Ordinal-response models for irregularly spaced transactions: A forecasting exercise

Dimitrakopoulos, Stefanos and Tsionas, Mike G. and Aknouche, Abdelhakim

Leeds University, Lancaster University, Qassim University and USTHB

1 October 2020

Ordinal-response models for irregularly spaced transactions: A forecasting exercise

Stefanos Dimitrakopoulos[1*], Mike Tsionas[**], and Abdelhakim Aknouche[***]

[*]Economics Division, Leeds University, UK

[**]Department of Economics, Lancaster University, UK

[***]Department of Mathematics, College of Science, Qassim University (Saudi Arabia) and Faculty of Mathematics, USTHB (Algeria)

Abstract

We propose a new model for transaction data that accounts jointly for the time duration between transactions and for the discreteness of the intraday stock price changes. Duration is assumed to follow a stochastic conditional duration model, while price discreteness is captured by an autoregressive moving average ordinal-response model with stochastic volatility and time-varying parameters. The proposed model also allows for endogeneity of the trade durations as well as for leverage and in-mean effects. In a purely Bayesian framework we conduct a forecasting exercise using multiple high-frequency transaction data sets and show that the proposed model produces better point and density forecasts than competing models.

JEL CODE: C11, C13, C18, C22, C50, G00

---

[1]Correspondence to: Stefanos Dimitrakopoulos, Economics Division, Leeds University Business School, Leeds University, UK, E-mail: s.dimitrakopoulos@leeds.ac.uk. All computations have been performed with the generous technical support of GOLCER and its staff. We would also like to thank the two referees, the comments of which greatly improved a first version of this paper.

1 Introduction

High frequency transaction data are important in studying market microstructure, as they exhibit unique characteristics that can not be found in lower frequencies. A major characteristic of trading data is that the price changes between successive trades tend to occur in multiples of a tick size (tick by tick data) and therefore can be considered as a discrete-valued variable. The relevant empirical literature on price discreteness is voluminous; see, for example, Gottlieb and Kalay (1985), Glosten and Harris (1988), Ball (1988), Harris (1990), Dravid (1991), Hasbrouck (1999), Rydberg and Shephard (2003), among others.

Among the various specifications that have been proposed for modelling high-frequency integer price changes are the ordinal-response (OR) models. These models recognise that the responses (price changes) are discrete and ordered. To capture better the behaviour of trading data, the ordinal-response models have also been extended to account for conditional heteroscedasticity. For example, Hausman et al., (1992) considered an ordered probit model for the analysis of transaction stock prices, where the time-varying conditional variance was a function of time between trades and of the lagged spread. Müller and Czado (2009) set up an ordered probit model with stochastic volatility, while Yang and Parwada (2012) and Dimitrakopoulos and Tsionas (2019) showed that there are forecast gains stemming from ordered probit models with generalized autoregressive conditional heteroscedasticity (GARCH).

Another empirical characteristic of stock trading is that it occurs in unequally spaced time intervals. In other words, the duration of the transactions varies. Duration was initially considered an exogenous variable. Under this assumption, Hausman et al., (1992) and Fletcher (1995) found evidence of correlation between duration and transaction price changes.

However, it has been documented that time duration contains information about trading intensity and price adjustments (Diamond and Verrechia, 1987; Admati and Pfleiderer, 1988, 1989; Easley and O'Hara, 1992; Dufour and Engle, 2000). Therefore, duration can not be considered exogenous to the stock price formation. Early research on modelling the evolution of positive transaction duration began with the seminal paper of Engle and Russell (1998), who proposed the autoregressive conditional duration (ACD) model. Since then, many variations of it have been put forward (Lunde, 1999; Bauwens and Giot, 2000; Grammig and Maurer, 2000; Zhang et al., 2001; Ghysels et al., 2004; Bauwens and Veredas, 2004).

Engle (2000) extended the Engle and Russell (1998) work by proposing a bivariate model of duration and returns, where the volatility of returns followed a GARCH-type specification (ACD-GARCH model). Long durations influenced negatively both returns and volatility of returns; see also Dufour and Engle (2000), Grammig and Wellner (2002), and Ait-Sahalia and Mykland (2003). Other similar joint models of returns and duration include those of Engle and Sun (2005), Ghysels and Jasiak (1998) and Liu and Maheu (2012).

In this paper we build upon the Dimitrakopoulos and Tsionas (2019) model to analyze irregularly spaced transaction data, by proposing a new joint model of price differences and transaction duration. In particular, the duration process is accounted for in the context of the stochastic conditional duration (SCD) specification (Bauwens and Veredas, 2004), in which the conditional duration is modeled as a latent variable instead of being observable as in the ACD model. Although more flexible alternatives exist (Fernandes and Grammig, 2006; Fernandes et al, 2016), the main goal of this paper is to present a parsimonious/representative model for forecasting irregularly spaced transaction price changes. Time duration is adjusted to net out diurnal patterns that are present in a normal trading day (Engle and Russell, 1998; Engle, 2000; Bauwens and Veredas, 2004).

Price discreteness is captured by an ordinal-response model with conditional heteroscedasticity, represented by a stochastic volatility (SV) process[1]. Furthermore, to take into account the unequally-spaced feature of the price changes, our SV equation is specified for the variance per time-unit. Therefore, our SV formulation implies a time-varying parameter SV equation for the total variance process. In the context of GARCH specifications, a similar approach has been followed by Engle (2000) and Meddahi et al. (2006). To be consistent with our analysis, the conditional mean of the OR model is also specified per time-unit, entailing a time-varying total conditional mean.

To control for the dynamic behaviour of price change observations as well as for microstructure effects, our OR-SV model is augmented to allow for an autoregressive moving average (ARMA) component, in-mean effects (M) and leverage effects (L). However, our model, as it stands, assumes that transaction intensity is exogenous.

Numerous studies have shown that this is not true, as there is an interdependence between

[1]OR-SV models have been previously considered in the econometrics literature by Müller and Czado (2009) and Dimitrakopoulos and Dey (2017), among others. Also, in modelling intraday returns, SV models have been found to produce lower forecasting errors than GARCH models (Stroud and Johannes, 2014).

trade duration and market price behavior (Easley and O'Hara, 1992; Dufour and Engle, 2000; Grammig and Wellner, 2002; Renault and Werker, 2011). We account for this endogeneity issue, by incorporating the latent process for price changes into the volatility equation of the SCD model. Therefore, duration affects the formulation of price changes and vice versa.

The estimation of the parameters is not an easy task, due to the high-dimensional integrals involved in the likelihood function of the model. To this end, we resort to purely Bayesian techniques for the update of all model parameters.

Our main goal is to conduct a forecasting exercise, using numerous (six) stocks traded at the New York Stock Exchange (NYSE). In particular, we compare our proposed model against alternative competing models, in terms of point and density forecasts. To test further the forecasting abilities of the models in question, we resort to the conditional predictive ability (CPA) test (Giacomini and White, 2006), as well as to the Model Confidence Set (MCS) approach (Hansen et al., 2011). Optimal prediction pools are also considered (Geweke and Amisano 2011). In addition, we quantify the trading gains that result from the forecasts by simulating trading rules.

The empirical duration literature on discrete transaction prices lacks such a forecasting exercise. Yang and Parwada (2012) considered an ACD-OR-GARCH model, which was estimated with a quasi-maximum likelihood method. They also performed an in-sample and out-of-sample forecasting analysis to calculate the percentage of correctly predicted price change directions. However, they did not perform a forecast comparison of ordinal-response type models. Moreover, the durations here are modelled stochastically (SCD model) rather than deterministically (ACD model). The same holds for the conditional heteroscedasticity of our OR model (SV versus GARCH). Finally, our specification accounts for endogenous financial durations and additional microstructure effects (ARMA, leverage, in-mean effects).

The paper is structured as follows. In section 2 we set up the main model and in section 3 we describe the posterior analysis. In section 4 we carry out our empirical study. Section 5 concludes. An Online Appendix accompanies this paper.

2 Econometric set up: The proposed model

Denoting by t_i the time associated with the transaction i, the duration between two consecutive transactions $i-1$ and i is given by $\tilde{x}_i = t_i - t_{i-1}$, while the price change from transaction $i-1$ to i is denoted by $y_i = P_i - P_{i-1}$, with $t_{i-1} < t_i$. Due to the presence of diurnal effects, we use throughout our analysis the adjusted time duration, denoted by x_i. In section 4.1, we explain in detail how $\tilde{x}_i$ is diurnally adjusted.

