Empirical Economics (2003) 28:709–731
DOI 10.1007/s00181-003-0155-7

EMPIRICAL ECONOMICS

Asymmetric ACD models: Introducing price information in ACD models

Luc Bauwens[1], **Pierre Giot**[1,2]

[1] CORE and Department of Economics, Université catholique de Louvain, Belgium
[2] Department of Business Administration & CEREFIM, University of Namur, Rempart de la Vierge 8, B-5000 Namur, Belgium (e-mail: pierre.giot@fundp.ac.be)

First version received: January 2001/Final version received: August 2002

Abstract. This paper proposes an asymmetric autoregressive conditional duration (ACD) model, which extends the ACD model of Engle and Russell (1998). The asymmetry consists of letting the duration process depend on the state of the price process. If the price has increased, the parameters of the ACD model can differ from what they are if the price has decreased. The model is applied to the bid-ask quotes of two stocks traded on the NYSE and the evidence in favour of asymmetry is strong. Information effects (Easley and O'Hara 1992) are also empirically relevant. As the model is a transition model for the price process, it delivers 'market forecasts' of where prices are heading. A trading strategy based on the model is implemented using tick-by-tick data.

Key words: Duration and transition model, high frequency data, market microstructure, forecasting

JEL classification: C10, C41, G10

While remaining responsible for any error in this paper, the authors would like to thank R. Anderson, G. Le Fol, C. Gouriéroux, J. Jasiak, W. Pohlmeier, A. Roell, O. Scaillet, S. Wei and three anonymous referees for useful remarks and suggestions on previous versions. The authors would also like to thank A. Ruttiens from KBC-CBC for useful discussions on practical issues related to trading.
Support of the European Commission Human Capital and Mobility Program through the network 'Econometric inference using simulation methods' is gratefully 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

High frequency duration models were first introduced in empirical finance by Engle and Russell (1997, 1998), with the autoregressive conditional duration (ACD) model. In these papers, the authors did not take the traditional viewpoint of looking at the volatility of the price process, but modelled instead the durations between successive market transactions of a stock. These durations exhibit clustering and overdispersion (the standard deviation is larger than the mean). The structure of ACD models is quite similar to the structure of GARCH processes, but rather than specifying a dynamic model on the conditional variance of the returns, they introduce a dynamic structure on the durations.

Duration models are particularly suited for applications in a high frequency setting, as one deals with irregularly spaced data. In this framework, the time elapsed between two market events (for example a trade, a specific type of trade such as a large buy, a revision of bid-ask quotes ...) conveys meaningful information and is the object of modeling. The importance of taking into account the time dimension of the price process is particularly stressed by the recent market microstructure literature; see e.g., Easley and O'Hara (1992), O'Hara (1995), and Easley et al. (1997).

Since the seminal papers by Engle and Russell (1997, 1998), several high frequency duration models have been put forward. Bauwens and Giot (2000) introduced a logarithmic version of the ACD model, called the Log-ACD model, which is more convenient than the ACD model when conditioning variables are included in the model in order to test microstructure effects.[1] As an alternative to the Weibull distribution used in the original ACD model, Grammig and Maurer (2000) use the Burr distribution, and Lunde (1998) uses the generalized gamma distribution (both distributions include the Weibull as a particular case). Ghysels et al. (1997) proposed the stochastic volatility duration model, which accounts for stochastic volatility in the durations. Bauwens and Veredas (2002) introduced the stochastic conditional duration (SCD) model, which uses a stochastic volatility-type model instead of a GARCH-type model to model the durations. Lunde (1997) has proposed a similar type of model. More information on high-frequency duration models and related issues can be found in Bauwens and Giot (2001).

All these models share a common feature: they only take into account the duration between market events, and do not include information given by the price process (at most they include additional variables drawn from the market microstructure literature). This may be an important drawback, as the information on the price process is of primary importance. Combining information given by the price process and the duration between market events is an important extension. This issue was first addressed by Engle (2000), who proposed an ACD-GARCH model: a marginal ACD model is specified for the durations, while the volatility of the returns is modelled by a GARCH process, conditionally on the duration. The ACD-GARCH model is also studied by Ghysels and Jasiak (1998), and Grammig and Wellner (2002).

---

[1] A similar model (under the name of "Nelson form ACD model") has been proposed independently by Lunde (1998) and Engle and Russell (1998), and used in Engle and Lunde (1998).

Other recent papers contain joint models for the price and duration processes. Meddahi et al. (1998) consider a continuous time framework for modelling volatility with irregularly spaced data. Darolles et al. (2000) model the transaction price dynamics by taking into account both the irregularly spacing of the data and the discreteness of the price. Russell and Engle (2002), with the autoregressive conditional multinomial (ACM) model, combine an ACD model on the durations and a generalized linear model on conditional transition probabilities of the price process. Hafner (1999) specifies an ACD model for the duration, and a non-parametric model for the return conditionally on the current duration, the lagged returns and the lagged volume. Prigent et al. (2001) focus on option pricing formulae in incomplete markets by making use of the Log-ACD model to describe the dynamics of the price process.

In this paper, we also focus on the joint modelling of the duration process and of the price process. We use a Log-ACD model for the durations between bid-ask quotes posted by a market maker, which takes into account the direction (upward or downward) of the change of the average of the bid and ask prices. The model can be viewed as a two-state transition model (in a competing risk framework) combined with a Log-ACD model, the two states being the upward and downward price movements. As such, the model is actually an asymmetric Log-ACD model, in the sense that the duration dynamics depends on the direction of the previous price change. Under appropriate parameter restrictions, the asymmetric model becomes the Log-ACD model of Bauwens and Giot (2000).

We apply the new model to intraday data of Disney and IBM, two stocks traded on the New York Stock Exchange (NYSE). Our empirical results indicate that the asymmetric model is an improvement over the simple model. We also study the influence of several characteristics of the trade process on the bid-ask quote revision process by introducing the trading intensity, the average volume per trade, and the spread in the specification of the asymmetric model. In order to illustrate a possible practical use of the model, we use it to forecast the price changes, and we implement a trading strategy based on such forecasts. The performance of our strategy is compared to that of a benchmark strategy.

The rest of the paper is organized in the following way. In Sect. 2, we briefly review ACD models, to set up the framework. In Sect. 3, we introduce the asymmetric Log-ACD model. In Sect. 4, we apply the model to the two stocks. In Sect. 5, we examine the forecasting performance of the model and we implement the trading strategy. Section 6 concludes.

2. Review of ACD models

Let x_i be the duration between two bid-ask quotes occurring at times t_{i-1} and t_i (a quote being a collection of data relative to a buy or a sell of a security on a stock exchange), i.e. $x_i = t_i - t_{i-1}$, which is usually measured in seconds. The assumption introduced by Engle and Russell (1998) is that the time dependence in the durations can be subsumed in their conditional expectations $\mathrm{E}(x_i|\mathcal{F}_{i-1})$, in such a way that $x_i/\mathrm{E}(x_i|\mathcal{F}_{i-1})$ is independent and identically distributed. $\mathcal{F}_{i-1}$ denotes the information set available before time t_i, which includes at least the past durations. The ACD model specifies the observed duration as

$$x_i = \Psi_i \epsilon_i, \tag{1}$$

where the ϵ_i are positive IID random variables, with $\mathrm{E}(\epsilon_i) = \mu$, so that $\mathrm{E}(x_i|\mathcal{F}_{i-1}) = \mu\Psi_i$. Notice that μ is not a parameter to be estimated, but a function of the parameters of the distribution of ϵ_i. The variable ϵ_i can be divided by μ to define a new innovation of expectation equal to 1, in which case a new Ψ_i replaces $\mu\Psi_i$.

A second equation specifies an autoregressive model for the conditional durations Ψ_i:[2]

$$\Psi_i = \omega + \alpha x_{i-1} + \beta\Psi_{i-1} \tag{2}$$

with the following constraints on the coefficients: $\omega > 0$, $\alpha \geq 0$, $\beta \geq 0$ (but $\beta = 0$ if $\alpha = 0$), and $\alpha + \beta < 1$ (assuming $\mu = 1$). The last constraint ensures the existence of the unconditional mean of the duration, the others ensure the positivity of the conditional durations.

