_{i=1}^{N} K\left(\frac{x-x_i}{h}\right)}$$

We have chosen the quartic kernel for our regression because it is fast to compute and it has a compact support which helps to reduce the complexity of the computations. For the results reported in the paper the values of the bandwidth parameters were chosen by their visual performance.

The component A_δ is somewhat different from χ and ζ. We estimate it as the mean duration for a given day of a week (formally, this is equivalent to setting $h = 1$, if we want to preserve the uniformity of exposition). Taking into account that we force the sample mean of each component of $\phi(\boldsymbol{\kappa})$ to zero, A_δ has four degrees of freedom, four parameters to be estimated, that is why we have mentioned that our approach can be called "semi-parametric".

The argument of $\chi(t)$, the life-cycle seasonal component, is the time to expiration in business days, $t \in \mathrm{N}$. We consider in our study the records with $1 \le t \le 130$. The argument of the diurnal component $\zeta(\tau)$ represents the time from the beginning of the trading day in seconds, $\tau \in [0, 24000)$.

We should point out that the whole data set is used to estimate the seasonal component of durations $\phi(\boldsymbol{\kappa})$. This will allow us to capture invariant properties of the seasonal components across the years spanned by the data set. We estimate SCD parameters individually for each contract in the date set. Part of the variation of trading intensity will be also accounted for by the parameter ω of the SCD model that we shall keep free.

4.2 QML estimation of the SCD model

After the seasonal component of the logarithm of durations has been estimated, we use adjusted series to estimate the SCD model. The sample mean of seasonally adjusted log durations and the average value of each of the adjustment factors are equal to zero over the whole sample of thirty seven contracts, but not for each individual contract. To account for this we allow the conditional duration to have a non-zero mean and we subtract from each of the seasonal components their average values over that specific sample. We assume that the seasonally adjusted log durations follow the model

$$\begin{aligned} d_i &= \mu(\gamma) + \psi_i + \xi_i \\ \psi_i &= \omega + \beta\psi_{i-1} + u_i,\ |\beta| < 1 \end{aligned}$$

where $\{\xi_i + \mu(\gamma)\}$ follows a log-Weibull distribution with parameter γ, $\mathbf{E}[\xi_i] = 0$, and we use QML and Kalman filter to estimate the parameters of the model. Note that

$$\mathbf{E}\left[\frac{1}{N}\sum_{i=1}^{N} d_i\right] = \mathbf{E}[d_i] = \mu(\gamma) + \frac{\omega}{1-\beta}.$$

Making the change of variable $\psi^* = \psi - \frac{\omega}{1-\beta}$, the model can be written as:

$$\begin{aligned} d_i^* &= \psi_i^* + \xi_i \\ \psi_i^* &= \beta\psi_{i-1}^* + u_i, \end{aligned} \tag{4}$$

where ω has been concentrated out of the likelihood, as has been suggested in Ruiz (1994) for estimating stochastic volatility models. The vector of parameters of the quasi-likelihood to be minimized is $\boldsymbol{\theta} = \left\{\sigma_u^2, \gamma, \beta\right\}$. The asymptotic theory for the QML estimate of $\boldsymbol{\theta}$ was developed in Dunsmuir (1979) yielding $T^{\frac{1}{2}}\left(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\right) \overset{d}{\sim} \mathrm{N}\left(0, C(\boldsymbol{\theta})\right)$. The expression for $C(\boldsymbol{\theta})$ and details of computation are given in the appendix.

The estimation results presented in this paper were obtained for the case when the data are treated as continuous series, i.e. the end of one day precedes directly the beginning of the next except that we have cut out the first 20 minutes in the beginning of each day. We use a Kalman filter initialized with the diffuse prior to compute the quasi likelihood function of the model (4).

4.3 Diagnostic methods

As far as it concerns the nonstochastic part of the model, our primary concern is whether the seasonal components of the additive model are close to being independent or not. We use graphical methods to investigate this: we estimate the diurnal component for each day of the week and draw them on the same graph; in the same manner we draw the life cycle seasonal component for different days of the week and for different periods of the day. If the is no interdependence between seasonal variables, the lines on each of these graphs will not be far apart one from another.

To assess how the SCD model describes the dynamics of seasonally adjusted durations we want to investigate two issues: first, how well the model accounts for the dependency properties of the process; second, how good are the parametric assumptions about the distributions of the innovations. Since the SCD model is a model with one latent variable, we can compute two series of residuals resulting from the model estimation:

$$\hat{\xi}_i = d_i^* - \hat{\psi}_i^* \text{ and}$$
$$\hat{u}_i = \hat{\psi}_i^* - \beta\hat{\psi}_{i-1}^*.$$

The serial dependence structures of the series $\left\{\hat{\xi}_i\right\}$ and $\{\hat{u}_i\}$ are similar one to another: their ACF are different only by a scalar under the model. Under the correct specification $\left\{\hat{\xi}_i + \mu(\gamma)\right\}$ has a log-Weibull distribution and $\{\hat{u}_i\}$ is distributed log-normally.

We use the traditional ACF and PACF analysis as well as the Spearman's coefficient to investigate the serial dependence in the residuals. Bauwens and Veredas (1999) argue that Ljung-Box statistic (or similar statistics based on the sample autocorrelations) would not be a correct measure of dependence in the context of irregularly spaced data which is why the Spearman's coefficient should be preferred. This argument has its merits. However, within a framework of a discrete time model, Ljung-Box statistic will detect the presence of linear dependence in the series, regardless whether the measurements are taken at equal physical time intervals or irregular.

Under the model, the residuals $\left\{\hat{\xi}_i\right\}$ are distributed as log-Weibull with parameters $(\gamma, 1)$, and $\{u_i\}$ have a standard normal distribution. In order to judge the compatibility of our parametric assumptions with the empirical observations we shall use p-value plots and p-value discrepancy plots as well as commonly used non-parametric goodness-of-fit tests such as the Anderson-Darling test and the Cramèr-von Mises test[7].

5 Estimation results: interpretation and discussion

5.1 Seasonality in trade and price durations

Figure 1 presents graphs of the seasonal components A_δ, χ and ζ of trade durations (left column) and of price durations (right column). We observe that

[7]These statistics, their properties and their critical values for normal and log-Weibull distributions can be found in (Stephens 1976) and (Stephens 1977).

over a week the trading is the least active on Mondays and then its intensity increases towards Friday. This is in line with the results reported in the financial literature. Over a day, the trading activity is high in the morning, then it is decreases gradually to its lowest level at some moment between 12:30CT and 13:00CT, and increases again near the end of the trading day. The observed diurnal pattern is similar to those described for stocks by Bauwens and Veredas (1999) or by Gouriéroux, Jasiak, and Fol (1999). The lowest levels of daily activity for the stocks studied, Boeing in the former case and Alcatel in the latter, has been observed between 13:30ET and 14:00ET (Central time is equal to Eastern time minus one). The conclusion that we can draw is that the there is a synchronicity between the life of the market in Chicago and the market in New York. The increase in the level of trading activity at the end of a day is more pronounced for the CME currency futures than for the stocks studied in the articles just mentioned. An explanation for this may be that the trading at the CME closes at 14:00CT, just one hour after the lunch break responsible presumably for the trough in trading activity, while the NYSE trades until 16:00ET. Therefore, the traders of CME currency futures have less time after the lunch break to take their end-of-the-day positions than those trading the stocks on the NYSE[8].

The seasonal pattern due to the contract life cycle is in accordance with our expectations. Trading is the most active for the contract closest to expiration, the level of trading activity is relatively flat from 65 business days to about 5 business days to expiration. The closest to expiration contract is traded less actively in the last few days of its existence because the traders switch to the next contract. When the nearest to expiration contract has six business days to expiration, the second to expiration contract has usually (not always because of holidays) sixty nine business days to expiration. The change in the life-cycle seasonal component of durations between 70 and 60 days to expiration correspond to a more that tenfold change in the expected trade duration and to a change of about four times in the expected price duration. Traders take positions in the contract that is going to become the nearest to expiration. We do not study the trading dynamics beyond 130 business days to expiration because the trading there is very sparse, typically just a few trades per day, if any.

5.2 Estimated parameters of the SCD model

Let us discuss now the results of estimation of the SCD log-Weibull model. Figure 2 shows estimated values of the parameters β, σ^2, γ for each of the contract in the calendar order with two asymptotic standard deviation error bars. The left column shows trade duration parameters, the right corner - the price duration parameters.

[8] The strength of this argument is mitigated by the fact that the CME futures are traded on GLOBEX 24 hours a day. The difference in the shapes of diurinal components for stocks and CME currency futures may be also due to institutional differences and/or in the mechanism how the trades are reported.

A prominent feature of estimated parameters of trade durations is that asymptotic standard errors are narrow. Given the sample sizes (the number of observations per contract is in the range from $53,000$ to $135,000$), we believe that the asymptotic standard errors provide a good indication of the true standard errors. This assertion is supported also by the results of Monte-Carlo experiments reported in Bauwens and Veredas (1999), where the authors simulated samples of the size $N = 50,000$. Their simulated standard errors are very close to the asymptotic values. If we allow for the possibility that the model is misspecified and we want to evaluate how far apart are the estimated model and the data-generating process, this distance (in some metric) will likely have two contributing factors: the statistical error of estimation and the error due to the model misspecification. Tight asymptotic standard errors of the estimated model parameters and the fact that the asymptotic standard errors are close to the finite sample standard errors may be interpreted as indicating that the first contributing factor, the model estimation error, is small. The argument just above is not so relevant for the model parameters of price durations because the sample sizes of price durations are smaller, hence, the asymptotic standard errors of parameter estimates are wider.