Since intraday price changes tend to occur in multiples of a tick size, the order of which matters, we model these price changes in the context of the following ordinal-response model

$$\begin{aligned}
y_i &= j \Leftrightarrow \zeta_{j-1} < y_i^* \le \zeta_j, \ 1 \le \ j \le J, \\
y_i^* &= \rho_i y_{i-1}^* + \tau_{1i}\sigma_i^2 + \varepsilon_i - \gamma_{1i}\varepsilon_{i-1} + \lambda_i \\
\rho_i &= \tfrac{x_i}{x_{i-1}}\rho, \ \tau_{1i} = \tfrac{\tau_1}{x_i}, \ \gamma_{1i} = \gamma_1 \tfrac{x_i}{x_{i-1}} \\
&\text{with } |\rho| < 1, \tau_1 \in \mathbb{R}, \ |\gamma_1| < 1
\end{aligned} \tag{1}$$

where ζ_j is the j-th cutpoint, y_i^* is a latent dependent variable, having a time-varying ARMA$(1,1)$ specification and (ε_i) is a zero-mean white noise process with a (total) conditional variance $E\left(\varepsilon_i^2|\sigma_i^2\right) = \sigma_i^2$. For the error term of the regression, it holds that $(\lambda_i) \sim iidN(0, \sigma_\lambda^2)$, with λ_i being independent of ε_i. The presence of this term facilitates posterior sampling. Model (1) is equipped with in-mean (M) effects that are reflected in the time-varying parameter τ_{1i}.

As can be seen from expression (1), the parameters in the ARMA(1,1) representation are time-varying. In this way, we take into account the irregularly spaced nature of the data. In fact, the total price difference y_i^* over the period $(t_{i-1}, t_i]$ should account for the irregular trading interval by having a specification "proportional" to the corresponding duration x_i. If $\widetilde{y}_i^* = \frac{y_i^*}{x_i}$ and $\widetilde{\varepsilon}_i = \frac{\varepsilon_i}{x_i}$ denote, respectively, the latent price change "*per time-unit*" and its corresponding ARMA error, then, it can be easily seen that $\widetilde{y}_i^*$ and $\widetilde{\varepsilon}_i$ would satisfy the following time-invariant ARMA(1,1)

$$\widetilde{y}_i^* = \rho\widetilde{y}_{i-1}^* + \tau_1\widetilde{\sigma}_i^2 + \widetilde{\varepsilon}_i - \gamma_1\widetilde{\varepsilon}_{i-1} + \widetilde{\lambda}_i.$$

where $\widetilde{\lambda}_i = \frac{\lambda_i}{x_i}$ and the conditional variance per time-unit $E\left(\widetilde{\varepsilon}_i^2|\sigma_i^2, x_i\right) = \widetilde{\sigma}_i^2 = \frac{\sigma_i^2}{x_i^2}$, is proportional to the inverse of the squared duration.

The error term ε_i exhibits conditional heteroscedasticity in the form of a SV representation

$$\begin{aligned}
&\varepsilon_i = \sigma_i\xi_i = x_i\widetilde{\sigma}_i\xi_i, \quad (\xi_i) \sim \text{iidN}(0,1) \\
&\log\left(\widetilde{\sigma}_i^2\right) = \alpha_0 + \alpha_1 \log\left(\widetilde{\sigma}_{i-1}^2\right) + \eta_i,\ \eta_i = \alpha_2\xi_i + e_i, \quad (e_i)\ \text{iidN}\left(0,1\right) \\
&\text{with } \ |\alpha_1| < 1,\ \alpha_0 \in \mathbb{R},\ \ \alpha_2 \in \mathbb{R} \\
&(\xi_i) \ \text{ and } \ (e_i) \ \text{ are independent}
\end{aligned} \tag{2}$$

where the noise (η_i) is iidN$(0, 1 + \alpha_2^2)$. Model (2) allows for a leverage (L) effect, since the SV innovation (ξ_i) is allowed to be correlated with the log-volatility error η_i with

$$Corr\left(\eta_i, \xi_i\right) = \tfrac{\alpha_2}{\sqrt{1+\alpha_2^2}} \in (-1, 1)\,.$$

Note that the log-volatility equation for the total variance process σ_i^2 would become

$$\begin{aligned}
&\log\left(\sigma_i^2\right) = \alpha_{0i} + \alpha_1 \log\left(\sigma_{i-1}^2\right) + \eta_i,\ \eta_i = \alpha_2\xi_i + e_i, \quad (e_i)\ \text{iidN}\left(0,1\right) \\
&\alpha_{0i} = \alpha_0 + \log\left(x_i\right) - \alpha_1 \log\left(x_{i-1}\right)
\end{aligned}$$

For x_i, we assume the stochastic conditional duration (SCD) model (Bauwens and Veredas, 2004), defined as

$$\begin{aligned}
&x_i = \Psi_i z_i, \quad \Psi_i = e^{\psi_i},\ (z_i)\ \text{ iid with } E\left(z_i\right) = 1 \\
&\psi_i = \omega_1 + \omega_2\psi_{i-1} + \omega_3\widetilde{y}_{i-1}^* + v_i, \quad \text{with } |\omega_2| < 1,\ \omega_1 \in \mathbb{R},\ \omega_3 \in \mathbb{R} \\
&(v_i) \sim iidN\left(0, \sigma_v^2\right),\ \sigma_v^2 > 0
\end{aligned} \tag{3}$$

where the logarithm of the latent variable Ψ_i follows a stationary AR(1) process[2] and (z_i) is the duration innovation which is an iid sequence with a distribution $f_z(\cdot)$ having a nonnegative support and a unit mean. Here, we assume an exponential distribution for $f_z(\cdot)$ with unit parameter. Also, the inclusion of the term $\widetilde{y}_{i-1}$ in the log-conditional mean equation of model (3) is made to render the duration endogenous to the price change variable. The error terms (z_i) and (v_i) of the SCD model are independent (endogeneity).

We name the proposed model, described by (1), (2) and (3), the SCD-ARMA-OR-SVLM model. For identification purposes, we set $\alpha_0 = 1$ (Hausman et al., 1992; Müller and Czado,

[2]In expression (3), we use $\widetilde{y}_{i-1}^*$ instead of y_{i-1}^* to maintain the stationarity condition $|\omega_2| < 1$.

2009) and for the cutpoints we assume the usual order constraint restrictions $\zeta_0 = -\infty < \zeta_1 < \zeta_2 < ... < \zeta_J = +\infty$.

Furthermore, our model for irregularly spaced price changes touches upon the Engle (2000) specification. Engel (2000) proposed an ACD$(1,1)$ for duration x_i and a time-varying GARCH$(1,1)$ (not for price changes but) for the returns y_i with the GARCH coefficients depending on x_i. However, Engle (2000)'s model does not account for the endogeneity of durations as well as for other interesting effects, such as ARMA, leverage and in-mean effects. In addition, since our primary goal is to model the price changes, we divided y_i^* by x_i (and not by $\sqrt{x_i}$ as in Engle, 2000) and accordingly σ_i^2 by x_i^2 (and not by x_i as in Engle, 2000). Therefore, volatility is diurnally adjusted.

The proposed specification encompasses a large number of submodels. For example, assuming a continuous dependent variable y_t (i.e, returns) the SCD-ARMA-OR-SVLM can be easily reduced to the SCD-ARMA-SVLM or the SCD-SVLM, which, in turn, can further be decomposed into the SCD-ARMA-SVL and SCD-ARMA-SVM or the SCD-SVL and SCD-SVM, respectively. Also, if duration is ignored, the proposed model reduces to the ARMA-OR-SVLM model.

3 Posterior analysis

3.1 MCMC

It is straightforward to update the parameters $\omega_1, \omega_2, \omega_3, \alpha_1, \alpha_2, \rho, \tau_1, \gamma_1, \sigma_v^2, \sigma_\lambda^2$ from well-defined Gibbs conditionals. The cutpoints are updated as in the Dimitrakopoulos and Tsionas (2019) paper. The volatilities in model (2) are updated using the Kim et al., (1998) approach. It remains to sample the ψ_is and the errors ε_is.