A variant of the ACD model that does not require the positivity of the coefficients is the Log-ACD model of Bauwens and Giot (2000). It is obtained by defining

$$\Psi_i = e^{\psi_i} \tag{3}$$

in (1), and replacing (2) by

$$\psi_i = \omega + \alpha\epsilon_{i-1} + \beta\psi_{i-1}, \tag{4}$$

i.e., the autoregressive equation bears on the *logarithm* of the conditional expectation. The autoregressive structure in Eqs. (2) and (4) allows the model to account for the clustering of the durations. Moments and stationarity conditions for the Log-ACD model are provided by Bauwens et al. (2002). One important condition for stationarity is that $|\beta| < 1$.

In order to estimate the model by the maximum likelihood (ML) method, a parametric hypothesis on the distribution of ϵ_i is necessary. Possible choices, among others, are the Gamma and the Weibull families, which nest the exponential one. Let us assume for example that ϵ_i is distributed as a Weibull(1,γ) random variable, such that $\mu = \Gamma(1 + 1/\gamma)$, where $\Gamma(.)$ is the usual Gamma function. The hazard function of ϵ_i, which is equal to the density divided by the survivor function, is given by

$$h(\epsilon_i) = \gamma\epsilon_i^{\gamma-1}. \tag{5}$$

If $\gamma < 1$, the hazard function is decreasing, if $\gamma > 1$, it is increasing, and if $\gamma = 1$, it is constant (and the distribution is exponential). While (5) gives the hazard function for ϵ_i, it is straightforward to compute the (conditional) hazard function of $x_i = \Psi_i\epsilon_i$, which is given by

$$h(x_i|\mathcal{F}_{i-1}) = \frac{\gamma}{\Psi_i}\,(\frac{x_i}{\Psi_i})^{\gamma-1}. \tag{6}$$

By definition, when Ψ_i is small, the expected duration of x_i is small. In this case, as implied by Eq. (6), the corresponding hazard of x_i is high: the instantaneous probability of witnessing the end of duration x_i is high.

[2] This model is the ACD(1,1). More lags of x_i and Ψ_i can be added. As we use only one lag, we use the short notation ACD.

As such, these duration models are self contained: the durations between two consecutive bid-ask quotes, or more generally between two market events, are modelled conditionally on the past durations. The information given by features linked to the market events is not used. Such observable features may include the traded volume and the corresponding transaction price, the bid and ask prices quoted by the market makers, released news, or functions thereof. Other features, such as the volatility, may be relevant but unobservable.

A natural extension of ACD models consists of combining them with a model bearing on such features of the market. As briefly explained in the introduction, this has been proposed by several authors. At a high level of generality, such models can be described as follows. Let us consider the marked point process (x_i, y_i), where y_i (the marks) is a set of variables (observed or latent) at the end of duration x_i, usually referred to as the 'end state' of duration x_i. The ACD models described so far in this section can be conceived as marginal models of x_i. Durations can be modelled conditionally on a variable in y_i in order to introduce a microstructure effect, as in Engle and Russell (1998), Bauwens and Giot (2000) or Coppejans and Domowitz (1998). The ACD-GARCH model of Engle (2000) is a model for a subset of y_i (the return and its volatility) conditional on x_i (the duration between two market transactions), combined with a marginal ACD model for x_i.

The model introduced in the next section is a joint model for a duration x_i and the direction of the price change y_i (a binary variable) that occurs at the end of x_i. Clearly, different choices of x_i and y_i, combined with different marginalization and conditioning operations can generate a vast range of statistical models. Future research is likely to focus on the specification of such models, and on their relevance for different objectives, such as (ex-post) market analysis, prediction, tests of microstructure hypotheses ...

3. Introducing asymmetry in ACD models

In this section, we introduce a new way of modelling the durations between two bid-ask quotes. The new asymmetric ACD model includes the main features of 'standard' ACD models but also makes use of the information given by the upside or downside movements in the bid-ask price process. Indeed, it turns out that existing ACD or Log-ACD models can be extended into models that combine information on the duration and direction of the price movement. We begin by presenting a simple two-state competing risk model that highlights the feature we want to include in the model, in order to get the asymmetric ACD model, which is then presented (in the second part of this section) and extended to introduce microstructure effects (in the third part).

3.1. A two-state transition model

A multi-state transition model has been applied to data from the Paris Bourse by Bisière and Kamionka (2000), who model the possible transitions between the different types of orders that can be submitted by traders: each type of order (for example a large buy, a small buy, a small sell, ...) represents a given state and the succession of these states is the object for their transition model (which is set up as a competing risk model).

In our framework, we consider the marked point process consisting of the pairs (x_i, y_i), where x_i is the duration between two bid-ask quotes posted by a market maker, and y_i is a variable indicating the direction of change of the mid price defined as the average of the bid and ask prices, i.e.,

- $y_i = 1$ if the mid price increased over duration x_i;
- $y_i = -1$ if the mid price decreased over duration x_i.

At the end of duration x_i, there are only two possible end states:[3] $y_i = 1$ or $y_i = -1$. For simplicity, let us assume that the hazard function is constant, which is equivalent to assume an exponential distribution for x_i. The idea of the transition model is to let the hazard vary with the end state, in this case to have two hazards since y_i is binary. Moreover in each case the hazard can be a function of the previous state y_{i-1} (letting the hazard depend on the previous durations is done in the next subsection). If we define λ_i^+ (respectively λ_i^-) as the hazard of duration x_i when the end state is $y_i = 1$ (respectively -1), then

$$\lambda_i^+ = \beta_1 I_{i-1}^+ + \beta_2 I_{i-1}^- \tag{7}$$

$$\lambda_i^- = \beta_3 I_{i-1}^+ + \beta_4 I_{i-1}^-, \tag{8}$$

where

$$I_{i-1}^+ = \begin{cases} 1 & \text{if } y_{i-1} = 1 \\ 0 & \text{if } y_{i-1} = 0, \end{cases} \tag{9}$$

$$I_{i-1}^- = 1 - I_{i-1}^+. \tag{10}$$

At the end of duration x_i, either state $y_i = 1$ or state $y_i = -1$ is realized. In the framework of a competing risk model, the duration corresponding to the state that is not realized is truncated, since the observed duration is the minimum of two possible durations: the one which would realize if $y_i = 1$, and the one which would realize if $y_i = -1$. The realized state contributes to the likelihood function via its density function, while the truncated state contributes to the likelihood function via its survivor function. Assuming independence (conditionally on the past state) between the durations ending at states $y_i = 1$ and $y_i = -1$, the joint 'density' of duration x_i and state y_i, given the previous state, is equal to

$$f(x_i, y_i | y_{i-1}) = \left(\lambda_i^+\right)^{I_i^+} e^{-\lambda_i^+ x_i} \left(\lambda_i^-\right)^{I_i^-} e^{-\lambda_i^- x_i}. \tag{11}$$

For example, if state $y_i = 1$ is observed ($I_i^+ = 1$ and $I_i^- = 0$), x_i contributes to the likelihood function via the density function $\lambda_i^+ e^{-\lambda_i^+ x_i}$ and via the survivor function $e^{-\lambda_i^- x_i}$. A possible refinement of the model would consist of introducing interdependence between the two 'competing' durations, as suggested in the recent literature in labor econometrics.[4]

---

[3] In principle, a third end state is possible: the state where there is no price change in the mid price over duration x_i. We do not take this state into account as it does not occur in the data we use. The model could easily be extended to a multi-state transition model (see also Subsect. 3.3). Another reason for considering more than two states arises if one wants to account for several levels of price increase or decrease. Notice, however, that a continuous price change is excluded because of the discreteness of price changes.

[4] For example, in Lindeboom and Van den Berg (1994) and in Carling (1996), dependence between durations arises because the hazard rates depend on stochastically related unobserved components.