The estimates of β of the latent process lie between 0.961 and 0.987 for trade durations series and between 0.882 and 0.981 for price durations: i.e. both the price and the traded durations the processes are very persistent. The values of β are significantly less than 1 at conventional levels of significance. For trade durations, the estimated values of the parameter γ of the Weibull distribution are significantly less than one for the contracts before June, 1995 and significantly greater than one for later contracts. The price durations $\hat{\gamma}$ are between 1.1 and 1.33 and are always significantly greater than one. The estimated values of σ_u^2 vary from 0.011 to 0.040 for trade durations and from 0.0057 to 0.019 for price durations. One can observe that the values of the model parameters β, γ and σ_u^2 are significantly different, based on the asymptotic standard errors, across the contracts considered, both for trade durations series and for price durations. Qualitatively, however, we may conclude the behavior of the durations process as described by the SCD model is similar for all the contracts studied: we observe high persistence and low signal-to-noise ratio (SNR)[9], ranging from 1.18 to 1.36 for trade durations and from 1.49 to 2.94 for price durations. As we shall see in a moment, low SNR will have interesting implications from the point of view of model identification.

5.3 SCD parameters and the horizon to expiration

It is interesting to investigate the question of stability of model parameters across different horizons to expiration. There is a natural split in the trading data for each contract: the records when the contract is the nearest to expiration and the records when the contract is the second nearest. If we find out that the model parameters are stable over the horizon to expiration, this will provide

[9]We define SNR here as the ration of the unconditional variance of $\{d_i\}$ to that of $\{\psi_i\}$.

some evidence for the correctness of the seasonal adjustment algorithm used and of the model itself; this shows that the model captures some invariant properties of the dynamics of the data.

Let us first consider trade durations series. The left column of Figure 3 shows the estimated values of the SCD model parameters for contracts with horizon to expiration from ranging from 70 to 130 business days. The values of the corresponding parameters estimated using the whole sample are given on the same graphs as a reference. A futures contract is traded much less actively when it is the second closest to expiration than when it is the closest to expiration, which is why the subsample corresponding to the trades with time to expiration from 70 to 130 days comprises only a small fraction of all trading records of a contract (less than 5% for some of the contracts).

The estimates of β and σ_u^2 for trade durations based on records with 1 to 130 days to expiration and on records with 70 to 130 days to expiration are close. Informal analysis suggests that the estimates of β in the former case are higher than in the latter; the estimates of σ_u^2 are smaller in the former case than in the latter. This difference in the estimated parameter values can be explained, in part at least, by our approach to modelling. We treat the data as continuous series which is merely an approximation. The second to expiration contract, as has been mentioned, is traded much less actively than the closest to expiration; hence, the links between intraday spells of trades constitute a larger proportion of the data in the former case. The conditional distribution of a duration will intuitively depend less on the previous measurement if this measurement has been taken at the end of the previous trading day but we do not account for this in the model. That is why we can expect to observe lower persistence for contracts with longer horizons of expiration within our modelling framework, and also a higher variance of innovations of the latent process.

The estimates of trade durations γ based on records with 70-130 days to expiration are definitely lower than those based on contracts with 1-130 days to expiration.

The SCD parameters of price durations for different horizons to expiration differ less than those of trade durations. We still observe that the estimates of the persistence parameter are lower at longer expiration horizons, and we can use the same rationale as for the trade duration to explain why. The estimates of σ_u^2 and of γ of price durations based on records with longer expiration horizons are not statistically different from those based on all records.

Price durations change less with the expiration horizon than the trade durations. Thus, the proportion of links between trading days in the price durations data does not increase as much, when we consider only trades with 70 to 130 business days to expiration, as in the trade durations data. This may be one of the explanations why the estimated SCD parameters of price durations differ less across expiration horizon than that of trade durations. Another part of the explanation may be that the dynamics of trade durations depend more on the specifics of the trading system than do the dynamics of price durations; the model used in this study accounts for the properties of the trading mechanism only in very general terms.

The dynamics of price durations do depend on the trading mechanism however. Futures price follows very closely (virtually one-to-one) the price of the underlying asset. One could expect that price durations would not change as a function of the horizon to expiration, but we definitely observe life-cycle seasonality in the price durations, albeit weaker than in trade durations. The reason that we observe life-cycle seasonality is that the number of futures contracts in circulations is smaller when the contract is the second to expiration that when it is the closest to expiration. This latter property, the number of contracts in circulation, is more closely related to the transmission mechanism that to the information process determining the dynamics of the "latent" futures price.

5.4 Specification diagnostics

Figure 4 depicts the estimated additive contributions of the life-cycle and diurnal component for each day of a week and the contributions of the life-cycle components for three two-hour periods of a trading day. We can conclude from the analysis of these graphs that the assumption of additivity is not grossly incompatible with the data, either for trade durations or (to a lesser extent) the price durations.

Figures 5 and 6 illustrate how well the SCD model accommodates the dependence properties of the (seasonally adjusted) trade and price durations series correspondingly[10]. The measures of dependence for the seasonally adjusted log durations are in the left columns of figures 5 and 6, the measures of dependence for the model residuals are in the right columns[11]. Both trade and price durations exhibit strong dependence, as has been documented in the financial literature. The SCD model fails to describe fully the dependence properties of the data: the residuals retain a degree of dependence.

Bauwens and Veredas (1999) have also found that the residuals of the SCD model with $AR(1)$ latent process estimated using trade durations of a stock are not independent (the authors used Spearman's ρ statistic as a measure of dependence). They mentioned as a possible explanation, citing Jasiak (1998), that trade durations may be fractionally integrated. Our preliminary analysis also suggests the presence of long memory in the duration series (the estimates of the fractional integration parameter x of the FISCD model given in the appendix are between 0.4 and 0.6).

One may think that a mechanical increase of the order of the latent process will allow the model to better accommodate the dependency properties of durations. This simple approach does not work as well as one might have expected. The reason "Why?" lies in the analysis of the structure of the model. The resolution a system measuring a mixture of signal and noise, i.e. the ability to distinguish between different signals (latent processes), as it is known in the theory of signal processing, depends among other factors on the geometry of

[10]We use the data of one specific contract to illustrated the dependence properties of the data and the model residuals and the goodness of fit of the parametric distributional assumptions. The results are representative of all contracts considered.

[11]We use $\{\hat{u}_i\}$ residuals to computed the ACF, PACF and Spearman coefficient shown.

the space of the solutions and on the signal-to-noise ratio. Given the parameter values typical for our data the SNR is low, especially for the trade duration series. If we allow an increase of the order of the latent process to $AR(2)$, we increase the domain of the possible solutions. Moreover, in the presence of white noise the QML estimation algorithm used loses its resolution abilities primarily at higher frequencies, but this is exactly where the $AR(1)$ and $AR(2)$ differ one from another.

To illustrate the argument above let us compute the inverse of the information matrix (see 5 in the Appendix) of ML estimates of the parameters of a Gaussian $AR(2)$ process measured with Gaussian noise (we ignore the correction for non-normality of the measurement noise for the sake of transparency of the exposition). The values of the parameters used in this example are $\beta_1 = 0.95, \beta_2 = 0.02, \sigma_u^2 = 0.02$ and $\gamma = 1$ (γ is used to compute the variance of the measurement noise, $\sigma_\xi^2 = \frac{\pi^2}{6\gamma^2}$) β_i are the parameters of the $AR(2)$ latent process[12] The model is parametrized as $\boldsymbol{\theta} = \left\{\beta_1, \beta_2, \sigma_u^2, \gamma\right\}$:

$$\mathcal{IF}^{-1} \cong \begin{pmatrix} 5323356.3 & -5166653.3 & -208937.10 & -32404.133 \\ -5166653.3 & 5014563.3 & 202786.57 & 31450.200 \\ -208937.10 & 202786.57 & 8200.6707 & 1271.8867 \\ -32404.133 & 31450.200 & 1271.8867 & 197.83256 \end{pmatrix}$$

We observe that the estimates of β_1 and β_2 have very high variance and the correlation between them is almost -1 (this will be especially true when β_1 is close to one and β_2 is relatively small). In practical terms this means that we cannot discriminate changes in β_1 from the changes in β_2 (we can estimate the quantity $(\beta_1 + \beta_2)$ well however).

Let us compare the inverse information matrix above to that of the $AR(1)$ process measured with white noise. We assume that $\beta = 0.95$, the values of γ and σ_u^2 remain unchanged from the above:

$$\mathcal{IF}^{-1} \cong \begin{pmatrix} 0.177 & -0.0885 & -0.0591 \\ -0.0885 & 0.0778 & 0.0555 \\ -0.0591 & 0.0555 & 0.584 \end{pmatrix}$$

The difference is striking. Keeping in mind that the inverse of the Fisher matrix puts the lower bound the norm of the variance-covariance matrix of the estimates, we see how much uncertainty is introduced by extending the class of possible latent processes from $AR(1)$ to $AR(2)$.

The analysis above shows that the problem of identifying the structure of the latent process of the SCD model using the QML approach has properties which make it similar to an ill-posed problem. It is not an ill-posed problem in the strict sense of the definition because the unique solution exists, and for any given accuracy there is a sample size at which this accuracy can be achieved. For practical purposes, however, acceptable variance of the parameter estimates

[12] It is possible to compute the components of the inverse Fisher matrix analytically but the expressions are very bulky. Therefore we present a numerical illustration here.

given available sample sizes can be achieved by restricting the space of possible solutions to the class of $AR(1)$ models which is a standard approach to solving ill-posed problems.

It has been noticed above that the asymptotic standard errors of the SCD model when the latent process is $AR(1)$ are very narrow, and the goodness of fit is determined primarily by how well the model is specified. We see that the class of SCD models with the $AR(2)$ is too wide given the information available which results in large asymptotic errors. Intuition suggests to look for a model in a class more flexible than the SCD with $AR(1)$ latent process but which would have a different structure than the SCD with $AR(2)$ latent process The FISCD model introduced in the Appendix is an attempt to find such class of models.