The contribution of the ith observation is

$$\begin{gathered}
e^{-\psi_i - \exp(-\psi_i x_i)}. \\
\sigma_v^{-1} \exp\left\{-\tfrac{1}{2\sigma_v^2}\left(\psi_i - \omega_1 - \omega_2\psi_{i-1} - \omega_3\widetilde{y}_{i-1}^*\right)^2\right\}\cdot \\
\sigma_\lambda^{-1} \exp\left\{-\tfrac{1}{2\sigma_\lambda^2}\left(y_i^* - \rho_i y_{i-1}^* - \tau_{1i}\sigma_i^2 - \varepsilon_i + \gamma_{1i}\varepsilon_{i-1}\right)^2\right\}\cdot \\
\exp\left\{-\tfrac{1}{2}\left(\log\left(\widetilde{\sigma}_i^2\right) - \alpha_0 - \alpha_1\log\left(\widetilde{\sigma}_{i-1}^2\right) - \alpha_2\xi_i\right)^2\right\}\cdot \\
\tfrac{1}{x_i\widetilde{\sigma}_i}\exp\left(-\tfrac{1}{2x_i^2\widetilde{\sigma}_i^2}\varepsilon_i^2\right)\cdot \\
y_i = j \Leftrightarrow \zeta_{j-1} < y_i^* \leq \zeta_j. \\
-\infty = \zeta_0 < \zeta_1 < \ldots < \zeta_6 = \infty.
\end{gathered} \tag{4}$$

The conditional posterior of each ψ_i is log-concave. Its mode satisfies:

$$-1 + x_i e^{-\psi_i x_i} - \tfrac{1}{\sigma_e^2}(2\psi_i - a_{i1} - a_{i2}) = 0, \tag{5}$$

where $a_{i1} = \omega_1 + \omega_2\psi_{i-1} + \omega_3\widetilde{y}_{i-1}^*$, $a_{i2} = \left(\psi_{i+1} - \omega_1 - \omega_3\widetilde{y}_i^*\right)\omega_2$. The second derivative is $-\left(x_i^2 e^{-\hat{\psi}_i x_i} + \frac{2}{\sigma_e^2}\right)$ where $\hat{\psi}_i$ is the solution to (5). In turn, one can use rejection sampling for log-concave densities. The conditional posterior of $h_i = \log\left(\widetilde{\sigma}_i^2\right)$ is also log-concave. Its mode solves the equation:

$$-h_i\left(1 + \alpha_1^2 + \tau_1^2(x_i/x_{i-1})^2\right) + \tfrac{1}{\sigma_\lambda^2}q_i\tau_1(x_i/x_{i-1}) + a_{i1} - a_{i2} + \tfrac{1}{2}, \tag{6}$$

where $q_i = y_i^* - \frac{x_i}{x_{i-1}}\left(\rho y_{i-1}^* + \varepsilon_i - \gamma_1\varepsilon_{i-1}\right)$, $a_{i1} = \alpha_0 + \alpha_1\log\left(\widetilde{\sigma}_{i-1}^2\right) + \alpha_2\xi_i$, $a_{i2} = \log\left(\widetilde{\sigma}_{i+1}^2\right) - \alpha_0 - \alpha_2\xi_{i+1}$ and the second derivative at the mode is $-\left[1 + \alpha_1^2 + \tau_1^2(x_i/x_{i-1})^2\right]$. Finding the mode does not require solving nonlinear equations.

Once a density is log-concave we do not have to use full rejection. If the mode and the second derivative are known, we can use a normal candidate-generating distribution whose mean is the mode and its variance is the inverse of the second derivative at the mode. Rarely this method needs more than three or fourth rejections per acceptance.

The conditional posterior of ε_i can be written as

$$p(\varepsilon_i|\cdot) \propto e^{-(\varepsilon_i - \hat{\varepsilon}_i)^2/(2s_i^2)}, \tag{7}$$

where

$$\hat{\varepsilon}_i = -\sigma_\lambda^2 \left(R_i^2 + R_{i+1}^2 + \frac{\sigma_\lambda^2}{x_i \sigma_i^2} + R_i^2 + \sigma_\lambda^2 \alpha_2^2 \right)^{-1} \left(-R_i a_{i1} - R_{i+1} a_{i2} + \sigma_\lambda^2 a_{i3} \alpha_2 \right),$$

$$s_i^2 = \sigma_\lambda^2 \left(R_i^2 + R_{i+1}^2 + \frac{\sigma_\lambda^2}{x_i \sigma_i^2} + R_i^2 + \sigma_\lambda^2 \alpha_2^2 \right)^{-1},$$

$$R_i = \frac{x_i}{x_{i-1}}, a_{i1} = y_i^* - \rho_i y_{i-1}^* - \tau_{1i} \sigma_i^2 + \gamma_{1i} \varepsilon_{i-1},$$

$$a_{i2} = y_{i+1}^* - \rho_{i+1} y_i^* - \tau_{1(i+1)} \sigma_{i+1}^2 + \gamma_{1(i+1)} \varepsilon_{i+1},$$

$$a_{i3} = \log\left(\widetilde{\sigma}_i^2\right) - \alpha_0 - \alpha_1 \log\left(\widetilde{\sigma}_{i-1}^2\right) - \alpha_2 \xi_i.$$

3.2 Point and density forecasts

Since the main goal of this paper is the evaluation of the SCD-ARMA-OR-SVLM model, in terms of its forecasting performance, we compute out-of-sample point and density r-step-ahead forecasts with $r = 1$, 5 and 10.

Let us collect the parameters of the model into the vector $\mathbf{\Theta}$, and let $DATA_t = \{(y_i, x_i), i = 1, ..., t\}$. The 1-step ahead conditional predictive density for y_{t+1} is given by

$$p(y_{t+1}|DATA_t) = \int p(y_{t+1}|DATA_t, \mathbf{\Theta}) p(\mathbf{\Theta}|DATA_t) d\mathbf{\Theta}, \tag{8}$$

where the integration is executed with respect to the posterior distribution of the parameters based on the information set $DATA_t$ and

$$\begin{aligned} p\left(y_{t+1} = j|DATA_t, \mathbf{\Theta}\right) \quad = \quad & \Phi\left(\frac{\zeta_j - \left(\rho_{t+1} y_t^* + \tau_{1,t+1} \sigma_{t+1}^2 + \varepsilon_{t+1} - \gamma_{1,t+1} \varepsilon_t\right)}{\sigma_\lambda}\right) - \\ & \Phi\left(\frac{\zeta_{j-1} - \left(\rho_{t+1} y_t^* + \tau_{1,t+1} \sigma_{t+1}^2 + \varepsilon_{t+1} - \gamma_{1,t+1} \varepsilon_t\right)}{\sigma_\lambda}\right) \end{aligned} \tag{9}$$

where $\Phi\left(.\right)$ denotes the cumulative distribution function of the standard Normal distribution.

Expression (8) is used as a density forecast. The density forecasts are evaluated using the sum of the log predictive likelihoods

$$LPS = \sum_{i=t_0}^{T-1} \log p(y_{i+1} = y_{i+1}^o | DATA_i), \tag{10}$$

where $p(y_{i+1} = y^o_{i+1}|DATA_i)$ is the predictive density of y_{i+1} evaluated at the observed value y^o_{i+1} and $i = t_0, ..., T-1$ is the evaluation period.

The root mean squared forecast error (RMSFE), which is a metric for the evaluation of point forecasts, is defined as

$$RMSFE = \sqrt{\frac{\sum_{i=t_0}^{T-1}(y^o_{i+1} - E(y_{i+1}|DATA_i))^2}{T-1-t_0+1}}. \tag{11}$$

Higher LPS values and lower RMSFE values indicate better forecasting performance. The r-step ahead forecast ($r > 1$) can be obtained by applying sequentially the 1-step ahead forecast from $k = t+1$ to $k = t+r$ while replacing at each time the y_{t+k} ($1 \leq k \leq r$) by the observed one y^o_{t+k}.