Table 1. Transition probabilities for two-state transition model (11)

	$y_i = +1$	$y_i = -1$
$y_{i-1} = +1$	$\pi_{1,1} = \frac{\beta_1}{\beta_1+\beta_3}$	$\pi_{1,-1} = \frac{\beta_3}{\beta_1+\beta_3}$
$y_{i-1} = -1$	$\pi_{-1,1} = \frac{\beta_2}{\beta_2+\beta_4}$	$\pi_{-1,-1} = \frac{\beta_4}{\beta_2+\beta_4}$

π_{y_{i-1},y_i} is the transition probability of moving from state y_{i-1} to state y_i. See also Table 3.

Based on the density function given in (11), the transition probabilities for given past and end states can be obtained by integrating this expression with respect to the duration x_i. The results are given in Table 1. The computations are detailed in the Appendix, where more information about the statistical properties of the model are provided. Although this model is too simple to capture the dynamics of the duration and price processes inherent in high-frequency quote data, it puts forward a simple way of combining a duration model and a transition model for the direction of the price process.

3.2. An asymmetric Log-ACD model

Instead of assuming a constant static hazard[5] for the two-state transition process, we can use a non-constant hazard of the Weibull-ACD type, i.e., we let the hazard depend not only on the previous state, but also on the previous durations. Combining a two-state competing risk model with a Log-ACD model yields an asymmetric Log-ACD model, which is defined by the following equations:

– If the end state of duration x_i is $y_i = 1$, the hazard for x_i, assuming a Weibull distribution, is given by

$$h(x_i|y_i = 1, \mathcal{F}_{i-1}) = \frac{\gamma^+}{\Psi_i^+}\left(\frac{x_i}{\Psi_i^+}\right)^{\gamma^+-1} \tag{12}$$

with $\Psi_i^+ = e^{\psi_i^+}$, and the autoregressive process on ψ_i^+ is defined as

$$\psi_i^+ = (\omega_1 + \alpha_1\epsilon_{i-1}^+)I_{i-1}^+ + (\omega_2 + \alpha_2\epsilon_{i-1}^+)I_{i-1}^- + \beta^+\psi_{i-1}^+, \tag{13}$$

with

$$x_i = e^{\psi_i^+}\epsilon_i^+. \tag{14}$$

When $\gamma^+ = 1$, the hazard simplifies to $h(x_i|y_i = 1, \mathcal{F}_{i-1}) = 1/\Psi_i^+$, and thus Ψ_i^+ is similar to the inverse of λ_i^+ of the previous subsection (the difference being that Ψ_i^+ depends on the previous duration and its own lag).

– If the end state of duration x_i is $y_i = -1$, the hazard for x_i is given by

$$h(x_i|y_i = -1, \mathcal{F}_{i-1}) = \frac{\gamma^-}{\Psi_i^-}\left(\frac{x_i}{\Psi_i^-}\right)^{\gamma^--1} \tag{15}$$

with $\Psi_i^- = e^{\psi_i^-}$, and the autoregressive process on ψ_i^- is defined as

$$\psi_i^- = (\omega_3 + \alpha_3\epsilon_{i-1}^-)I_{i-1}^+ + (\omega_4 + \alpha_4\epsilon_{i-1}^-)I_{i-1}^- + \beta^-\psi_{i-1}^-, \tag{16}$$

with

$$x_i = e^{\psi_i^-}\epsilon_i^-. \tag{17}$$

[5] Static with respect to previous durations.

This model exhibits the main feature of the Log-ACD model (i.e., a dynamic structure on the conditional expectation of the duration) while including conditioning information on the underlying bid-ask mid price change. The autoregressive structure of the conditional expectation of the duration process can differ according to the direction (upward or downward) of the mid price. In this framework, the density function of x_i and y_i, given the past information (that now includes past states and durations) is given by

$$f(x_i, y_i|\mathcal{F}_{i-1}) = \left[\frac{\gamma^+}{\Psi_i^+}\left(\frac{x_i}{\Psi_i^+}\right)^{\gamma^+-1}\right]^{I_i^+} e^{-\left(\frac{x_i}{\Psi_i^+}\right)^{\gamma^+}} \left[\frac{\gamma^-}{\Psi_i^-}\left(\frac{x_i}{\Psi_i^-}\right)^{\gamma^--1}\right]^{I_i^-} e^{-\left(\frac{x_i}{\Psi_i^-}\right)^{\gamma^-}} . \tag{18}$$

We can estimate the parameters by maximizing the likelihood which is the product of the densities (18) over all observations, or equivalently, we can estimate the parameters relative to the price increase and decrease separately (by splitting the log-likelihood function into two parts which depend on different parameters).

In the model, we could assume an identical (Weibull) distribution for ϵ^+ and ϵ^-, i.e. that $\gamma^+ = \gamma^-$, a testable hypothesis. In the Appendix, it is shown that this hypothesis implies that the transition probabilities, conditionally on the past durations, do not depend on the current duration x_i, but that otherwise, this conditional independence does not hold.[6] Because Ψ_i^+ and Ψ_i^- depend on the past through the autoregressive process of the conditional durations, the transition probabilities change through time: at the start of each duration x_i, the transition probabilities can be updated based on the information set available at the start of this duration. This is an important feature of the asymmetric ACD model, as it provides evolving forecasts regarding the next price movement. In Sect. 5, we illustrate and use this property using empirical data for two stocks traded on the NYSE.

3.3. Comparison with the ACM model

Russell and Engle (2002) propose a joint model for the duration between two trades and the state of the price change between these trades. Because of the tick rule, price changes can occur only in a small number of values (like 0, $\pm\$1/8$, $\pm\$1/4, \ldots$). They specify a marginal ACD model for the duration x_i, and a conditional model for the discrete valued price change. Furthermore, both parts of the model are conditional on past durations and price changes (and possibly other past information, see also the next subsection). Because the price change is a discrete random variable with a finite number (say k) of possible states, it is convenient to define a $k \times 1$ vector z_i whose j-th component is equal to 1 if the price change is equal to the j-th possible value of the price change, while the other components are equal to 0. To each component is associated a probability of being equal to 1, and the vector z_i is assumed to have a multinomial distribution parameterized by the $k-1$ probabilities

[6] However, when γ^+ and γ^- are not equal, we cannot obtain the conditional (on the past) transition probabilities analytically.

(obviously the last probability is redundant). Russell and Engle (2002) propose to specify convenient transformations of the probabilities as a function of their past history and of the past values of z_i, i.e., ARMA-type specifications. They call their conditional model of z_i given x_i and the past history the autoregressive conditional multinomial (ACM) model.

Obviously, the ACM-ACD model of Russell and Engle (2002) shares several features of the asymmetric ACD model of this paper. Both models specify the joint distribution of the duration and a discrete variable describing the price change. Both are competing risk models (or transition models), which specify the conditional hazard rate $h_j(x|\mathcal{F}_{i-1})$ of a transition from a given state (the previous price change) to the next state., i.e., the instantaneous probability of exiting to state j at time $t_{i-1} + x$ given that the next event has not occurred by $t_{i-1} + x$.

There are two main differences between the models. The first difference concerns the functional form of the hazard functions: in Russell and Engle's model, adapted to the case of two states (price increase and price decrease), the two hazard functions are proportional to each other, such that their ratio is equal to the ratio of the probability of a price increase to the complementary probability of a price decrease. In that respect, our model is more flexible: its hazard functions, given by Eqs. (12) and (15), are not proportional to each other. The second main difference, which is the 'mirror' of the first one, is that the specification of the transition probabilities of the ACM-ACD model is more flexible than in our model: in Eq. (26) (given in the Appendix of the paper), which is our most flexible specification, the transition probabilities (conditional on x_i) are completely determined by the two hazard functions. In the ACM-ACD model, they can be specified as flexibly as one wishes in terms of their functional form (in particular in the way they depend on x_i). These differences reflect the way the model are built: we focus on the duration given the price change, they focus on the price change given the duration.

There are also minor differences. We consider a Log-ACD specification, whereas they consider an ACD specification, but each model can be easily adapted to the alternative specification. Another small difference concerns the number of states (k): we consider two states ($k = 2$) whereas they consider the possibility of more than two states. As explained earlier (see Footnote 3), our model can be easily extended to cover several states.