As far as the parametric distributional assumptions concerned, analysis of the p-value plots and of the p-value discrepancy plots of the empirical distribution of $\left\{\hat{\xi}_i\right\}$ against the log-Weibull distribution shown on figures 7 and 9 does not indicate gross incompatibilities of the adopted parametric form either for the trade durations data or for the price durations data. The shape of the p-plots is very similar to the shape observed in Bauwens and Veredas (1999) for trade durations of Boeing stocks.

Analysis of the p-value plot of the empirical distribution of trade and price durations $\{\hat{u}_i\}$ against the normal distribution suggests a distinct departure from normality (Figure 8) In both cases the empirical distribution has fatter tails than the normal distribution. Our observations with respect to the tails of the empirical distribution of $\{u_i\}$ are opposite to those reported in Bauwens and Veredas (1999) for the trade and price durations of the Boeing stocks; in the latter case, the empirical distribution had thinner than normal tails.

Formal goodness-of-fit tests based on the Anderson-Darling statistic and Cramèr-von Mises statistics reject the parametric distributional assumptions of the model at any conventional level. We expected that the parametric assumptions would be rejected by these tests because they were very powerful given a typical size of the sample and because our model was capable, by design, of capturing only the most general features of the data. However, a closer look at the behaviour of the goodness-of-fit statistics illuminates directions for improvement of the model and for further research.

We expect that our model which treats the data as continuous series, does not describe well transitions from one day to another. The values of the goodness-of-fit the statistics decrease multiple times if we censor from the samples the residuals corresponding to initial and final moments of trading in every day (the statistics still remain in the rejection region however). These quantitative results confirm the graphical analysis above: the rejection of the assumption of normality of $\{\hat{u}_i\}$ is overwhelmingly stronger than the rejection of the assumption about the parametric form of $\left\{\hat{\xi}_i\right\}$, the latter being still significant on the 1% level.

5.5 Discussion

Our empirical analysis suggests several ways to improve the performance of the model. When estimating the seasonal components non-parametrically, we chose the bandwidth parameters based on visual analysis of the graphs. A more formal approach based, for example, on cross-validation, is of course possible. The difficulty with automated choice of the bandwidth parameters is related to the fact that the methods of automated selection break down if the errors are dependent. There are few methods that can handle the dependent data (for example, Francisco-Fernandez, Opsomer, and Vilar-Fernandez (2001)) but they are relatively complicated algorithmically and have been developed only for univariate nonparametric regression.. Thus, even if we adopt an automated algorithm to choose the bandwidths, we shall still have to accept a degree of sub-optimality in this choice.

Treating the data as continuous series leaves may be inappropriate when we study the trade durations of contracts with longer horizons to expiration, where we typically have a few records per day, or when we study price durations. The asymptotic theory developed in Dunsmuir (1979) is directly applicable to Kalman QML estimation provided that the filter is initialized with the diffuse prior. We have amended the asymptotic theory of QML estimation in a way that it is applicable to the case when the data consists of a set of independent subsamples (see part 7.2 of the Appendix). We have also tried two alternative approaches: we initialized the Kalman filter at 8:00 each day using the average logarithm of trade durations (seasonally adjusted) between 7:40 and 8:00 and we initialized the Kalman filter in the beginning of each day with the sample mean. The estimation results were very close for all the three methods when we analyze the whole data set, especially in the case of trade durations. This is because the records corresponding to the nearest to expiration contract constitute a larger part of the data, and the number of records per day is large, hence, the initialization of the Kalman filter affects only marginally the value of the quasi-likelihood function. We shall include the estimation of the SCD model with daily re-initialization of the Kalman filter in the later versions of this study. We would like to notice in the end that the initialization of Kalman filter in the beginning of every day provides us with at tool to introduce in a non-trivial way the information accumulated overnight.

We have observed that the empirical distribution of $\{\hat{u}_i\}$ departs from the normal distribution, especially in trade durations series. We expect that this problem will be mitigated if daily initialization of the Kalman filter is used: large overnight innovations may be responsible in part for fatter tails for $\{u_i\}$. One can attempt however using a different parametric form of the distribution of $\{u_i\}$: a Student t distribution, for example. The model with normal innovations of the latent process will be nested into Student t parametrization.

Finally, it would be interesting exercise to investigate forecasting abilities of our model.

6 Conclusions

We have mentioned in the introduction that, when we talk about a financial asset, the latent economic process can be viewed as manifesting itself through the transmission mechanism of the institutional structure of trading. A relatively simple model used in this study is far from describing the economic process and the transmission mechanism in details; it captures dynamic properties of the data without revealing the structure of the data-generating mechanism. Nonetheless, if we adopt a constructive approach to modeling the generating mechanism for high-frequency financial series, a parsimonious description of the output signal given by the SCD model will be a valuable resource in synthesising the transmission function which given as an input a signal described by one of existing models of asset price dynamics will generate the output with dynamics similar to that of empirical point processes investigated in this study.

We believe that the synthesis of models of trading mechanism which would bridge the gap between the dynamic financial models and the empirical models of high-frequency financial series will be a promising area of research. Design of a realistic model of trading would be a very complex task requring substantial resourses. However, even a simple stylized model may provide further insights into the microstructure of financial markets. Imagine, for example, that the latent price process follows a stochastic volatility model and that a new transaction occurs when the latent price deviates from the last observed price by a given margin[13]. This model has a continuous-time process as an input and a point process as an output and is probably the simpliest imaginable model of the trading mechansm. Empirical evidence suggests the presence of long memory in the volatility process and in the durations process. We conjecture that in the model just described long memory in the volatility of the laten price process will translate into long memory in the durations process.

Summarising empirical findings of this study, we observed that while the estimated parameters of the SCD process are statistically different from contract to contract, qualitatively the process does not change much over the years. This is an indication that the model captures some invariant properties of the economic process and the transmission mechanism. We observed that the SCD parameters, with a certain leeway, are also stable across the expiration horizon, the last statement being more accurate with respect to the price durations process. Price of a future has almost a functional relationship with the price of the underlying asset. From this it follows that the life-cycle seasonality of futures durations depends on more on characteristics of the trading mechanism than of the properties of the latent economic process.

Preliminary analysis shows that the fit of the model improves noticably if the overnight and the weekend interruptions in trading are taken into account. The simpliest way to do this is to assume that in the beginning of every trading day the durations process is initialized with the unconditional mean. This approach has the advantage that the existing asymptotic theory of the QML estimators

[13]The framework with informed and uninformed traders can be used to explain the liquidity of our stylized market.

can be applied with minor modifications (see 7.2 of the Apendix). Alternatively, the durations process can be initialized in the beginning of every trading day using auxiliary information available to the econometrician. Designing various initialization procedures is probably not very interesting for an academic researcher because these procedures would vary depending on the specific exchange and the asset considered, and other factors. However, the initialization of the process will be crucial in any practical application of the model. The difficulty which one faces when using an informative prior to initialize the process in the beginning of every trading day is that the asymptotic theory for the estimates cannot be derived from the results of Dunsmuir (1979) because the essential assumption of ergodicity is violated, and the theory would have to be developed from scratch.

The science of signal processing has traditions of the analysis of maximal achievable resolution of a system and of informational analysis of a transmission channel. We believe that the econometrics of high-frequency financial data can develop on these traditions. A simple example of informational analysis of the SCD model given in this study is a modest contribution to this interesting direction of research. We illustrate the practical limitation, given the model structure and the information available, of our abilities to estimate and/or identify the signal (the latent process), and how these limitations can be discovered through the informational analysis of the system comprising the model and the estimation algorithm" (our analysis of the SCD model applies also to stochastic volatility models that have a similar mathematical structure). Our estimation and identification capabilities can be improved either by introducing new a priori information[14] or by using an estimation algorithm, if there exists one, that makes a better use of the existing information. In economics, choosing an alternative model of the transmission channel and of the signal itself can often be a productive, in contrast with the natural sciences where due to the established methodological paradigm the acceptable choice of models is much more restricted.

Finally, a dynamic model of durations should accommodate the possibility of long memory because the empirical evidence strongly suggests the presence of long-range dependence in the durations process It is not a very difficult task to estimated the long-memory parameter; it is more difficult to estimate both the high-frequency dynamics and the long-memory of the process. The FISCD model provides a parametric framework that allows, in theory, modelling both the high-frequency dynamics and the low frequency (long-memory) properties. The advantage of QML estimation in the spectral domain of the FISCD model is its algorithmic and computational simplicity. Investigation of the properties of these estimates beyond the strong consistency, which is known to hold, remains a challenging theoretical problem.

[14] For example, this a priori information can be in a form that the latent process is an $AR(1)$ process.