3.3 Additional forecasting tools

To verify the validity of the forecasting results, we use additional econometric tools. One such tool is the Conditional Predictive Ability (CPA) test of Giacomini and White (2006). This pairwise Wald-type test is based on a sequence of loss functions used to evaluate a corresponding sequence of out-of-sample forecasts, obtained from rolling samples. A failure to reject the null hypothesis implies that the two competing models have equal predictability. In our empirical applications, the CPA test is implemented to point forecasts, where as a loss function we use the squared error loss and rolling windows are made up of 100 observations.

For testing the predictive performance of multiple competing models, we resort to the Model Confidence Set (MCS) approach of Hansen et al., (2011). Since it is a multiple-wise test, there is no need to specify a benchmark model. The logic of this approach is rather simple. We start from the full set of models and through a sequential testing process we eliminate the worst performing models, until we reach a set of models, the MCS, for which the null hypothesis of equal predictive ability is not rejected at a certain confidence level a. The construction of MCS test statistic is based on the Diebold-Mariano (1995) test statistic. For the calculation of the MCS p-values we used the block bootstrap method, setting the length of the block equal to ten.

Last but not least, we considered weighted linear combinations of prediction models in order to obtain the optimal prediction pool. Such a pool is obtained by maximizing the log pooled predictive score function (predictive likelihood)

$$\underset{w_l, l=1,...,L}{\operatorname{argmax}} \sum_{i=\tau_1}^{\tau_2} \log \left(\sum_{l=1}^{L} w_l p(y_i | y_1, ..., y_{i-1}, \mathcal{L}_l) \right), \tag{12}$$

where $w_l \geq 0, l = 1, ..., L$ is the weight under the restriction that $\sum_{l=1}^{L} w_l = 1$, and $p(y_i|y_1, ..., y_{i-1}, \mathcal{L}_l)$ is the predictive density for model $\mathcal{L}_l$ evaluated at the observed value y_i. From this maximization, which is done using a standard nonlinear solver (Nash, 1984), we obtain the optimal weight vector, which we report in our empirical results. The calculation of the weights is based on the last 500 observations.

4 Empirical analysis

4.1 Data

We conduct a formal forecasting exercise, using 6 stocks traded at the New York Stock Exchange (NYSE); see Table 1. The transaction data sets were obtained from the Trades and Quotes (TAQ) database, published by the NYSE. Each stock was recorded on five normal trading days of the week and all trades before 9:30 AM and after 4:00 PM were discarded.

In table 1, we have chosen stocks from time periods with different volatility characteristics. For example, the first three stocks refer to the years 2008, 2010, 2011, respectively that reflect different phases of the 2008 global economic crisis, with the most volatile one being that of 2008. Also, the selected stocks differ in the magnitude of the tick size. The tick size on NYSE was $\$\frac{1}{18}$ before June 24, 1997, $\$\frac{1}{16}$ before January 29, 2001 and since January 29, 2001, the price changes tend to occur in multiples of one cent.

In high-frequency transaction data sets it is possible to have multiple transactions, even with different prices, within a single second. We aggregated all the transactions done on the same second into a single transaction, with its corresponding price being the volume weighted average price of the aggregated transactions. Trade durations along with their corresponding price changes were, then, computed.

However, it is well known that intraday data exhibit a strong diurnal pattern. The trading intensity of a stock varies during the day; there are more transactions at the beginning and at the end of the trading hours, while there are less transactions around lunch time. Consequently, after the open and prior to the close of the market, time durations are very short, whereas around lunch break, durations are much longer. This empirical fact creates a daily cyclical pattern for the time durations between transactions that resembles an inverted "U" shape.

To this end, durations are diurnally adjusted. The diurnal effect is removed by defining $x_i = \tilde{x}_i/\hat{x}_i$, where $\tilde{x}_i$ is the original duration, $\hat{x}_i$ is the corresponding fitted value of duration and x_i is the adjusted duration. The estimated value $\hat{x}_i$ is obtained by regressing the durations on the time of day and on the day of the week, using a cubic spline specification with knots set on each hour (10:00, 11:00, 12:00, 13:00, 14:00, 15:00, 16:00). Hence, each weekday is allowed to have its own diurnal factor (Bauwens and Giot, 2000).

All our empirical analysis makes use of the diurnally adjusted durations. Since the volatility process also exhibits a daily seasonal component similar to that of the duration process, intradaily volatility is also diurnally adjusted, as we mentioned in Section 2.

4.2 Modelling strategies

Our main specification, the SCD-ARMA-OR-SVLM model (without covariates) is compared against two other models. The first one is its continuous counterpart, the SCD-ARMA-SVLM model, and the second one is the ARMA-OR-GARCHLM model, proposed by Dimitrakopoulos and Tsionas (2019).

For ease of exposition, throughout the paper, we report the ratio of the LPS value of a given model to that of the autoregressive ordinal-response (AR-OR) model, which is used as the baseline model. Hence, values greater than one indicate better forecasting power compared to the baseline model. We also report the ratio of the RMSFE of a given model to that of the baseline model. Therefore, values lower than one indicate better forecasting performance than the AR-OR model.

For illustration purposes and due to space constraints, we present in the next section (Section 4.3) the full estimation and forecasting results for only one data set[3], the choice of which is purely random. In Section 4.4 we summarize the forecasting results obtained from the rest of

[3]The estimation and forecasting results for the rest of the empirical data sets are available upon request.

the data sets.

4.3 Data set II: The JNJ stock

The tick-by-tick data set of JNJ is recorded on normal trading days of the week from October 4 to October 8, 2010, within normal trading hours (9:30 am to 4:00 pm). In total, there were 43150 adjusted durations, the time series plot of which is given in Figure 1. The spline estimates of durations for every day of the week are displayed in Figure 2. There is, as expected, a strong intradaily pattern, where the daily spline for durations exhibits an inverse "U" shape.

The histogram of transaction price changes is plotted in Figure 3[4]. As these changes tend to be clustered on multiples of 1 cent, they can be viewed as an ordinal variable y_t. We also noticed that price changes of less than 2 cents and of more than 2 cents were very few, constituting only 0.61% of the total changes. Therefore, the discretisation of the price changes is done as follows. Prices falling more than two cents (< -2) are grouped under one category and are assigned the value of "1", price changes falling in the interval $[-2, -1)$ are assigned the value of "2" and so on up to price changes that were larger than 2 cents (> 2), in which case y_t is assigned the value of "6". Therefore, y_t in our analysis takes on 6 categories. In Table 2 we display the intervals of price changes, the frequencies of price changes falling in each interval and the values of the ordinal dependent variable for the JNJ stock data.

4.3.1 Estimation results

We run the MCMC algorithms for 150000 iterations, after discarding the first 50000 draws and reported the posterior means and posterior standard deviations for the parameters of the SCD-ARMA-OR-SVLM and SCD-ARMA-SVLM models in Table 3.

The estimates of ω_2 are high in magnitude (0.810 in the proposed model and 0.721 in its continuous version), meaning that the time trade between transactions is a persistent process that exhibits the stochastic volatility-type property of clustering; long (short) trading durations tend to be followed by long (short) trading durations. The same holds for the evolution of the log-volatility of ε_i, the persistence of which is reflected in the parameter α_1.

[4]From the histogram of Figure 3 it can be seen that a large number of price changes concentrates on the zero value. A zero-inflated version of the proposed could potentially tackle this problem that we will consider in a future paper.

In addition, the volatility of duration is affected by the formulation of latent price changes and this effect is captured by the parameter ω_3. This empirical fact verifies the literature on transaction durations that supports that market price behaviour and duration are not independent.

The positive values of the posterior means of the autoregressive parameters ρ in both models indicate strong state dependence in the latent process of transaction price changes. The sign of this parameter is also in agreement with that of previous studies (Müller and Czado, 2005; Dimitrakopoulos and Tsionas, 2019). Yet, when duration is accounted for, the magnitude of this parameter increases notably.

The conditional variance of y_i^* affects in a positive way the conditional mean of it, as can be seen from the reported in-mean parameters τ_1 that control for the volatility feedback effect. Moreover, the posterior means of γ_1 signal some presence of market microstructure dynamics in the discretization of stock price changes that should not be neglected.

We also observe that the SCD-ARMA-SVLM model that ignores the discrete nature of price changes underestimates the effects, captured by the parameters ω_2, τ_1, and α_1, as well as the leverage effects (α_2). The coefficient on α_2 was found to be negative, a result supported by previous relevant studies (for example, Dimitrakopoulos and Tsionas (2019)).