3.4. Additional explanatory variables

The specification of the conditional expectation of duration x_i as given by Eqs. (13) and (16) does not use possible relevant information that may be included in the information set (as argued at the end of Sect. 2). Indeed, it only includes information on the past duration, the past expected duration and the past state of the price process (increase or decrease). Although this specification is an improvement over the basic ACD model, it can be further extended.

Recently, several authors have proposed to include in such equations explanatory variables drawn from the market microstructure literature, in order to better specify the expected conditional duration. The variables that are taken into account are usually related to the information models originally developed by Glosten and Milgrom (1985). These models assume that the market maker posting the bid and ask quotes faces non informed (or liquidity traders) and informed traders. Because the informed traders know the outcome of the price process when dealing with the market maker, they have an informational advantage and the market maker loses on the trades made with this type of traders. This key feature has led to the development of the so-called information models and to a whole range of predictions for the behaviour of the market maker regarding the way he updates his quotes, increases or decreases his spread, and reacts to the trade process.

With respect to the possible additional variables that can be included in the specification of the autoregressive process, some of the proposals that can be made are:

– The trading intensity: Bauwens and Giot (2000) introduce the lagged trading intensity as an explanatory variable for the expected conditional duration. Over each bid-ask quote duration, the trading intensity is defined as the number of trades divided by the bid-ask duration. The model of Easley and O'Hara (1992) implies that an increase in the trading intensity (due for example to a release of news) should lead to more frequent revisions of the quotes and thus to a shortening of the durations between the posting of bid and ask quotes by the market maker. This should lead to a negative coefficient for the trading intensity when introduced in the specification of the expected conditional duration.
– The volume: introducing a link between the price process and the volume is one of the main features of technical analysis, an empirical tool much favored by market practitioners. This issue is also closely studied by Easley and O'Hara (1992), and by Blume et al. (1994). In the Easley and O'Hara (1992) paper, the argument presented regarding the influence of the spread (see below) and trading intensity is extended to the volume. More precisely, it is the unexpected volume that matters most.[7] To capture the influence of volume on the bid-ask quote revision process, we use the average volume per trade. As it is the unexpected component of the variable that matters, the variable included in the model is the time-of-day adjusted average volume per trade (see Sect. 4 for the definition of the variable with respect to the time-of-day effect). If unexpected volume is related to informed trading, its increase should shorten the next expected duration: a negative coefficient is expected for this variable when included in the specification of the model.
– The spread: Easley and O'Hara (1992) extend the information model introduced by Glosten and Milgrom (1985) by focusing on the role of time in price adjustment. Easley and O'Hara argue that the duration between two trades conveys information. In their model, a no-trade outcome (a long duration) means that no new information has been released. Thus the probability of dealing with an informed trader is small, relative to the case where the duration would be small. Consequently, with a low probability of dealing with an informed trader, the market maker decreases his bid-ask

[7] The unexpected volume is defined as the volume in excess of the normal volume or liquidity volume. It is thus the 'deviation' of the actual trading volume from the time-of-day adjusted volume. The source of the unexpected volume is the arrival of informed traders in the market.

spread. This should lead to a negative coefficient for the lagged spread when introduced in the specification of the expected conditional duration. Note that this argument should also extend to the case of price durations (as studied in this paper, see Sect. 4.1 for a description of price durations). Indeed, shorter price durations mean increased volatility (see Giot 2000) and Jones et al. (1994) suggest that increased trading activity (i.e., small durations between trades) is a main determinant of volatility. Therefore, it is expected that price durations should qualitatively exhibit the same dependence on the lagged spread as durations between trades.

By adding these additional explanatory variables, the specification for the asymmetric Log-ACD model becomes

$$\psi_i^+ = (\omega_1 + \alpha_1 \epsilon_{i-1}^+) I_{i-1}^+ + (\omega_2 + \alpha_2 \epsilon_{i-1}^+) I_{i-1}^- + \beta^+ \psi_{i-1}^+ + \delta_j^+ E_{i-1} \tag{19}$$

for the end state $y_i = +1$, where δ_j^+ is the coefficient of the additional variable E_{i-1}, which can be the past trading intensity ($tint_{i-1}$ and δ_1^+), the past average volume per trade (vol_{i-1} and δ_2^+), and the past spread (sp_{i-1} and δ_3^+). The specification for the end state $y_i = -1$ is

$$\psi_i^- = (\omega_3 + \alpha_3 \epsilon_{i-1}^-) I_{i-1}^+ + (\omega_4 + \alpha_4 \epsilon_{i-1}^-) I_{i-1}^- + \beta^- \psi_{i-1}^- + \delta_j^- E_{i-1}, \tag{20}$$

where δ_j^- is the coefficient of the additional variable E_{i-1}. As explained above, according to the market microstructure information models, all six coefficients δ_j^+ and δ_j^- should be negative. These hypotheses are investigated in the next section, where we apply the asymmetric Log-ACD model to actual bid-ask quote data for two stocks traded on the NYSE.

4. Application to stocks traded on the NYSE

Using data over a three-month period for Disney and IBM, two stocks traded on the NYSE,[8] we estimate the asymmetric Log-ACD model. We start by a description of the data, and we proceed with a discussion of the results.

4.1. Data description

The data were extracted from the Trades and Quotes (TAQ) database made available by the NYSE. This database consists of two parts: the first reports all trades, while the second lists the bid and ask prices quoted by the specialists (at the NYSE). Some information about the data is given in Table 2.

As in other papers (Engle and Russell 1998; Bauwens and Giot 2000; Bauwens and Veredas 2002; Giot 2000), we estimate the model on price durations because we are interested in significant price changes and the corresponding durations implied by these price changes. Price durations are defined by thinning the bid-ask quote durations using the mid price, i.e., we compute the durations leading to a significant cumulated change in the mid price of the specialist's quote. A significant change in the mid price of the

[8] We did not take into account the data from the other regional stock exchanges where these stocks are traded too.

Table 2. Information on the data

	Disney	IBM
Number of trades	33,146	61,063
Number of b/a price durations	2,160	6,728
Mean of x_i and X_i	0.98–646.98	0.99–215.92
Standard deviation of x_i and X_i	1.25–966.34	1.50–365.46
Minimum value x_i and X_i	0.0050–3	0.0030–1
Maximum value x_i and X_i	12.29–9739	35.82–7170

Data extracted from the September, October and November 1996 TAQ database. The number of bid-ask price durations is the number obtained after thinning the data (the original number of bid-ask quotes is equal to 37,325 for Disney and 34,321 for IBM). x_i is the time-of-day adjusted duration, see (21), and is measured in seconds. X_i is the non-adjusted price duration.

specialist is defined as a change leading to at least a \$0.125 cumulated change in the mid price. Thus we did not take into account the numerous \$0.0625 changes in the mid price, which are due to a \$0.125 price change of the bid or of the ask. Of course, two successive \$0.0625 changes in the same direction yield a cumulative \$0.125 change (and thus lead to a retained duration). The thinning can be justified by the presumption that the \$0.0625 changes are transitory, i.e., are mainly due to the short term component of the bid-ask quote updating process. Indeed, Biais et al. (1995) provide evidence that information effects in the order process lead to similar (successive) changes in quotes on both sides of the market (i.e., information events quickly lead to movements of the bid and ask quotes in the same direction). More information on price durations is given in Bauwens and Giot (2001).

For the IBM data, the number of original bid-ask quotes (34,321) is approximately half the number of trades (61,063). After thinning, the number of price durations is 6,728. For the Disney data, the number of original bid-ask quotes is equal to 37,325. It is reduced to 2,160 by the thinning. The number of trades is equal to 33,146.