7 Appendix

7.1 Computation of asymptotic standard errors for the log-Weibull SCD model

Application of the Dunsmuir's (1979) asymptotic theory to the case when the quasi-likelihood function is estimated using the Kalman filter is described in Harvey (1989), pp.220-221. The QML estimate of the parameter vector $\boldsymbol{\theta}$ of the model is asymptotically normal, unbiased and has the variance-covariance matrix $\mathbf{C} = 2\mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{B}\mathbf{A}^{-1}$ where

$$\mathbf{A} = \frac{1}{2\pi}\int_{-\pi}^{\pi}\frac{\partial\log\left(g\left(\lambda\right)\right)}{\partial\boldsymbol{\theta}}\frac{\partial\log\left(g\left(\lambda\right)\right)}{\partial\boldsymbol{\theta}'}d\lambda$$

is proportional to the information matrix of the process and

$$\mathbf{B} = \kappa\left[\frac{1}{2\pi}\int_{-\pi}^{\pi}\frac{\partial\log\left(g\left(\lambda\right)\right)}{\partial\boldsymbol{\theta}}d\lambda\right]\left[\frac{1}{2\pi}\int_{-\pi}^{\pi}\frac{\partial\log\left(g\left(\lambda\right)\right)}{\partial\boldsymbol{\theta}'}d\lambda\right].$$

Here, $g\left(\lambda\right)$ is the spectral-generating function of the process. It is easy to write down the spectral-generating function of the process $\hat{d}_i$, the spectrum of this process will be different form the spectrum of the AR process with normal innovations only by a constant. The spectrum of the log-durations of the SCD model when the latent process is $AR\left(1\right)$ and the observation error has a log-Weibull distribution is

$$g\left(\lambda\right) = \frac{\pi^2}{6\gamma} + \frac{\sigma^2}{\left(1-2\beta\cos\lambda+\beta^2\right)^2}.$$

Analytical computation of the matrices $\mathbf{A}$ and $\mathbf{B}$ is a conceptually straightforward but tedious task. We opted to use approximate discrete representations of these integrals in our computations. Given our sample sizes, the discrete approximation is very accurate and much easier to implement.

As in (Harvey 1989), we define discrete analogs of the matrices $\mathbf{A}$ and $\mathbf{B}$,

$$\mathbf{A}_T = \frac{1}{T}\sum_{m=0}^{T-1}\frac{\partial\log g\left(\frac{2\pi m}{T}\right)}{\partial\boldsymbol{\theta}}\frac{\partial\log g\left(\frac{2\pi m}{T}\right)}{\partial\boldsymbol{\theta}'}$$

$$\mathbf{B}_T = \kappa\left[\frac{1}{T}\sum_{m=0}^{T-1}\frac{\partial\log g\left(\frac{2\pi m}{T}\right)}{\partial\boldsymbol{\theta}}\right]\left[\frac{1}{T}\sum_{m=0}^{T-1}\frac{\partial\log g\left(\frac{2\pi m}{T}\right)}{\partial\boldsymbol{\theta}}\right],$$

$\lim_{T\to\infty}\left(\mathbf{A}_T\right) = \mathbf{A}$, $\lim_{T\to\infty}\left(\mathbf{B}_T\right) = \mathbf{B}$.

The information matrix that we are considering in the text of the paper is defined as (we omit the component which is due to the non-normality for the sake of simplicity):

$$\mathcal{IF} = \frac{1}{4\pi}\int_{-\pi}^{\pi}\frac{\partial\log\left(g\left(\lambda\right)\right)}{\partial\boldsymbol{\theta}}\frac{\partial\log\left(g\left(\lambda\right)\right)}{\partial\boldsymbol{\theta}'}d\lambda \qquad (5)$$

7.2 Asymptotic theory when the sample consists of T independent subsamples

Claim 1 *Assume that the sample from the process $\{z_i(\theta_0)\}$ satisfying the conditions of Theorem 2.1 of (Dunsmuir 1979) consists of T independent subsamples $\left\{Z_i^j\right\}_{i=1}^{N_j}$, $N = \sum_{j=1}^{T} N_j$ (we consider only scalar processes to simplify the exposition). Assume that $N \to \infty$ in such way that $\frac{N_j}{N} \to \lambda_i$. Denote $g_\theta(\lambda)$ the spectral density of the process $z(\theta)$ and $I^j(\lambda)$ - the periodogram of the subsample Then if $\bar{\theta}_N = \arg\max \bar{L}_N(\theta)$, $\bar{L}_N(\theta) = \sum_{j=1}^{T} L^j(\theta)$, where*

$$L^j(\theta) = \log\left(\frac{1}{2\pi}\int g_\theta(\lambda)\, d\lambda\right) + \frac{1}{2\pi}\int \frac{I^j(\lambda)}{g_\theta(\lambda)} d\lambda$$

then the quantity $N^{1/2}\left(\bar{\theta}_N - \theta_0\right)$ has an asymptotic normal distribution with zero mean and the covariance matrix

$$\left(T^{-2}\sum_{j=1}^{T}\lambda_j^{-1}\right)\Omega^{-1}(2\Omega + \Pi)\,\Omega^{-1}$$

where

$$\Omega = \frac{1}{2\pi}\int \frac{\partial \ln g_\theta(\lambda)}{\partial \theta'}\frac{\partial \ln g_\theta(\lambda)}{\partial \theta} d\lambda$$

and

$$\Pi = \kappa\left(\frac{1}{2\pi}\int \frac{\partial \ln g_\theta(\lambda)}{\partial \theta'} d\lambda\right)\left(\frac{1}{2\pi}\int \frac{\partial \ln g_\theta(\lambda)}{\partial \theta} d\lambda\right).$$

κ is the forth cumulant of the innovations.

Proof. The proof of the claim requires only a slight modification of the proof of the theorem 2.1 of (Dunsmuir 1979). We can see that under the conditions of the claim the quantity $\frac{\partial^2}{\partial\theta\partial\theta'}\bar{L}_N(\theta) \to_p T\Omega$ (each $L^j(\theta)$ coverges to Ω), and the quantity $N^{1/2}\frac{\partial}{\partial\theta}\bar{L}_N(\theta)$ is asymptotically normal with the variance-covariance matrix

$$\sum_{j=1}^{T}\lambda_j^{-1}(2\Omega + \Pi)$$

The result of the claim immediately follows. ■

7.3 Long-memory in the dynamics of durations

7.3.1 FISCD model

The fractionally-integrated stochastic conditional duration model is specified as follows (we give here the general specification but we shall later consider only

$FISCD\left(1,x,0\right)$):

$$\begin{aligned} d_i &= \mu\left(\gamma\right) + \psi_i + \xi_i \\ \phi\left(L\right)\left(1-L\right)^x\left(\psi_i - \bar{\psi}\right) &= \eta\left(L\right)u_i, \end{aligned} \tag{6}$$

where $\mathbf{E}\left[\psi_i\right] = \bar{\psi}$, all roots of the polynomials $\phi\left(L\right)$ and $\eta\left(L\right)$ are outside the unit circle. As before, $\{\exp\left(\xi_i + \mu\left(\gamma\right)\right)\}$ are i.i.d. with a distribution having a positive support, and $\{u_i\}\tilde{\ } n.i.i.d.\left(0,\sigma_u^2\right)$. The spectral density of this process exists provided that $x < \frac{1}{2}$ and it has the following form:

$$g\left(\lambda\right) = \frac{1}{2\pi}\left(\sigma_\xi^2\left(\gamma\right) + \frac{\sigma_u^2\left|\eta\left(e^{-i\lambda}\right)\right|^2 2^{-x}\left(1-\cos\lambda\right)^{-x}}{\left|\phi\left(e^{-i\lambda}\right)\right|^2}\right),$$

where $\sigma_\xi^2\left(\gamma\right)$ is the variance ξ_i. The spectrum has a singularity at $\lambda = 0$ which is integrable for the stationary range of the fractional integration parameter x. When the process is stationary one can easily compute the autocovariances of the process knowing the spectrum (they cannot be expressed in elementary functions but it is not important for our exposition):

$$\gamma_k = \int_{-\pi}^{\pi} g\left(\lambda\right)\exp\left(i\lambda k\right)d\lambda,\ k = 0, 1, ...$$

where γ_k denotes the k-th autocovariance. The expressions for the autocovariances of the process $\{D_i\}$ are bulky and we do not give them here.

7.3.2 QML estimation of FISCD in spectral domain

The asymptotic theory of Dunsmuir (1979) is not applicable to fractionally-integrated processes and the quasi-likelihood can not be computed using the state-space representation. Several approaches have been suggested to estimating models with latent variables that have the structure similar to that of the equation 6. These approaches, either based on the generalized method of moments or on computing the likelihood using simulations, have the common property: they are very computationally intensive. It is possible to compute the quasi-likelihood function of the $FISCD(p,x,q)$ process based on the sample spectrum of the process (this approach has been suggested in the context of stochastic volatility processes with long memory). The applicability of the spectral QML estimation technique (which gives essentially a Whittle-type estimator) to the problem in hand is based on results of Breidt, Crato, and de Lima (1998) where it has been shown that the maximizer of the expression 7 is a strongly-consistent estimator of $\boldsymbol{\theta}$ provided that the parameter space is compact and the parameter is uniquely identified at the true value. There is no asymptotic theory available for the spectral QML estimator.

The first steps in estimating the FISCD model are the same as those in the estimation of the SCD model. After the seasonal adjustment and subtracting the mean the model to be estimated has the following form:

$$d_i^* = \psi_i^* + \xi_i$$
$$\psi_i^* = (1-\beta L)^{-1}(1-L)^{-x} u_i$$

The logarithm of the spectral likelihood function of the process $\{d_i^*\}$ is

$$L(\boldsymbol{\theta}) = -\frac{2\pi}{n}\sum_{k=1}^{[N/2]}\left(\log g_{\boldsymbol{\theta}}(\lambda_k) + \frac{I_N(\lambda_k)}{g_{\boldsymbol{\theta}}(\lambda_k)}\right) \tag{7}$$

where $g_{\boldsymbol{\theta}}(\lambda)$ is the spectrum of the $FISCD(1,x,0)$ process with the parameter vector $\boldsymbol{\theta} = (\sigma_u^2, \gamma, x, \beta)$,

$$g(\lambda) = \frac{1}{2\pi}\left(\frac{\pi^2}{6\gamma^2} + \frac{\sigma_u^2 2^{-x}(1-\cos\lambda)^{-x}}{\beta^2 - 2\beta\cos\lambda + 1}\right)$$

and $I_N(\lambda)$ is the sample periodogram. $\lambda_k = \frac{2\pi k}{N}$, $k = 0, 1, 2, ..., N$.