The priors and the posterior distributions of all the parameters of the two models are displayed in the Online Appendix.

4.3.2 Forecasting results

We begin the forecast evaluation of our competing models. To this end, we first computed point and density forecasts for r-step ahead transactions ($r = 1, 5, 10, 20$). The density forecast results are presented in Table 4. The first observation is that the proposed model outperforms all its competitors across all forecast horizons, producing the largest relative log-predictive score. The model that produced the second best density forecasts overall was the SCD-ARMA-SVLM model. The ARMA-OR-GARCHLM model that ignores duration had the worst performance apart from $r = 5$, where it did better than the SCD-ARMA-SVLM model.

The results for the point forecasts are given in Table 5. The dominant model is the SCD-ARMA-OR-SVLM once again, with the SCD-ARMA-SVLM model occupying the second best position. The point forecasts produced by these two models are significantly different at 5% level from those produced by the baseline model (AR-OR), a result supported by the p-values of the Giacomini-White (2006) test (see Table 6). However, this is not the case for the ARMA-OR-GARCHLM model.

In the same table (Table 5) we report the superior set of models with equal predictability, based on the MSC-p values (see Table 7). This set includes the proposed model along with its continuous version for all horizons at the 95% significant level. From this set is excluded the model that ignores transaction duration (ARMA-OR-GARCHLM). Overall, from Tables 4 and 5, it can be seen that accounting for duration, time-varying parameters in the conditional mean of the OR model as well as for unobserved components (SV and SCD), we achieve better forecasting ability than the ARMA-OR-GARCHLM model that was found to be the winner among several alternative submodels in the Dimitrakopoulos and Tsionas paper (2019).

As a second step in our forecasting analysis, we computed the weights (Geweke and Amisano, 2011) for each of the last 500 transactions and we did that for each model and for each horizon. These weights were obtained by considering weighted linear combinations of the three competing models (pools) for forecasting r-step ahead transactions ($r = 1, 5, 10, 20$). The evolution of these weights is presented in Figure 4. It is clear that most of the weight is allocated to the proposed model, irrespective of the horizon over the entire out-of-sample forecasting period. The second largest weight is received by the SCD-ARMA-SVLM model, especially for longer horizons ($r = 10, 20$), throughout the forecast period. For shorter horizons ($r = 1, 5$), the ARMA-OR-GARCHLM model appears to "weighs a bit more" than the SCD-ARMA-SVLM model.

In addition, we maximized expression (12) over the last 500 observations to obtain the optimal weights for the 3-model pool for each horizon. We report these weights in Table 11 (third row-Data set II). For $r = 1$, the proposed model was assigned almost all the optimal weight (0.92). The rest of the (negligible) weight was mostly given to the ARMA-OR-GARCHLM (reported in red as M3). A similar story is repeated for $r = 5$, albeit more optimal weight is assigned to the SCD-ARMA-OR-SVLM model (0.98). For $r = 10$ and $r = 20$, the winner is still the proposed model but now the "least weighted" model is the ARMA-OR-GARCHLM model.

In Figure 5, we demonstrate graphically how close the series of the predicted cumulative price changes generated by each of the three competing models are to the series of actual cumulative price changes. We conduct that experiment for $r = 1$ and for four different time intervals of 10000 observations each. The proposed model has a remarkable fit to the real path of cumulative price changes (panels (a)-(c)). On the contrary, the predicted series obtained from the SCD-ARMA-SVLM model exhibits repeated deviations from the actual series and these deviations are more obvious in panel 5(b). The ARMA-OR-GARCHLM fails substantially to trace the evolution of the real price changes in all time intervals (panels (a)-(c)). In Figure 5(d) we have generated from the optimal pool of models the predicted values for the last 5000 data points of the JNJ data set. The posterior model pool is able to follow well the actual path.

As a last step in our analysis, we quantified the gains of predictions in terms of their profitability, based on a trading rule. This rule states that we sell when we predict a decrease in price and buy otherwise. The decrease or increase is evaluated as "significant" if the price difference changes category. Based on the predictions generated from the optimal pool of models as well as from the SCD-ARMA-SVLM model, we calculated the average excess returns for each forecast horizon ($r = 1, 5, 10, 20$). The results are presented in the third row of Table 12. The pool of models produced out-of-sample average excess returns that were greater than those of the SCD-ARMA-SVLM by 1.89% for $r = 1$, by 2.52% for $r = 5$, by 2.97% for $r = 10$ and by 3.310% for $r = 20$. We also repeated the same experiment but this time we restricted the number of trades to three per day (see Table 13). We observe that the gains of predictions based on the optimal pool of models over the SCD-ARMA-SVLM model increase with the horizon.

4.4 General findings for the rest of the data sets

For the rest of the data sets of Table 1, we repeated the same analysis as before. We present summary tables (Tables 8-13) only for the forecasting results related to these data sets. In these tables, the models SCD-ARMA-OR-SVLM, SCD-ARMA-SVLM and ARMA-OR-GARCHLM are reported as M1, M2 and M3 models, respectively.

Table 8 presents the model that outperforms all the others in term of the LPS criterion across different horizons (r = 1, 5, 10 and 20) and across different data sets. As can be seen, the proposed model (M1) produced the best density forecasts. The same story is repeated in Table 9 that presents the best model in terms of point forecasts.

Table 10 displays the set of models with equal predictability, based on the MCS-p values. In any case, this set includes always the M1 model. The M3 model appears only in the first data set and only for $r = 1$. The SCD-ARMA-SVLM model (M2) appears more frequently into the MCS. In particular, for $r = 1$, the M2 model belongs to the MCS in four out of six times, for $r = 5$ in two out of six times and for $r = 10$ and $r = 20$ only 1 time.

We also calculated the optimal weights for the pool of models optimized over the last 500 observations; see Table 11. In this table, the reported model in black received the largest weight, whereas the reported model in red is the second best model in terms of optimal weights. Almost all the weight (90% and above) is allocated to the M1 model , irrespective of the time horizon and the data set. We also see that for $r = 1$, the second largest weight is allocated to the M3 model (apart from data set I). For all other horizons and data sets, the second "most weighted" model is the M2 model.

To sum up, based on the above analysis, the first best forecasting model is the proposed one, followed by its continuous version, with the ARMA-OR-GARCHLM model occupying the worst position. This empirical fact highlights the importance of accounting for unobserved versions of the GARCH and the ACD (that is, SV and SCD) as well as for time-varying parameters in OR models, when analyzing high-frequency trading data.

Last but not least, the trading rule used for the JNJ data set was also applied to rest of the data sets. The results, which are presented in Tables 12 (all trades) and 13 (restricted trades), indicate trading gains obtained from the optimal pool of models over those obtained from the SCD-ARMA-SVLM model. For the restricted trading rule these gains increase as we move further in the future (from $r = 1$ to $r = 20$) for every data set.

5 Conclusions

In this paper we propose a new joint model of intraday price changes and duration. Time duration between successive trades follows a stochastic conditional duration model, while the price discreteness is captured by an ordinal-response time series model with stochastic volatility and time-varying parameters. The ordinal-response model controls for additional microstructure effects, related to intraday transaction stock prices, such as leverage effects, autoregressive moving average components and conditional heteroscedasticity in mean. The proposed specification accounts also for endogeneity issues related to trade durations. Using Bayesian techniques, we conducted an empirical forecasting exercise and the results indicated the dominance of our proposed model over competing specifications.

References

Admati, A.R., & Pfleiderer, P. (1988). A theory of intraday patterns: Volume and price variability. *Review of Financial Studies*, 1, 3–40.

Admati, A.R., & Pfleiderer, P. (1989). Divide and conquer: A theory of intraday and day-of-the-week mean effects. *Review of Financial Studies*, 2, 189–223.

Ait-Sahalia, Y., & Mykland, P.A. (2003). The effects of random and discrete sampling when estimating continuous-time diffusions. *Econometrica*, 71, 483–549.

Andrieu, C., Doucet, A., & Holenstein, R. (2010). Particle Markov chain Monte Carlo methods (with discussion). *Journal of the Royal Statistical Society: Series B*, (2), 1–33.

Ball, C.A. (1988). Estimation bias induced by discrete security prices. *Journal of Finance*, 43, 841–865.