As explained by Engle and Russell (1998), it is necessary to transform the raw durations before estimating the ACD model. Indeed these durations consist of two parts: a stochastic component (which is explained by the ACD model) and a deterministic part, which is called a ‘time-of-day’ effect. This time-of-day effect arises from the systematic variations of the quote arrivals over the trading day and mainly depends on the institutional features of the exchange and the behavior of traders and market makers (Bauwens and Giot 2001). This effect should be removed from the raw durations before estimating the model. We follow Engle and Russell (1998) in defining this deterministic effect as a multiplicative component, i.e.,

$$X_i = x_i \phi(t_i) \tag{21}$$

where X_i is the raw duration, $\phi(t_i)$ is the time-of-day function, and x_i denotes the time-of-day (tod) adjusted duration. Two methods have been used in previous studies to compute the time-of-day function: spline smoothing and kernel smoothing. In the first case (Engle and Russell 1998), the tod function is defined as the expected duration conditioned on time-of-day, the expectation being computed by averaging the durations over thirty-minute intervals. Cubic splines are then used to smooth the time-of-day function on the

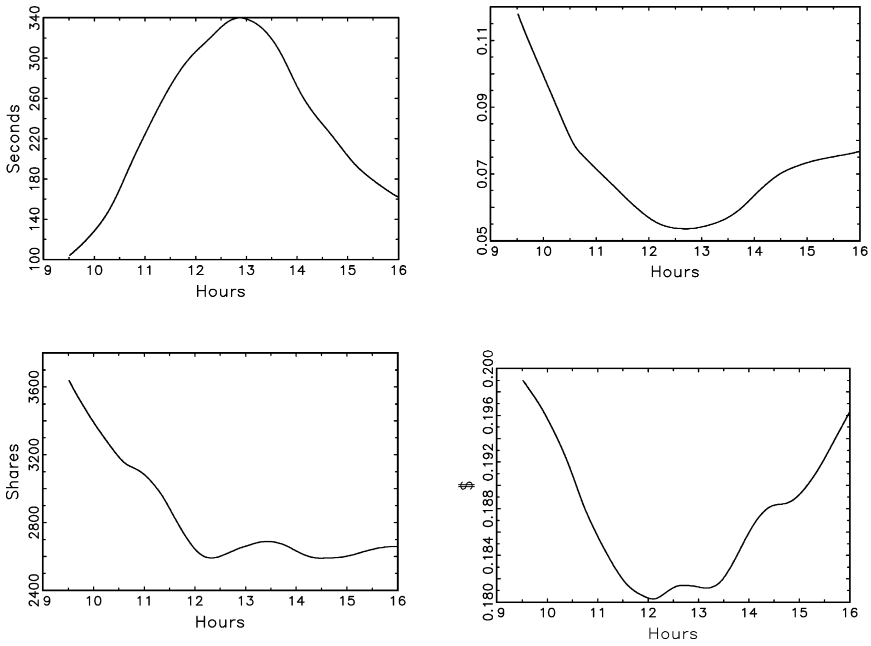

Fig. 1. Time-of-day pattern for the price durations and the market microstructure variables (IBM). From *left* to *right* and *up* to *down*: Price durations, trading intensity, average volume per trade and spread

thirty-minute intervals.[9] In the second case, a non-parametric regression of the raw durations on the time of day is used. Veredas et al. (2001) advocate the use of a gamma kernel with the Nadaraya-Watson estimator. This is the method used in this paper.[10]

Because the trading intensity $tint_i$ and the average volume per trade vol_i depend on the number of trades and duration X_i, they also feature a time-of-day effect. As far as the spread is concerned, it is well known that it also displays a strong intraday seasonality (see for example Chung et al. 1999). The time-of-day function for the price durations and the three market microstructure variables is shown in Fig. 1 (for the IBM stock). For all those variables, we removed the time-of-day effect by the same method as for X_i. The adjusted trading intensity (average volume per trade, spread) can then be interpreted as measuring the unexpected component of the trading intensity (average volume per trade, spread), which can be larger or smaller than the 'normal' level of a particular time. Finally, in order to illustrate the price transitions, we compute the empirical transition probabilities for the mid price changes associated with the price durations. The results, given in Table 3, show that the succession of similar states (either increase or decrease) is more likely

[9] An alternative method is to estimate jointly the parameters of the duration model and the tod function; Engle and Russell (1998) indicate that this method gives almost equivalent results to those of the simpler two step method.

[10] We have also used the spline method. Our empirical results and conclusions are not qualitatively different according to the method used.

than price reversals (increase followed by decrease or vice-versa) for IBM, but not for Disney (for which a price decrease was observed on the sample period). The figures in Table 3 provide estimates of the probabilities in Table 1 and can obviously be transformed into estimates of the parameters β_i ($i = 1, \ldots 4$) of the static transition model of Subsect. 3.1. All the reported probabilities are significantly different from 0.5, except in the case of Disney for a transition from +1 to +1 (and of course to −1).

4.2. Estimation results

Estimates of the asymmetric Log-ACD model are reported in Table 4 for the IBM and Disney stocks. The estimation has been performed by maximizing

Table 3. Empirical transition probabilities

	$y_i = +1$	$y_i = -1$	z-statistic	n
$y_{i-1} = +1$	0.575 (0.512)	0.425 (0.488)	8.93 (0.81)	3462 (1148)
$y_{i-1} = -1$	0.451 (0.553)	0.549 (0.447)	−5.63 (3.13)	3265 (1011)

$y_i = +1$ corresponds to an increase of the midprice from t_{i-1} to t_i, $y_i = -1$ corresponds to an decrease. Numbers in parentheses are for Disney, others are for IBM. The z-statistics in the last column are the asymptotic $N(0, 1)$ statistics for the hypothesis that each probability is equal to 0.5. The last column gives the number of observations such that $y_{i-1} = +1$ (first row) and $y_{i-1} = -1$ (last row).

Table 4. ML results for the asymmetric Log-ACD model

	End state: price increase			End state: price increase	
	IBM	Disney		IBM	Disney
Without market microstructure variables					
ω_1	−0.103 (0.018)	−0.069 (0.019)	ω_3	−0.077 (0.013)	−0.115 (0.021)
ω_2	−0.074 (0.015)	−0.017 (0.024)	ω_4	−0.100 (0.016)	−0.004 (0.022)
α_1	0.139 (0.031)	0.101 (0.036)	α_3	0.218 (0.033)	0.202 (0.051)
α_2	0.257 (0.043)	0.128 (0.041)	α_4	0.182 (0.032)	0.104 (0.033)
β^+	0.976 (0.009)	0.974 (0.013)	β^-	0.984 (0.005)	0.988 (0.007)
γ^+	0.988 (0.012)	0.929 (0.018)	γ^-	0.990 (0.012)	0.973 (0.020)
With the lagged trading intensity					
δ_1^+	−0.056 (0.012)	−0.047 (0.021)	δ_1^-	−0.050 (0.012)	−0.047 (0.019)
With the lagged average volume per trade					
δ_2^+	−0.013 (0.007)	−0.012 (0.010)	δ_2^-	−0.027 (0.007)	−0.030 (0.013)
With the lagged spread					
δ_3^+	−0.076 (0.037)	−0.127 (0.049)	δ_3^-	−0.094 (0.030)	−0.083 (0.036)
With all market microstructure variables included simultaneously					
δ_1^+	−0.056 (0.012)	−0.038 (0.022)	δ_1^-	−0.047 (0.013)	−0.039 (0.020)
δ_2^+	−0.009 (0.008)	−0.002 (0.012)	δ_2^-	−0.022 (0.008)	−0.026 (0.014)
δ_3^+	−0.053 (0.036)	−0.112 (0.059)	δ_3^-	−0.059 (0.029)	−0.014 (0.048)

For a definition of the model, see Eqs. (14) with (13) or (19) for the price increase, and Eqs. (17) with (16) or (20) for the price decrease. δ_1^+ and δ_1^- are the coefficients of the lagged trading intensity, δ_2^+ and δ_2^- are the coefficients of the lagged average volume per trade, and δ_3^+ and δ_3^- are the coefficients of the lagged spread. QML asymptotic standard errors are given in parentheses.

the likelihood function defined in Sect. 3, (using the BHHH algorithm implemented in GAUSS). The first three columns of Table 4 are for the part of the model when the end state is a price increase – see Eqs. (13) and (19) – and the last three columns are for the part when the end state is a price decrease – see (16) and (20). In the upper block of the table, we present the results for Eqs. (13) and (16). In the lower blocks, we report only the results for the coefficients of the market microstructure variables when they are included in the specification of the model (first separately, and then simultaneously). Standard errors are computed using the QML 'sandwich' formula. The main features of these results are:

1) Strong autoregressive effects (β coefficients) for all specifications and for both stocks, indicating a strong persistence in the dynamics of the duration process. This result is similar to those of Bauwens and Giot (2000) for the symmetric Log-ACD model and of Engle and Russell (1998) for the ACD model. The inclusion of the explanatory variables hardly changes the estimates of the β coefficients.
2) The estimates of the γ parameters are very close to 1, indicating that exponential distributions for the innovations are acceptable for the data we use. Moreover, the hypotheses that $\gamma^+ = \gamma^-$ and $\gamma^+ = \gamma^- = 1$ are not rejected at conventional levels of significance for the IBM stock, while the hypotheses that $\gamma^+ = \gamma^-$ and $\gamma^- = 1$ are not rejected for the Disney stock (the hypothesis that $\gamma^+ = 1$ is rejected for this stock).
3) When the end state is a price increase ($y_i = 1$), α_1 is significantly smaller than α_2 for the IBM stock: the expected (conditional) duration of witnessing a transition from state $+1$ to state $+1$ is smaller than the expected duration of witnessing a transition from state -1 to state $+1$. For Disney, the two coefficients are close. When the end state is a price decrease ($y_i = -1$), α_3 is larger than α_4 for both stocks although the difference is not significant for IBM; the expected duration of witnessing a transition from state -1 to state -1 is smaller than the expected duration of witnessing a transition from state $+1$ to state -1.
4) The estimates of δ_1^+ and δ_1^- are negative and significant for both stocks: an increase of the trading intensity (defined relatively to its normal level measured by the time-of-day function) decreases the next expected duration between the bid-ask quotes, both for the exit at state $+1$ and the exit at state -1.
5) The estimates of δ_2^+ and δ_2^- are negative, with δ_2^- being significant for both stocks: when the average volume per trade (defined relatively as above) increases, the next expected duration becomes smaller.
6) The estimates of δ_3^+ and δ_3^- are also significantly negative for both stocks. Once more, these coefficients have the expected sign: an increase of the spread (defined relatively as above) decreases the next expected duration between the bid-ask quotes. This is in agreement with the Easley and O'Hara (1992) model as explained earlier in the paper.

The estimation results obtained after including the additional variables altogether are given in the lower part of Table 4. Not surprisingly, the corresponding coefficients are all negative, although some of them are not significant. Among the three additional variables, the trading intensity seems to be the most relevant, as δ_1^+ and δ_1^- are significant in all cases and for both stocks. It is interesting to note that these results are quite similar to those of Jones et al. (1994). Indeed, in their study (using daily data) on a large sample of stock traded on the NASDAQ, they found that it is the number of trades, and not the average volume per trade, that has the most important impact on the volatility of the bid-ask quote process. With respect to the estimation of the model with the microstructure variables, we conclude that the empirical evidence in favour of the information model of Easley and O'Hara (1992) is rather convincing. In all cases, the coefficients are negative, supporting the idea that an assumed release of information decreases the next expected duration of the bid-ask quotes.

Furthermore, the results given in Table 4 indicate that, for each additional variable included in the model, the corresponding coefficients for the two states (price increase and price decrease) are quite close. Testing $\delta_1^+ = \delta_1^-$ (trading intensity), $\delta_2^+ = \delta_2^-$ (average volume per trade), and δ_3^+ and δ_3^- (spread) separately by a t-test for the IBM (Disney) stock, we do not not reject the null hypotheses since the statistics are equal to −0.55 (−0.01), 1.60 (1.29), and 0.40 (−0.73), respectively. Moreover, testing the three hypotheses jointly by a likelihood ratio test does not reject the null hypothesis either.

It is also possible to test the asymmetric Log-ACD model (without the additional explanatory variables) against the symmetric version of Bauwens and Giot (2000). The asymmetric model introduced in Sect. 3.2 can be constrained into a symmetric one by setting $\omega_1 = \omega_2 = \omega_3 = \omega_4$, $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4$, $\beta^+ = \beta^-$, and $\gamma^+ = \gamma^-$, so that $\psi_i^+ = \psi_i^-$. The constrained model is then equivalent to the original Log-ACD model of Bauwens and Giot (2000). The likelihood ratio test statistic, distributed as $\chi^2(8)$, is equal to 31.9 for the Disney stock and to 66 for the IBM stock: both values are highly significant. The symmetric model is strongly rejected against the asymmetric one.

5. Forecasting and trading with the model

As detailed in Sect. 3, an interesting property of the asymmetric Log-ACD model is to provide two hazards (one for the exit at a price increase, the other for the exit at a price decrease) that change over time as they depend on the previous state and durations. This feature implies a time-varying transition probability matrix for the mid price. In this section, we use this feature to assess the forecasting performance of the model on the available intraday data. We also combine the 'predictions' of the model regarding the next price movement with a possible trading strategy in order to present a possible empirical use of the model.

5.1. Forecasting intraday price changes

The sequence of time-varying transition probability matrices is based on the estimates of the asymmetric Log-ACD model imposing that $\gamma^+ = \gamma^- = 1$. As we have seen in Sect. 4, this restriction is far from being rejected by our data. We impose it because it allows to compute analytically the transition

Table 5. Forecasting and trading results

Threshold	Percentage of correct decisions	Gain/loss on portfolio	
	AACD model	Managed	Unmanaged
IBM:			
\$1/8	53.90	−212,031	1,284,944
\$1/4	55.23	−52,074	634,561
\$1/2	59.37	334,370	292,483
Disney:			
\$1/8	54.28	2,017,417	1,006,596
\$1/4	59.24	788,044	491,740

A correct decision is a correctly forecasted direction of price change. 'AACD Model' is the asymmetric Log-ACD model. Resulting gain/loss after the three-month period is in \$. 'Managed' means that the portfolio is managed using the model, and 'unmanaged' means it follows the buy/hold strategy.

probabilities (unconditional with respect to the current duration x_i). Using the constrained estimates,[11] at the start of each price duration, we compute the (conditional) probability of a price increase. If this probability is greater (lower) than 1/2, we predict a price increase (decrease). At the end of each duration, the end state is observed and the realized outcome is compared to the predicted outcome. By repeating this procedure for all available durations, we can compute the percentage of correct decisions implied by the model. For the two stocks and for several possible price thresholds,[12] the results are given in column 2 of Table 5. For the IBM stock and for a threshold of \$1/4, we show in Fig. 1 the evolution of the probability of a price increase, i.e. the conditional probability $P(y_i = 1|\mathcal{F}_{i-1})$, given by formula (6) specialized to the case $\gamma^+ = \gamma^- = 1$. When this probability is larger than 1/2, the next price movement is forecasted to be an increase. The dotted line in Fig. 2 gives the realized states at the end of the price duration. As can be seen on Fig. 2 and as expected from the structure of the model, the conditional probability $P(y_i = 1|\mathcal{F}_{i-1})$ is strongly influenced by the previous realized states (price movements). As indicated in Table 5, the model delivers a percentage of correct decisions greater than 50%, for both stocks and for all thresholds. Before drawing conclusions regarding the forecasting performance of the model, we conduct an evaluation based on trading strategies.

5.2. A trading strategy

We now combine the computed transition probabilities with a possible trading strategy. We stress however that this strategy is only meant to be illustrative. In a 'real life' situation, other parameters have to be taken into account: risk analysis of the positions, limitations on overnight positions,

[11] They are not reported because they are almost identical to the estimates in Sect. 4.

[12] We have estimated the model for each set of price durations corresponding to a different threshold of the price change used to thin the quote process. Estimates are not reported since they are qualitatively similar to those reported in Sect. 4.