7.3.3 Estimation results and discussion

The estimated $FISCD(1,x,0)$ parameters for trade and price durations are presented on Figure 11 in a graphic format similar to that used for depicting in this paper estimates of the SCD parameters. The estimates of x vary between 0.42 and 0.52 for trade durations and between 0.23 and 0.65 for price durations. We have several contracts where the QML point estimates of x are greater than 0.5, i.e. are outside the stationarity region. This does not mean, of course, that the durations process is non-stationary since the confidence intervals are not available.

The estimates of the parameter γ of the Weibull distribution lie between 0.95 and 1.26 for trade durations and between 1.2 and 2.65 for price durations. The estimates of σ_u^2 and β are very volatile and the data indicates that these two quantities have strong negative correlation. It seems also that we observe two types of estimates: those with higher $\hat{\sigma}_u^2$ and $\hat{\beta}$ close to 0 and those with smaller $\hat{\sigma}_u^2$ and estimates of β in the range of $0.4-0.6$. Again, we cannot draw definitive conclusions because we don't know the distribution of the estimated parameters even asymptotically. We can hypothesize that the instability of parameter estimates may be caused in part by problems with the data. A simulation study will probably provide a controlled environment and help to discover properties of estimated parameter of the FISCD model.

References

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., and D. Veredas (1999): "The Stochastic Conditional Duration Model: A Latent Variable Model for the Analysis of Financial Durations," forthcoming in Journal of Econometrics.

Breidt, F. J., N. Crato, and P. de Lima (1998): "The Detection and Estimation of Long Memory in Stochastic Volatility," *Journal of Econometrics*, 83, 325–348.

Breiman, L., and J. H. Friedman (1985): "Estimating Optimal Transformations for Multiple Regression and Correlation (with Discussion)," *Journal of American Statistical Association*, 80, 580–619.

Dunsmuir, W. (1979): "A Central Limit Theorem for Parameter Estimation in Stationary Vector Time Series and its Application to Models for a Signal Observed with Noise," *Annals of Statistics*, 7, 490–506.

Engle, R., and J. Russell (1998): "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data," *Econometrica*, 66(5), 1127–1162.

Fan, J., and I. Gijbels (1996): *Local Polynomial Modelling and its Applications*. Chapman and Hall, Suffolk.

Francisco-Fernandez, M., J. Opsomer, and J. Vilar-Fernandez (2001): "A Plug-in Bandwidth Selector for Local Polynomial Regression Estimator with Correlated Errors," .

Friedman, J. H., and W. Stuetzle (1981): "Projection Pursuit Regression," *Journal of American Statistical Association*, 76, 817–823.

Gerhard, F., and N. Hautsch (2002): "Volatility Estimation on the Basis of Price Intensities," *Journal of Empirical Finance*, 9, 57–89.

Ghysels, E., C. Gourieroux, and J. Jasiak (1997): "Stochastic Volatility Duration Model," *CREST working paper 9746*.

Ghysels, E., and J. Jasiak (1998): "GARCH for Irregularly Spaced Financial Data: The ACD-GARCH Model," *Studies in Nonlinear Dynamics and Econometrics*, 2, 133–149.

Gouriéroux, C., J. Jasiak, and G. L. Fol (1999): "Intra-Day Market Activity," *Journal of Financial Markets*, 2, 193–226.

Grammig, J., and K.-O. Maurer (1999): "Non-Monotonic Hazard Functions and the Autoregressive Conditional Duration Model," *The Econometrics Journal*, 3, 16–38.

Harvey, A. C. (1989): *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Jasiak, J. (1998): "Persistence in Intertrade Durations," *Finance*, 19, 166–195.

Linton, O., and J. P. Nielsen (1995): "A Kernel Method of Estimating Structured Nonparametric Regression Based on Marginal Integration," *Biometrika*, 82, 93–100.

Ruiz, E. (1994): "Quasi-Maximum Likelihood Estimation of Stochastic Volatility Models," *Journal of Econometrics*, 63, 289–306.

Russel, J., and R. Engle (1998): "Econometric Analysis of Discrete-Valued Irregularly-Spaced Financial Transactions Data Using a New Autoregressive Conditional Multinomial Model," .

Stephens, M. A. (1976): "Asymptotic Results for Goodness-of-Fit Statistics with Unknown Parameters," *The Annals of Statistics*, 4, 357–369.

——— (1977): "Goodness of Fit for the Extreme Value Distribution," *Biometrika*, 64, 583–588.

Veredas, D., J. Rodriguez-Poo, and A. Espasa (2002): "On the (Intraday) Seasonality and Dynamics of a Financial Point Process: A Semiparametric Approach," *CORE DP 2002/23*.

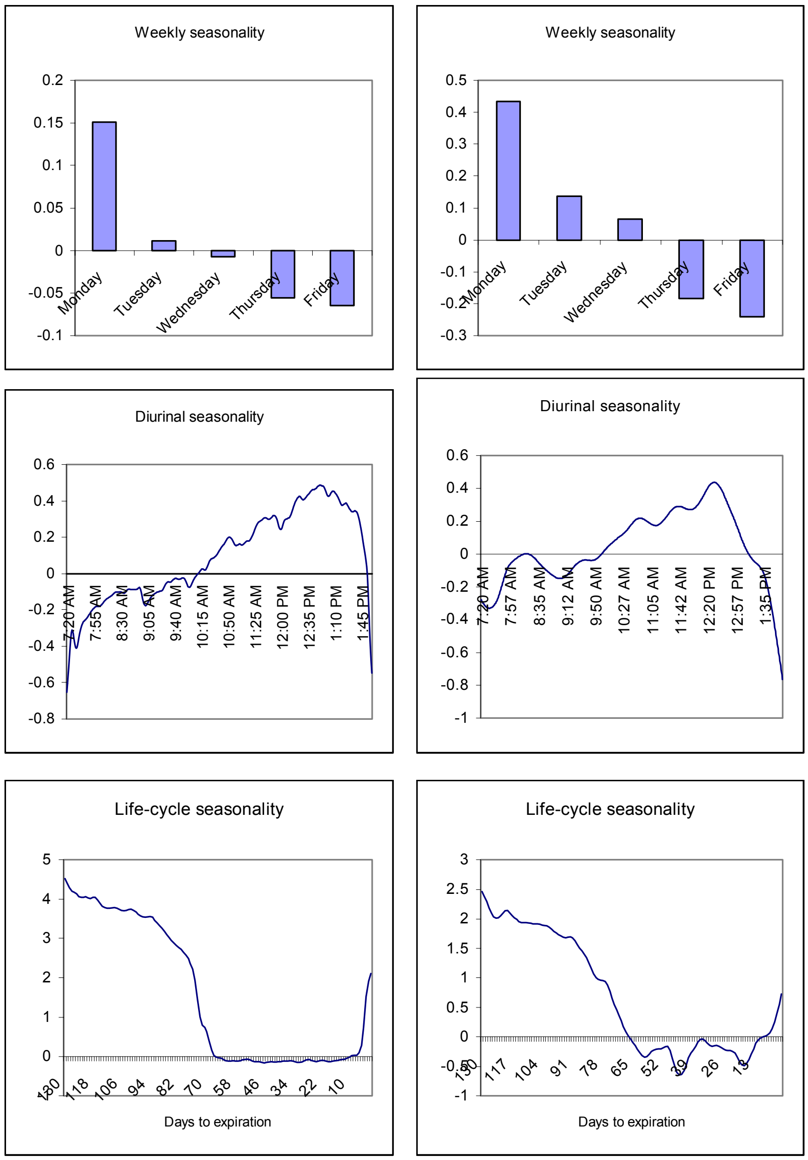

Figure 1: Estimates of the seasonal components $\hat{A}_\delta$, $\hat{\chi}(t)$ and $\hat{\zeta}(\tau)$ of trade durations (left column) and of price durations (right column)

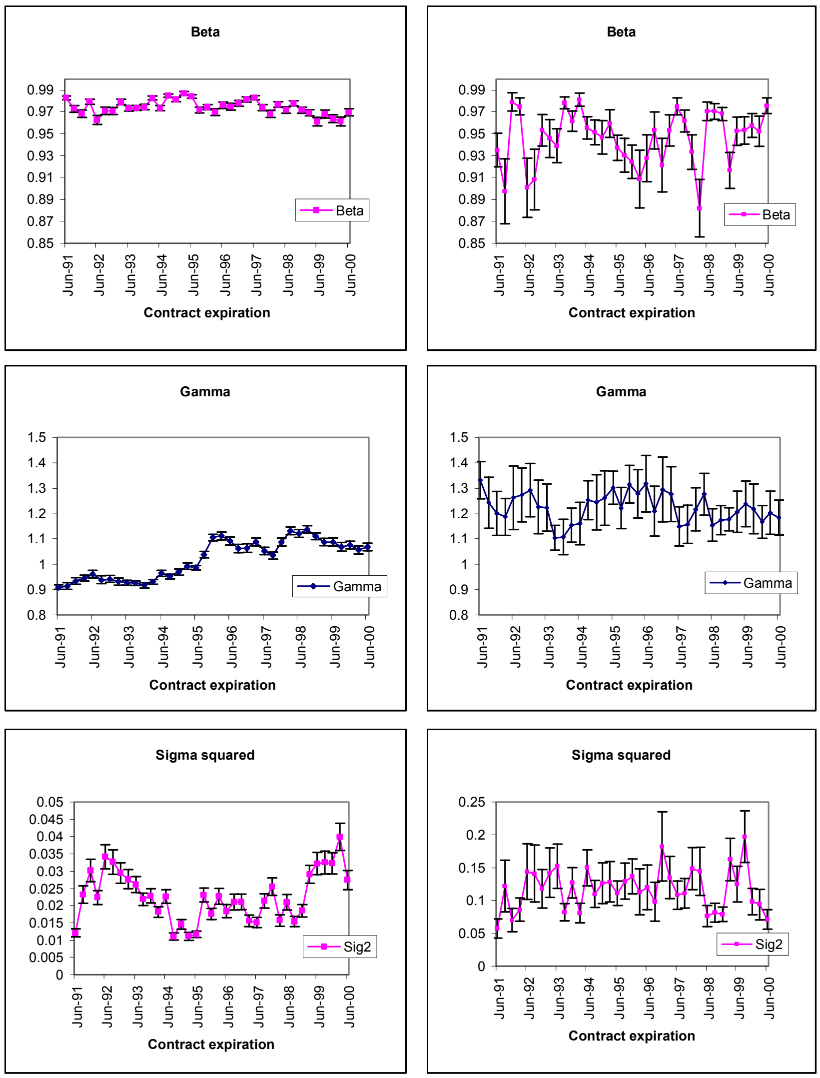

Figure 2: SCD parameters for contracts with different expiration dates. Left column - trade durations, right column - price durations. Error bars correspond to two SE.