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. (2004). Asymmetric ACD models: Introducing price information in ACD models. *Empirical Economics*, 28, 709–731.

Bauwens, L., & Veredas, D. (2004). The stochastic conditional duration model: A latent variable model for the analysis of financial durations. *Journal of Econometrics*, 119, 381–412.

Diamond, D.W., & Verrecchia, R.E. (1987). Constraints on shortselling and asset price adjustments to private information. *Journal of Financial Economics*, 18, 277–311.

Dimitrakopoulos, S., & Dey, D.K. (2017). Discrete-response state space models with conditional heteroscedasticity: An application to forecasting the federal funds rate target. *Economics Letters*, 154, 20 – 23.

Dimitrakopoulos, S., & Tsionas, M. (2019). Ordinal-response GARCH models for transaction data: A forecasting exercise. *International Journal of Forecasting*, 35, 1273 – 1287.

Creal, D.D., & Tsay, R.S. (2015). High dimensional dynamic stochastic copula models. *Journal of Econometrics*, 2, 335–345.

Dravid, A.R. (1991). Effects of bid–ask spreads and price discreteness on stock returns.

Working Paper, Wharton School, University of Pennsylvania.

Dufour, A., & Engle, R.F. (2000). Time and the price impact of a trade. *Journal of Finance*, 6, 2467–2498.

Easley, D., & O'Hara, M. (1992). Time and the process of security price adjustment, *Journal of Finance*, 47, 577–605.

Engle, R.F. (2000). The econometrics of ultra-high-frequency data. *Econometrica*, 68, 1–22.

Engle, R.F., & Russell, J.R. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 66, 1127 – 1162.

Engle, R.F., & Sun, Z.(2005). Forecasting volatility using tick by tick data. Working Paper. Department of Finance, New York University.

Fernandes, M., & Grammig, J. (2006). A family of autoregressive conditional duration models. *Journal of Econometrics*, 130, 1 – 23.

Fernandes, M., Medeiros, M.C., & Veiga, A. (2016). A (Semi)Parametric Functional Coefficient Logarithmic Autoregressive Conditional Duration Model. *Econometric Reviews*, 35, 1221–1250.

Fletcher, R.A. (1995). The role of information and the time between trades: An empirical investigation. *Journal of Financial Research*, 18, 239–260.

Geweke, J., & Amisano, G. (2011). Optimal prediction pools. *Journal of Econometrics*, 164, 130–141.

Giacomini, R., & White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74, 1545–1578.

Girolami, M., & Calderhead, B. (2011). Riemann manifold Langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73, 123–214.

Gottlieb, G., & Kalay, A. (1985). Implications of the discreteness of observed stock prices. *Journal of Finance*, 40, 135–153.

Grammig, J., & Maurer, K.-O. (2000). Non-monotonic hazard functions and the autoregressive conditional duration model. *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.

Ghysels, E., Gourieroux, C., & Jasiak, J. (2004). Stochastic volatility duration models. *Journal of Econometrics*, 119, 413–433.

Ghysels, E., & Jasiak, J. (1998). GARCH for irregularly spaced financial data: the ACD-GARCH model. *Studies in Nonlinear Dynamics and Econometrics*, 2, 133–149.

Harris, L. (1990). Estimation of stock price variances and serial covariances from discrete observations. *Journal of Financial and Quantitative Analysis*, 25, 291–306.

Hasbrouck, J. (1999). The dynamics of discrete bid and ask quotes. *Journal of Finance*, 54, 2109–2142.

Hansen, P.R., Lunde, A., & Nason, J.M. (2011). The model confidence set. *Econometrica*, 79, 453–497.

Hausman, J.A., Lo, A.W., & MacKinlay, A.C. (1992). An ordered probit analysis of transaction stock prices. *Journal of Financial Economics*, 31, 319–379.

Kim S., Shepherd N., & Chib S. (1998).Stochastic volatility: Likelihood inference and comparison with ARCH Models. *Review of Economic Studies*, 65: 361–393.

Liu, C., & Maheu, J.M. (2012). Intraday dynamics of volatility and duration: Evidence from Chinese stocks. *Pacific-Basin Finance Journal*, 20, 329–348.

Lunde, A. (1999). A generalized gamma autoregressive conditional duration model. *Discussion Paper*, Aarlborg University.

Meddahi, N., Renault, E., & Werker, B. (2006). GARCH and irregularly spaced data. *Economics Letters*, 90, 200–204.

Müller, G.J., & Czado, C. (2009). Stochastic volatility models for ordinal valued time series with application to finance. *Statistical Modelling*, 9, 69–95.

Nash, S.G. (1984). Newton–type minimization via the Lanczos method. *SIAM Journal on Numerical Analysis*, 21, 770–788.

Renault, E., & Werker, B.J.M. (2011). Causality effects in return volatility measures with random times. *Journal of Econometrics*, 160, 272–279

Rydberg, T.H., & Shephard, N. (2003). Dynamics of trade-by-trade price movements: Decomposition and models. *Journal of Financial Econometrics*, 1, 2–25.

Stroud, J.R., & Johannes M.S. (2014). Bayesian modeling and forecasting of 24-hour high-frequency volatility. *Journal of the American Statistical Association*,109, 1368–1384.

Yang, J.W., & Parwada, J. (2012). Predicting stock price movements: An ordered probit analysis on the Australian securities exchange. *Quantitative Finance*, 12, 791–804.

Zhang, M.Y., Russell, J.R., & Tsay, R.S. (2001). A nonlinear autoregressive conditional duration model with applications to financial transaction data. *Journal of Econometrics*, 104, 179–207.

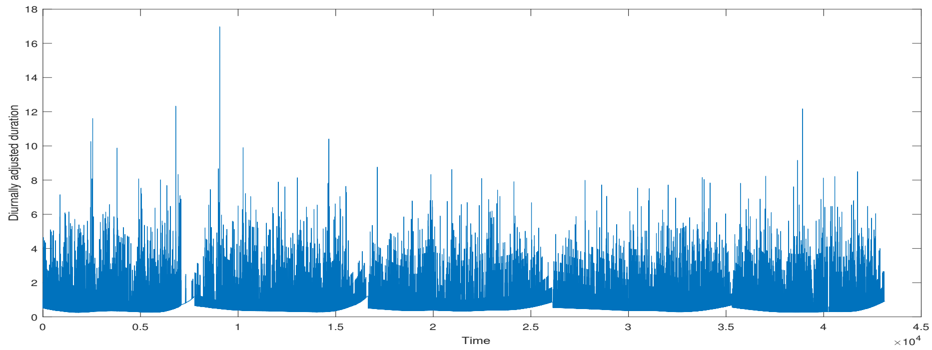

Figure 1: Empirical results. Time plot of diurnally adjusted trade durations for the JNJ stock.

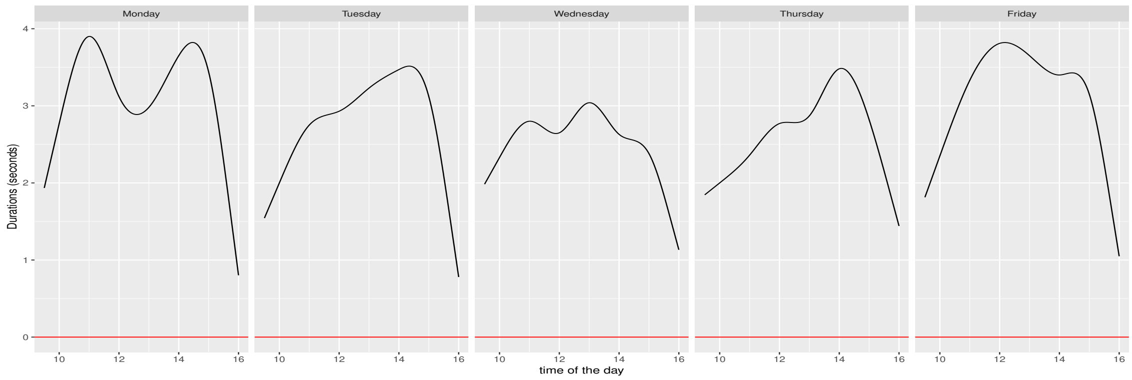

Figure 2: Empirical results. Daily Spline estimates for the JNJ duration.