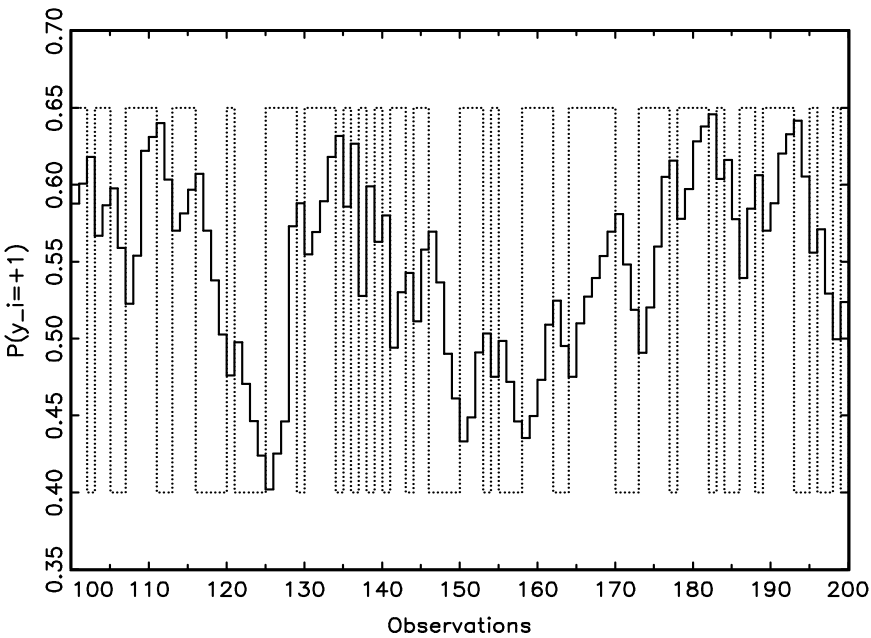

Fig. 2. Conditional Probability of a Price Increase $P(y_i = 1|\mathcal{F}_{i-1})$. The *solid* line shows the evolution of the probability of a price increase computed from the asymmetric Log-ACD model for IBM price durations (threshold at \$1/4). The *dotted* line shows the evolution of the realized price, with the value 0.40 indicating a decrease and the value 0.65 an increase

limited amount of short-selling or limit positions in trading. Moreover, the literature on technical analysis and corresponding trading strategies would suggest many more possibilities regarding the choice of additional explanatory variables that could be included in the specifications of the hazards: trading signals based on past volumes, trading signals based on the 'distance' between the last recorded price and moving averages, ... In our case, we define the trading rules as follows:

– If a price increase is forecasted, n shares of the stock are bought at the prevailing ask price and are added to the existing position; if the portfolio has an existing short position, it is settled at the prevailing ask price, prior to buying the n shares;
– If a price decrease is forecasted, n shares are short-sold at the prevailing bid price and are added to the existing short position; if the portfolio has an existing long position, it is settled at the prevailing bid price, prior to short-selling the n shares.

We choose $n = 1,000$ with a transaction cost equal to \$10. The first decision is taken at the start of the second duration and the game is ended at the end of the three-month period. The portfolio is self-financed, i.e. it borrows money (at a 10% interest rate) if shares are to be bought. When shares are short-sold or when a trading profit is realized, the amount of money received earns interests at a 10% rate.

For the two stocks and for several thresholds, the outcome of this trading strategy (which we call the managed portfolio) after a three-month period is given in column 3 of Table 5. Generally speaking, the results are quite in line with those given in column 2: because the percentage of correct decisions is higher than 50%, the model ends up with a gain in most cases. When it does not, the loss seems to be due to very large transaction costs because of the frequent trading imposed by the model (as for IBM and price durations at \$1/8 where 6,728 trades have to be made).

While this trading strategy seems to be profitable in some cases (even after taken into account the transaction costs, the financing costs and the losses due to the bid-ask bounce), it is worthwhile to compare it with a benchmark strategy, which we take to be the buy/hold strategy in this case. The buy/hold strategy is one of the standard benchmark models in finance when actively managed portfolio are evaluated; see for example Sharpe et al. (1999) for a general discussion regarding the performance of actively managed portfolio. Under the buy/hold strategy (which is equivalent to tracking an unmanaged portfolio), n shares are bought at the start of the three-month period and they are sold at the end, resulting in a profit or loss. To conduct a fair comparison between the two strategies, n is chosen such that the financing cost of the buy/hold strategy is equal to the total cost of the managed portfolio, i.e., the available amount of money needed at the start of the three-month period to buy the shares is equal to the amount that can be borrowed for an interest cost (at a 10% interest rate) equal to the total cost of the managed portfolio. As indicated in Table 5, the two trading strategies deliver a gain after the three-month period in three cases out of five. However, the results are less good for the actively managed IBM portfolio than for the unmanaged one. For the Disney stock, the active strategy based on the model outperforms the buy/hold strategy. Note that we do not take into account the risk of the two strategies as we only compare the nominal amounts gained or lost by the models at the end of the trading game.[13]

5.3. Some comments

In the two preceding subsections, we have presented a possible practical application of the asymmetric Log-ACD model. However, we want to stress that the results are highly dependent on the stock we study and the period under review. In the finance literature dealing with the assessment of the performance of trading strategies and actively managed portfolio, studies are usually conducted on a large sample of stocks and for a much longer timespan. Thus we view these results as an interesting illustration of what the model can do, but more work would be needed to evaluate the performance and riskiness of the model by considering a wide variety of cases. Furthermore, the performance of an actively managed portfolio strongly depends on the transaction costs involved in the trading of the shares. Regarding the Disney stock for example, if the transaction costs are increased from \$10 to \$100 per trade, then the buy/hold strategy becomes more profitable than the one based on the asymmetric Log-ACD model. Hence this kind of trading strategy should be applied to intraday stock futures data as the trading costs and bid-ask spread

[13] See also the comments in the next subsection.

are much lower for futures contracts than for individual stocks. In addition, active trading strategies are usually characterized by a high level of risk compared to passive investment strategies as end-of-period outcomes are usually more variable and therefore exhibit a larger variance than outcomes from passive investment strategies. Therefore, extensions of our study should focus on the performance of the model in a risk-adjusted framework.

Finally, the question may be raised whether our results contradict the efficient market hypothesis, since it seems possible to predict prices at the intraday level. This is not the first study that concludes in this way, see Alexander (2001, p. 389) for references. In that respect, we have to be careful. Testing of the efficient market hypothesis at the intraday level and over a relatively short time horizon is an issue that has to be investigated in much more detail. At the theoretical level, the hypothesis, which says (loosely speaking) that prices incorporate the relevant information, is not in contradiction with the statement that it takes some time for the information to be incorporated in prices. If the adjustment is not instantaneous, prices may be predictable in the short run.

6. Conclusion

ACD models seem to be a useful and promising tool not only to model durations arising from irregular events occurring on financial markets, but also as a building block of larger models bearing on features of the market events, such as price, volume, volatility. This paper has proposed a two-state transition model for the intraday price of a stock, where the two states correspond either to a price increase or to a price decrease, and where the durations are modelled by a Log-ACD process. Our empirical results favour the asymmetric specification as an enrichment of the symmetric one. Three additional explanatory variables drawn from the market microstructure literature related to information models were included in the model: the past spread, the past trading intensity and the past average volume per trade. The coefficients of these variables were negative and most were significant. This provides support to the information models and explains the quick revision of bid-ask prices when information has been released. As a possible practical application, we conducted a study bearing on the forecasting performance of the model and the resulting gain/loss of a trading strategy making use of the predictions of the model.

Extensions of this model could consist of considering a more refined state space (more than two states), other distributions than the Weibull for the innovations (possibly a semi-parametric specification), more refined parameterizations of the autoregressive equations for the conditional durations, and a more thorough investigation of the performance of the model on a large sample of stocks and over a much larger timespan. More fundamentally, theoretical foundations for the determinants of the asymmetry revealed by our empirical results should be uncovered.

Appendix

We consider the joint density of the static two-state exponential transition model set in the competing risk framework, i.e. the function defined in (11).