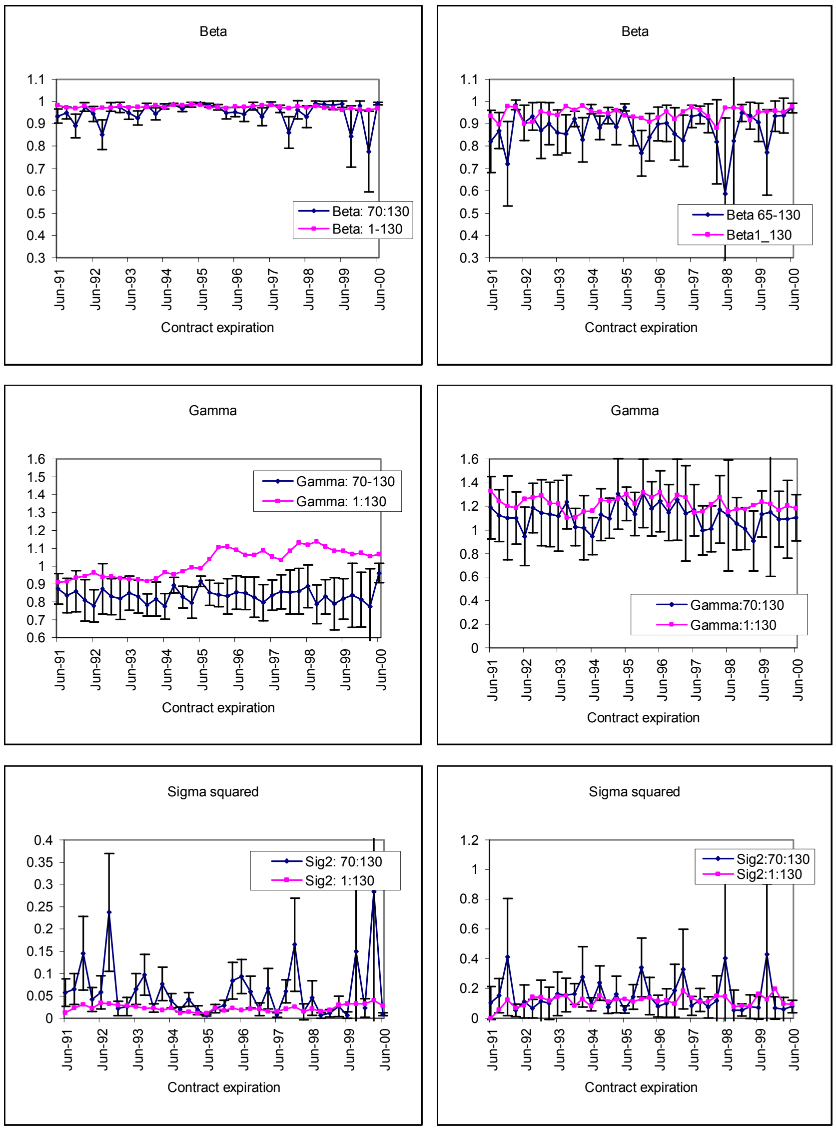

Figure 3: SCD parameters estimated using records with 70-130 business days to expiration. Left column - trade durations, right column - trade durations. Error bars correspond to two SE.

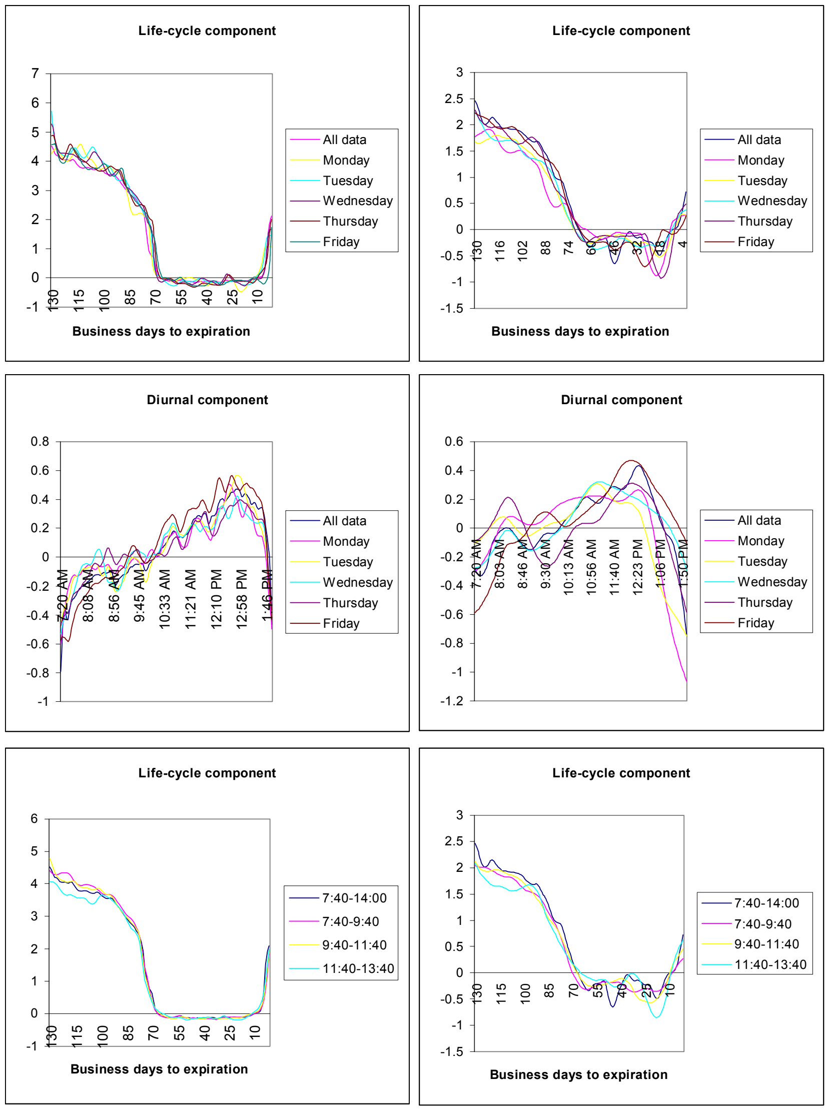

Figure 4: Graphic verification of the additivity property. Top two rows - life-cycle and diurnal contribution when the data is separated according weekday. Bottom row - life-cycle contribution estimated for different periods of trading day. Left column - trade durations, right - price durations.

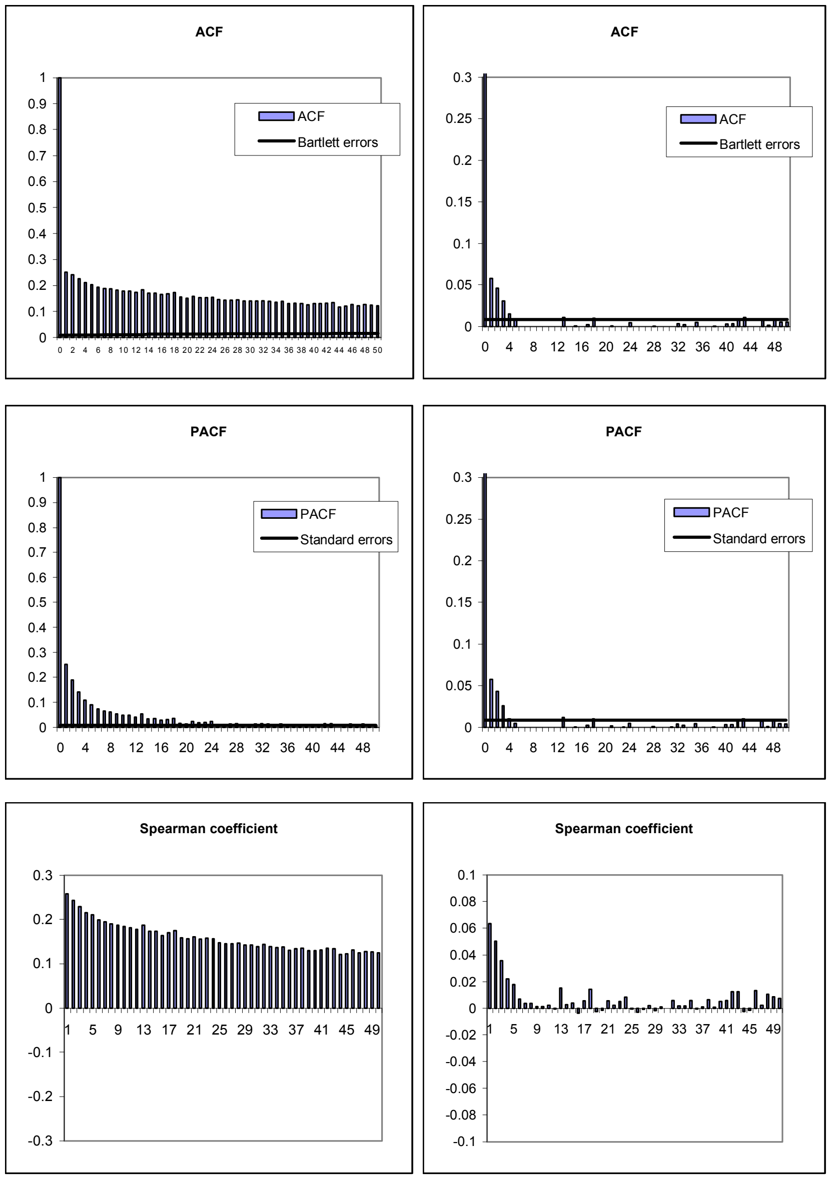

Figure 5: Measures of dependence for trade log-durations (left column) and for SCD residuals (right column)