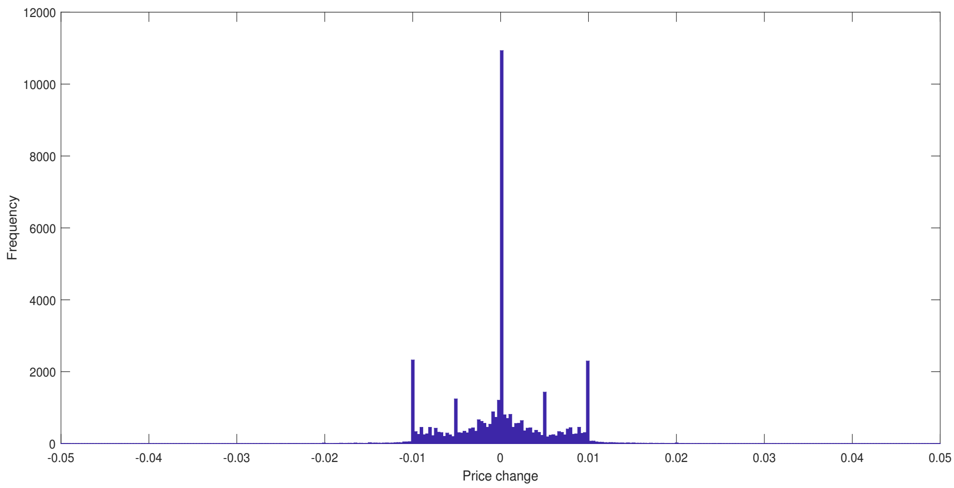

Figure 3: Empirical results. Histogram of transaction price changes (measured in cents) for the JNJ stock.

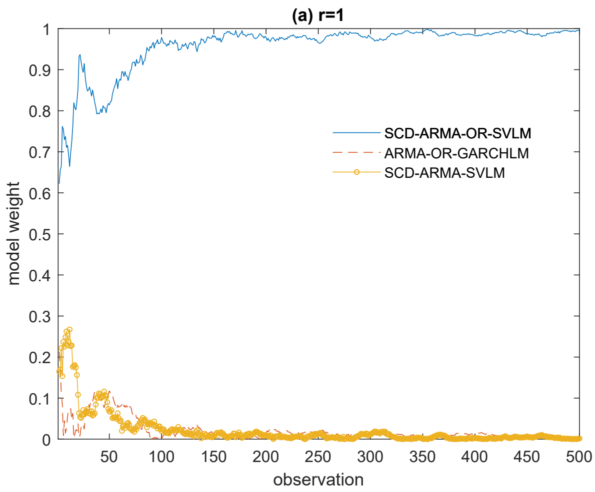

(a)

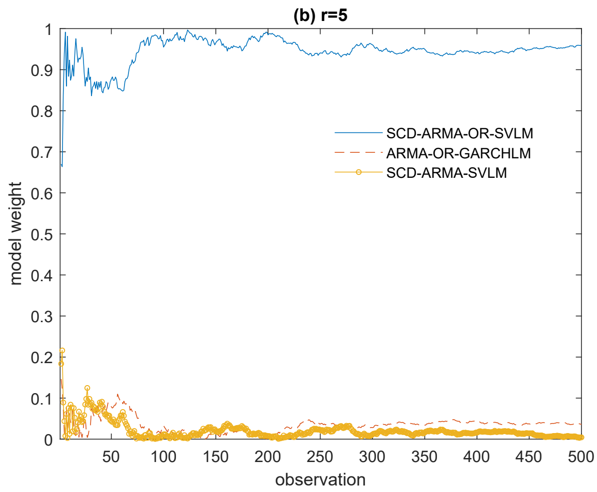

(b)

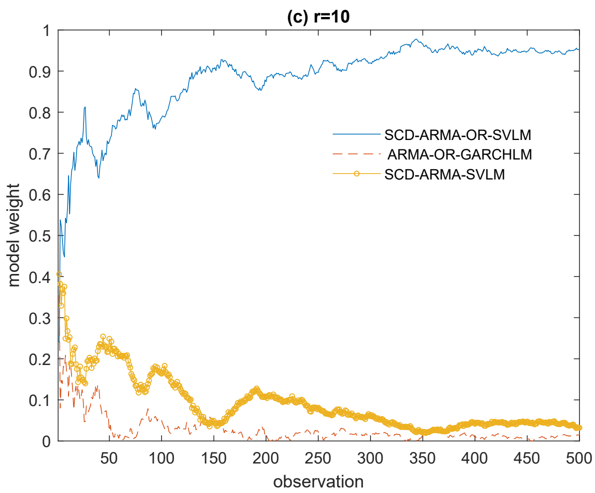

(c)

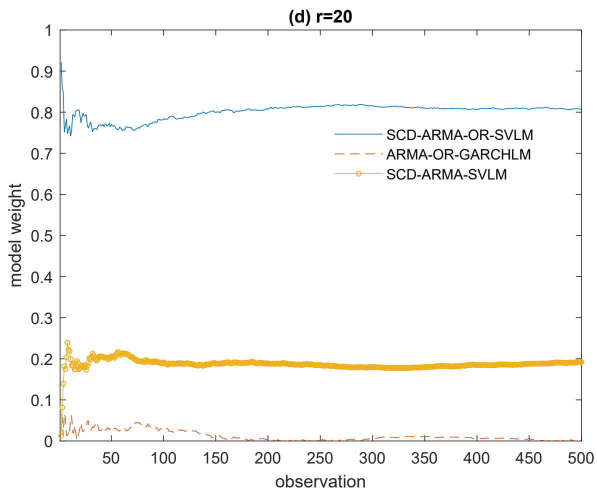

(d)

Figure 4: Empirical results. Evolution of model weights in the three-model pool of JNJ stock $r = 1, 5, 10$ and 20-step ahead predictive densities.

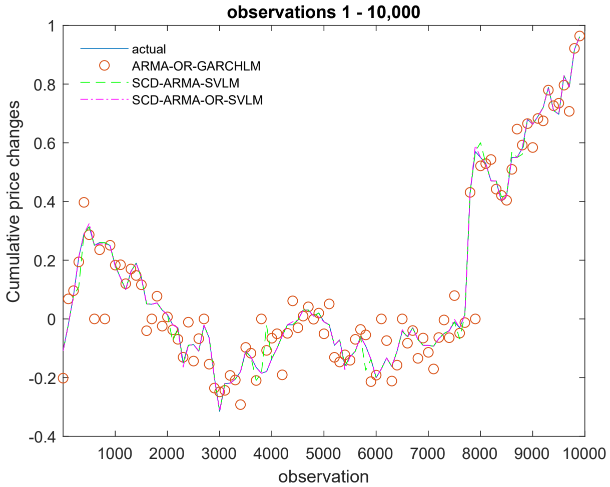

(a) Actual and predicted time series of cumulative price changes for competing models ($t = 1000, ..., 10000$).

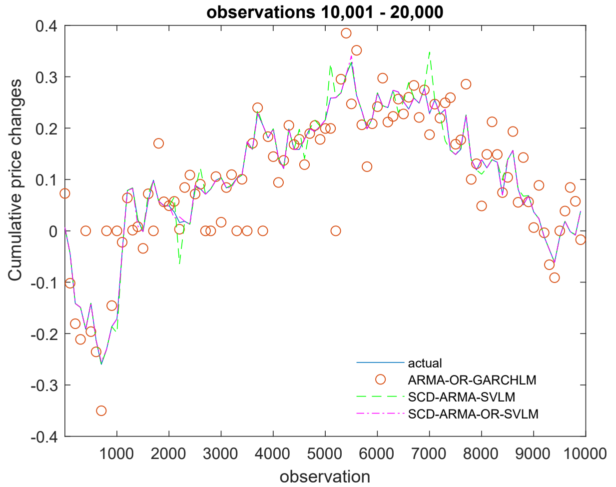

(b) Actual and predicted time series of cumulative price changes for particular models ($t = 10001, ..., 20000$).