Several features of the model can be highlighted. First, this density integrates to one, as it should:

$$\sum_{y_i=-1,1} \int_0^\infty f(x_i, y_i|y_{i-1})dx_i = \int_0^\infty \lambda_i^+ e^{-(\lambda_i^+ + \lambda_i^-)x_i} dx_i + \int_0^\infty \lambda_i^- e^{-(\lambda_i^+ + \lambda_i^-)x_i} dx_i = 1.$$

Integrating x_i in (11) yields the transition probabilities

$$f(y_i|y_{i-1}) = \left(\lambda_i^+\right)^{I_i^+} \left(\lambda_i^-\right)^{I_i^-} \int_0^\infty e^{-(\lambda_i^+ + \lambda_i^-)x_i} dx_i = \frac{\left(\lambda_i^+\right)^{I_i^+} \left(\lambda_i^-\right)^{I_i^-}}{\lambda_i^+ + \lambda_i^-}, \tag{22}$$

so that $P(y_i = 1|y_{i-1}) = \lambda_i^+/(\lambda_i^+ + \lambda_i^-)$ and $P(y_i = -1|y_{i-1}) = \lambda_i^-/(\lambda_i^+ + \lambda_i^-)$, which are the results given in Table 1. Furthermore, as

$$f(y_i|x_i, y_{i-1}) = \frac{f(x_i, y_i|y_{i-1})}{f(x_i|y_{i-1})} \tag{23}$$

and

$$f(x_i|y_{i-1}) = \sum_{y_i=-1,1} f(x_i, y_i|y_{i-1}) = (\lambda_i^+ + \lambda_i^-)e^{-(\lambda_i^+ + \lambda_i^-)x_i} \tag{24}$$

(an exponential distribution), we have that

$$f(y_i|x_i, y_{i-1}) = \frac{\left(\lambda_i^+\right)^{I_i^+} \left(\lambda_i^-\right)^{I_i^-}}{\lambda_i^+ + \lambda_i^-}, \tag{25}$$

which does not depend on x_i. Thus y_i is independent of x_i, conditionally on y_{i-1}.

When the model is extended to the Weibull-ACD case, the joint density $f(x_i, y_i|\mathcal{F}_{i-1})$ (where $\mathcal{F}_{i-1}$ includes the past states and durations) is given by (18). Similar computations as above give

$$f(y_i|x_i, \mathcal{F}_{i-1}) = \frac{\left[\frac{\gamma^+}{\Psi_i^+}\left(\frac{x_i}{\Psi_i^+}\right)^{\gamma^+-1}\right]^{I_i^+} \left[\frac{\gamma^-}{\Psi_i^-}\left(\frac{x_i}{\Psi_i^-}\right)^{\gamma^--1}\right]^{I_i^-}}{\frac{\gamma^+}{\Psi_i^+}\left(\frac{x_i}{\Psi_i^+}\right)^{\gamma^+-1} + \frac{\gamma^-}{\Psi_i^-}\left(\frac{x_i}{\Psi_i^-}\right)^{\gamma^--1}}, \tag{26}$$

which implies that generally y_i is not independent of x_i, given the previous state and all the previous durations. However, it is obvious that when $\gamma^+ = \gamma^-$, $f(y_i|x_i, \mathcal{F}_{i-1})$ no longer depends on x_i, meaning that in that case y_i is (conditionally) independent of x_i. In the empirical applications of this paper, we accept the hypothesis that $\gamma^+ = \gamma^-$, and we can even accept that $\gamma^+ = \gamma^- = 1$ (see Section 4). Moreover it can be shown that when $\gamma^+ = \gamma^- = \gamma$,

$$f(x_i|y_{i-1}) = [(\Psi_i^+)^{-\gamma} + (\Psi_i^-)^{-\gamma}]\gamma x_i^{\gamma-1} e^{-x_i^\gamma[(\Psi_i^+)^{-\gamma} + (\Psi_i^-)^{-\gamma}]}, \tag{27}$$

i.e., a Weibull distribution with parameters γ and $[(\Psi_i^+)^{-\gamma} + (\Psi_i^-)^{-\gamma}]^{1/\gamma}$.

References

Alexander C (2001) Market models. Wiley, New York
Bauwens L, Giot P (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, Giot P (2001) Econometric modelling of stock market intraday activity. Kluwer Academic Publishers, Dordrecht
Bauwens L, Veredas D (2002) The stochastic conditional duration model: a latent factor model for the analysis of financial durations. Journal of Econometrics (forthcoming)
Bauwens L, Galli F, Giot P (2002) The moments of Log-ACD models. Mimeo. CORE, Université catholique de Louvain
Biais B, Hillion P, Spatt C (1995) An empirical analysis of the limit order book and the order flow in the Paris Bourse. Journal of Finance 50: 1655–1689
Bisière C, Kamionka T (2000) Timing of orders, orders aggressiveness and the order book at the Paris Bourse. Annales d'Economie et de Statistique 60: 43–72
Blume L, Easley D, O'Hara M (1994) Market statistics and technical analysis: the role of volume. Journal of Finance 49: 153–181
Chung KH, Van Ness BF, Van Ness RA (1999) Limit orders and the bid-ask spread. Journal of Financial Economics 53: 255–287
Carling K (1996). Testing for independence in a competing risk model. Computational Statistics and Data Analysis 22: 527–535
Coppejans M, Domowitz I (1998) Stock and flow information as inputs to limit order book trading activity. Mimeo. Presented at the Conference on Microstructure and High Frequency Data in Paris, December 1998
Darolles S, Gouriéroux C, Le Fol G (2000) Intraday transaction price dynamics. Annales d'Economie et de Statistique 60: 207–238
Easley D, O'Hara M (1992) Time and the process of security price adjustment. Journal of Finance 19: 69–90
Easley D, Kiefer NM, O'Hara M (1997) The information content of the trading process. Journal of Empirical Finance 159–186
Engle RF (2000) The econometrics of ultra-high frequency data. Econometrica 68: 1–22
Engle R, Lunde A (1998) Trades and quotes. A bivariate point process. Discussion Paper 98-07. University of California, San Diego
Engle R, Russell J (1997) Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model. Journal of Empirical Finance 4: 187–212
Engle RF, Russell JR (1998) Autoregressive conditional duration: a new approach for irregularly spaced transaction data. Econometrica 66: 1127–1162
Ghysels E, Gouriéroux C, Jasiak J (1997) Stochastic volatility duration models. Working Paper 9746. CREST, Paris
Ghysels E, Jasiak J (1998) GARCH for irregularly spaced financial data: the ACD-GARCH model. Studies in Nonlinear Dynamics and Econometrics. 2: 133–149
Giot P (2000) Time transformations, intraday data, and volatility models. Journal of Computational Finance 4: 31–62
Glosten L, Milgrom P (1985) Bid, ask, and the transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 13: 71–100
Grammig J, Maurer K-O (2000) Non-monotonic hazard functions and the autoregressive conditional duration model. The Econometrics Journal 3: 16–38
Grammig J, Wellner M (2002) Modeling the interdependence of volatility and inter-transaction duration processes. Journal of Econometrics 106: 369–400
Hafner C (1999) Durations, volume and the prediction of financial returns in transaction time. Mimeo. Presented at the Econometric Society World Congress in Seattle, August 2000
Jones CM, Kaul G, Lipson ML (1994) Transactions, volume, and volatility. Review of Financial Studies 7: 631–651
Lindeboom M, Van Den Berg GJ (1994) Heterogeneity in models for bivariate survival: the importance of the mixing distribution. Journal of the Royal Statistical Society B 56: 49–60
Lunde A (1997) A conjugate gamma model for durations in transactions. Working Paper. Department of Economics, University of Aarhus

Lunde A (1998) A generalized gamma autoregressive conditional duration model. Working Paper. Department of Economics, University of Aarhus

Meddahi N, Renault E, Werker B (1998) Continuous time models for high frequency data. Proceedings of the Second High-Frequency Data in Finance Conference, Olsen & Associates, Zurich

O'Hara M (1995) Market microstructure theory. Basil Blackwell

Prigent JL, Renault O, Scaillet O (2001) An autoregressive conditional binomial option pricing model. In: Geman, Madan, Pliska, Vorst (eds.), Selected papers from the first world congress of the Bachelier Finance Society. Springer, Berlin Heidelberg New York

Russell J, Engle R (2002) Econometric analysis of discrete-valued, irregularly-spaced financial transactions data. Revised version of Discussion Paper 98-10, University of California, San Diego

Sharpe WF, Alexander GJ, Bailey JV (1999) Investments, 6th ed. Prentice-Hall

Veredas D, Rodriguez-Poo J, Espasa A (2001) On the (intradaily) seasonality of a financial point process. WP 01-33 Statistics and Econometrics Series. University Carlos III de Madrid