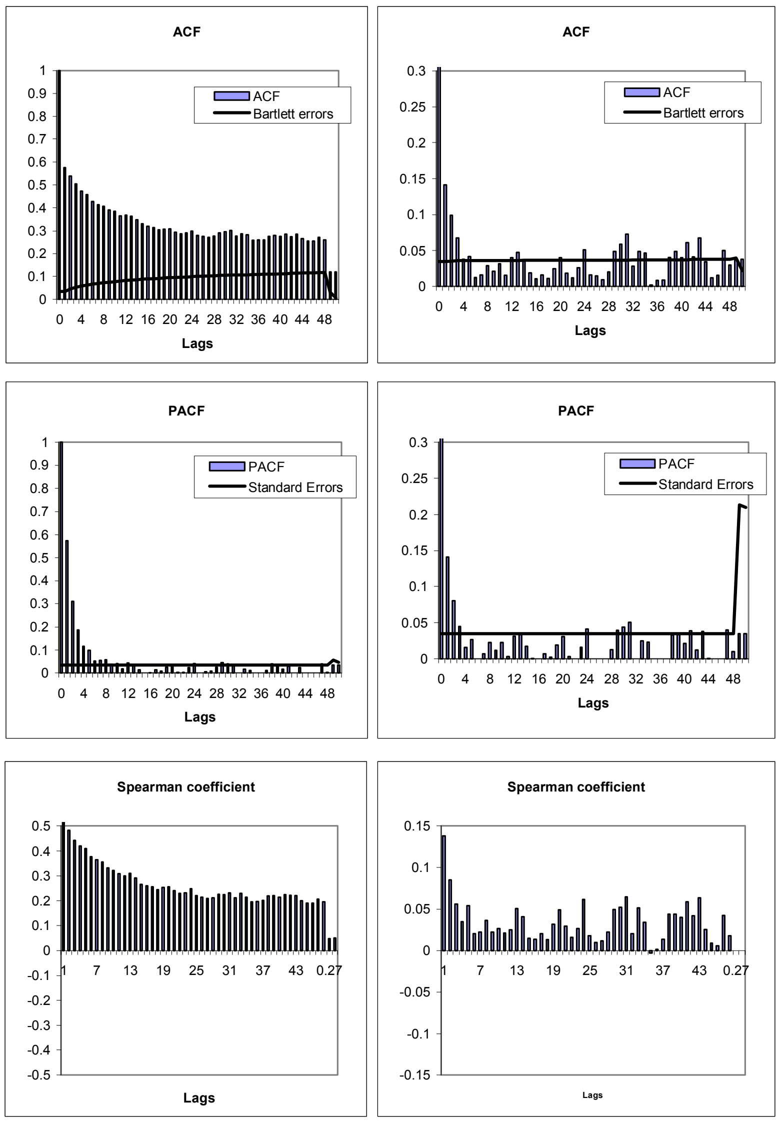

Figure 6: Measures of dependence for price log-durations (left column) and the SCD residuals (right column)

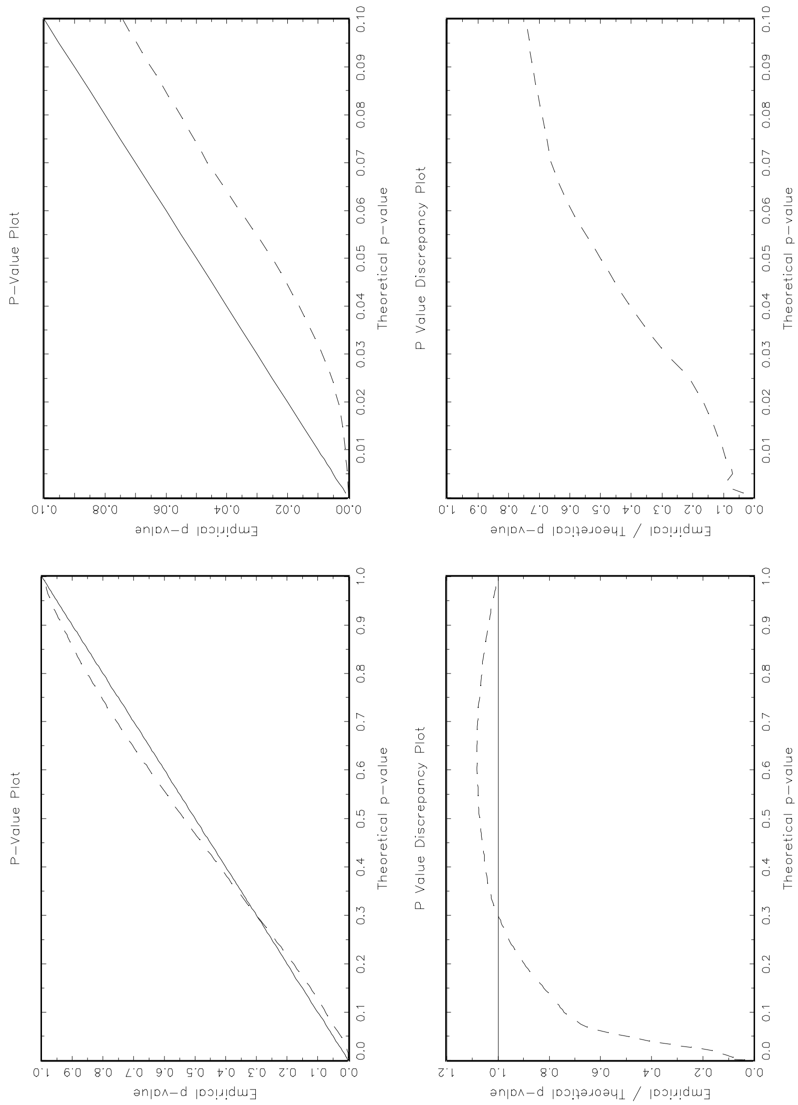

Figure 7: Trade durations. Goodness of fit of the empirical distribuiton of $\hat{\xi}$.

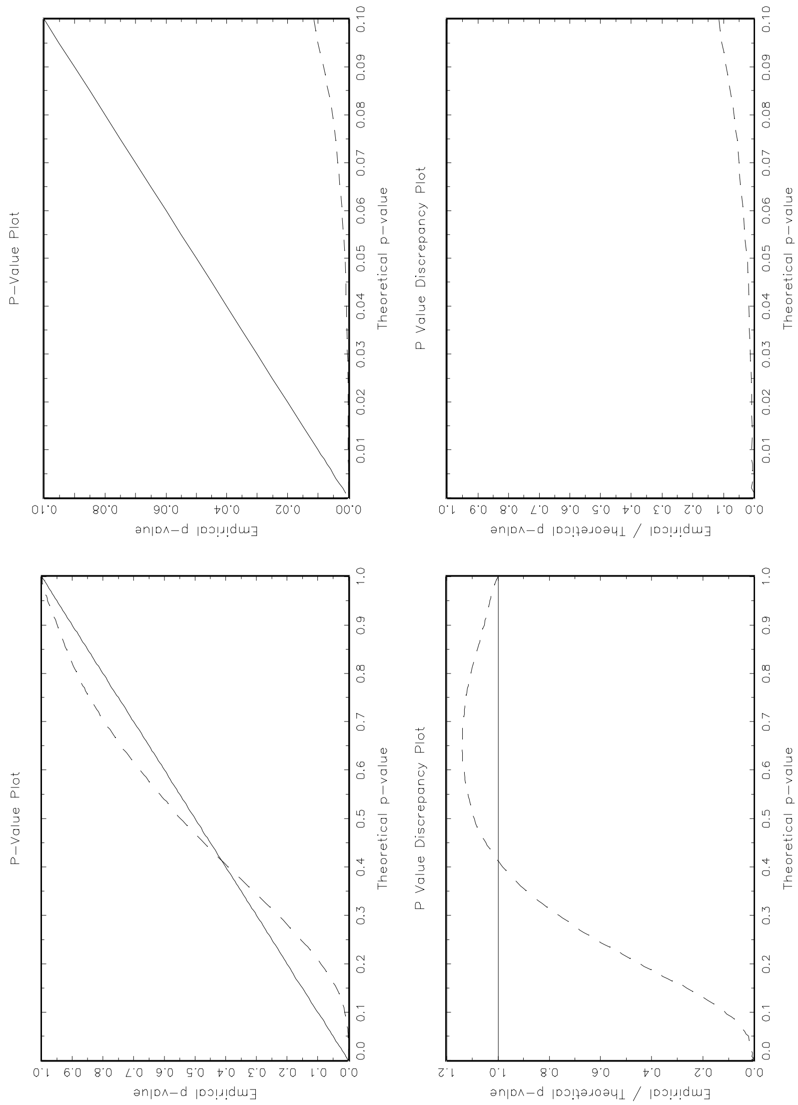

Figure 8: Trade durations. Goodness of fit of empirical distribution of $\hat{u}$.

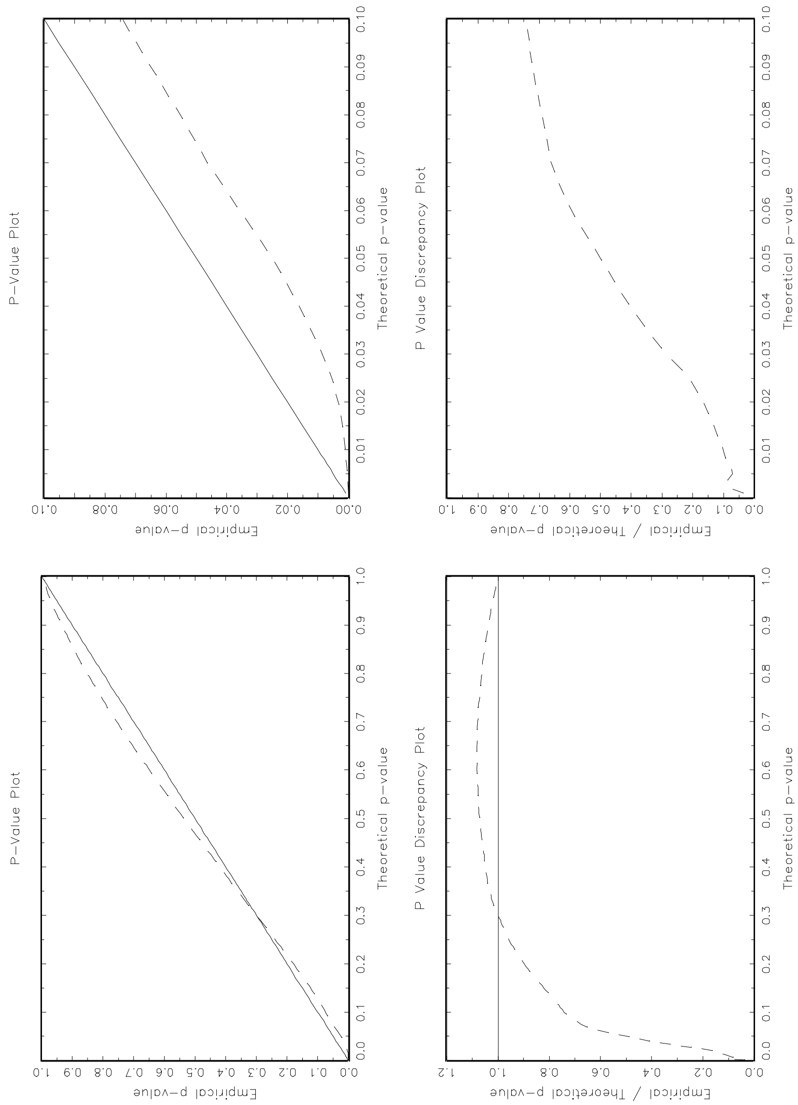

Figure 9: Price durations. Goodness of fit of empirical distribution of $\hat{\xi}$

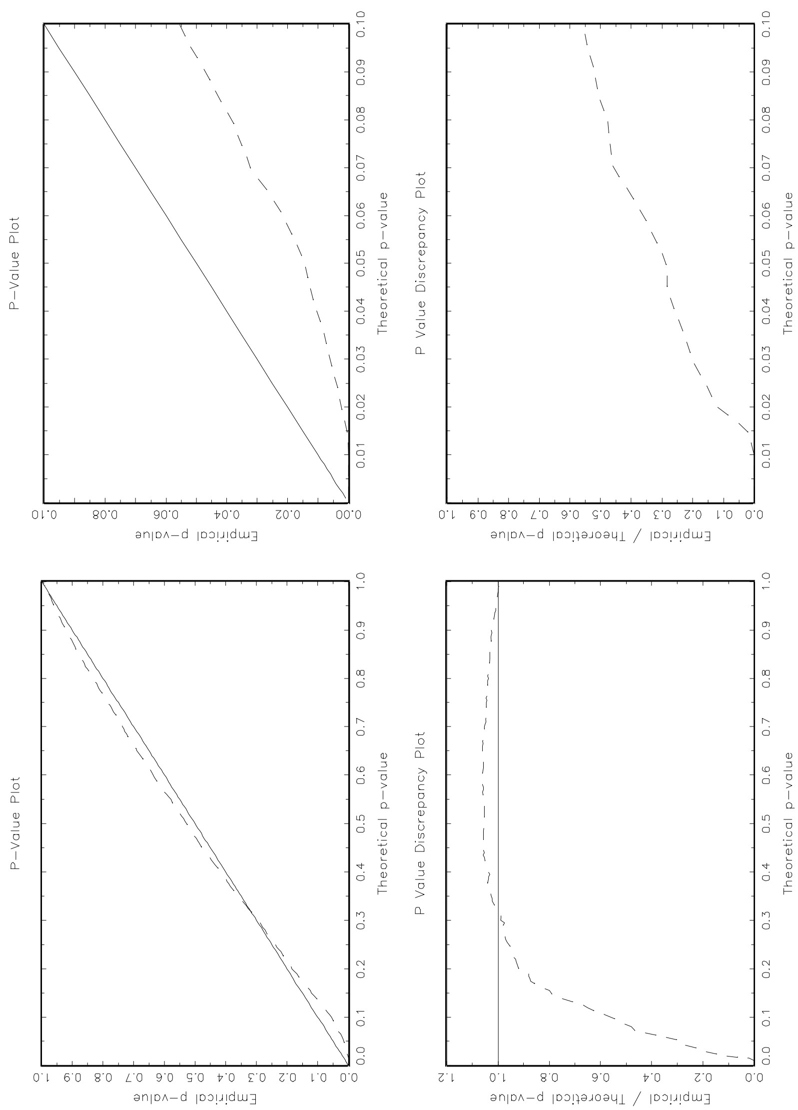

Figure 10: Price durations. Goodness of fit of empirical distribution of $\hat{u}$

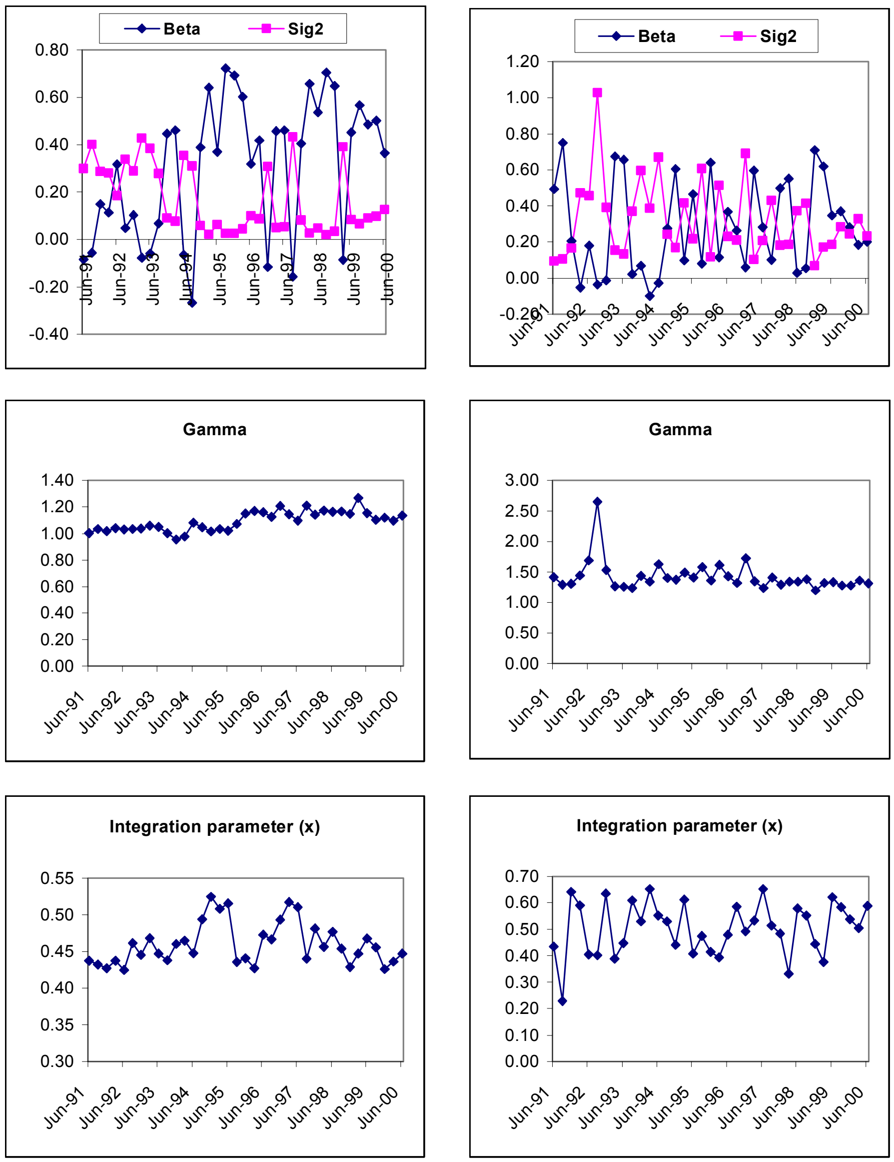

Figure 11: FISCD(1,x,0) parameters for contracts with varying expiration dates. Trade durations are in the left column, price durations - in the right column.