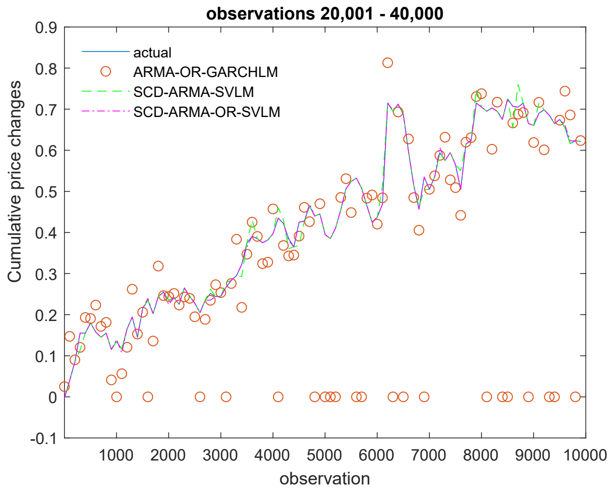

(c) Actual and predicted time series of cumulative price changes for particular models ($t = 20001, ..., 40000$).

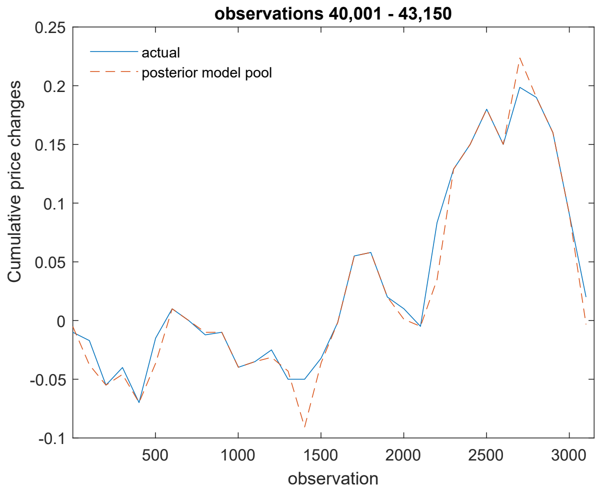

(d) Actual and predicted time series of cumulative price changes for the optimal pool of models ($t = 40001, ..., 43150$).

Figure 5: Empirical results: Actual and predicted cumulative price changes for $r = 1$- step ahead transactions (JNJ stock).

Table 1: Empirical data.

Data set	Description	Number of adjusted durations
I	Boeing stock from December 1 to December 5, 2008.	54794
II	Johnson and Johnson stock from October 4 to October 8, 2010	43150
III	Starbucks stock from July 25 to July 29, 2011.	38475
IV	IBM stock traded from January 21 to January 25, 1991.	6413
V	IBM stock traded from December 13 to December 17, 1999.	23528
VI	General Electric stock from December 1 to December 5, 2003.	44633

Notes: The transaction data sets were obtained from the Trades and Quotes (TAQ) database of the New York Stock Exchange (NYSE). The stocks were recorded on normal trading days and hours.

Table 2: JNJ stock: Intervals of price changes, response categories, and frequencies of price changes falling in each category.

Price change intervals (cents)	< -2	$[-2, -1)$	$[-1, 0]$	$(0, 1]$	$(1, 2]$	> 2
Responses (y_t)	1	2	3	4	5	6
Frequencies	136	1081	25046	15589	1168	130

Table 3: Empirical estimation results for the JNJ stock data.

Parameters/Models	SCD-ARMA-OR-SVLM	SCD-ARMA-SVLM
ρ	0.910	0.810
	(0.013)	(0.024)
τ_1	0.315	0.282
	(0.024)	(0.043)
γ_1	0.017	0.021
	(0.005)	(0.007)
α_1	0.917	0.833
	(0.012)	(0.025)
α_2	-0.315	-0.410
	(0.030)	(0.053)
ω_1	1.27	1.10
	(0.014)	(0.022)
ω_2	0.810	0.721
	(0.017)	(0.011)
ω_3	0.440	0.392
	(0.021)	(0.019)
σ_λ^2	0.012	
	(0.0013)	
σ_v^2	0.151	0.133
	(0.03)	(0.055)
ζ_1	-0.150	
	(0.011)	
ζ_2	-0.082	
	(0.004)	
ζ_3	0.004	
	(0.001)	
ζ_4	0.133	
	(0.017)	
ζ_5	0.162	
	(0.010)	

Reported posterior means. Standard errors in parenthesis.

Table 4: Empirical results. Relative log predictive scores (LPS) of competing models for the JNJ stock data.

Models	$r = 1$	$r = 5$	$r = 10$	$r = 20$
SCD-ARMA-OR-SVLM	**1.44**	**1.73**	**1.54**	**1.20**
SCD-ARMA-SVLM	1.32	1.70	1.52	1.18
ARMA-OR-GARCHLM	0.87	0.72	0.67	0.53

Notes: Values in bold indicate that the corresponding models have the best forecasting performance. The benchmark model is the AR-OR model.

Table 5: Empirical results. Relative root mean squared forecast errors (RMSFEs) of competing models for the JNJ stock data.

Models	$r = 1$	$r = 5$	$r = 10$	$r = 20$
SCD-ARMA-OR-SVLM	**0.62***	**0.55***	**0.67***	**0.63***
SCD-ARMA-SVLM	0.70*	0.72*	0.75*	0.68*
ARMA-OR-GARCHLM	1.32	1.35	1.44	1.61

Notes:
1) Values in bold indicate that the corresponding models have the best forecasting performance. The benchmark model is the AR-OR model.

2) * indicates that the respective model performs significantly different from the AR-OR model at the 5% level, based on the p-values of the Giacomini-White (2006) test. These values are given in Table 6.

3) Underlined numbers indicate that the corresponding models are included in the Model Confidence Set (Hansen et al., 2011). The confidence level for MCS is 95%. The MCS p-values are given in Table 7.

Table 6: Empirical results (JNJ stock). Giacomini-White (2006) test for the RMSFEs results.

Model	$r = 1$	$r = 5$	$r = 10$	$r = 20$
SCD-ARMA-OR-SVLM	0.003	0.001	0.00	0.00
SCD-ARMA-SVLM	0.005	0.007	0.009	0.022
ARMA-OR-GARCHLM	0.054	0.060	0.089	0.155

Notes: This table reports the p-values from the two-sided Giacomini-White (2006) test. An entry of 0.00 means that the p-value was less than 0.001. The benchmark model is the AR-OR model.

Table 7: Empirical results (JNJ stock). Model Confidence Set (MCS) of Hansen et al. (2011).

Model	$r = 1$	$r = 5$	$r = 10$	$r = 20$
SCD-ARMA-OR-SVLM	0.445*	0.513*	0.313*	0.275*
SCD-ARMA-SVLM	0.210*	0.303*	0.189*	0.203*
ARMA-OR-GARCHLM	0.001	0.001	0.00	0.00

Notes: This table reports the p-values from the MCS approach. An asterisk denotes that the model is included in the 95%-MCS. An entry of 0.00 means that the p-value was less than 0.001.

Table 8: Summary table for the LPS results.

Data	$r=1$	$r=5$	$r=10$	$r=20$
I	M1	M1	M1	M1
II	M1	M1	M1	M1
III	M1	M1	M1	M1
IV	M1	M1	M1	M1
V	M1	M1	M1	M1
VI	M1	M1	M1	M1

Notes: This table reports the best forecasting model for each data set and for each horizon r, based on the LPS values. M1 stands for the SCD-ARMA-OR-SVLM model, M2 stands for the SCD-ARMA-SVLM model, and M3 stands for the ARMA-OR-GARCHLM model.

Table 9: Summary table for the RMSFE results.

Data	$r=1$	$r=5$	$r=10$	$r=20$
I	M1	M1	M1	M1
II	M1	M1	M1	M1
III	M1	M1	M1	M1
IV	M1	M1	M1	M1
V	M1	M1	M1	M1
VI	M1	M1	M1	M1

Notes: This table reports the best forecasting model for each data set and for each horizon r, based on the RMSFE values. M1 stands for the SCD-ARMA-OR-SVLM model, M2 stands for the SCD-ARMA-SVLM model, and M3 stands for the ARMA-OR-GARCHLM model.

Table 10: Summary table for the MCS results.

Data	$r=1$	$r=5$	$r=10$	$r=20$
I	M1, M2, M3	M1	M1	M1
II	M1, M2	M1,M2	M1,M2	M1,M2
III	M1, M2	M1, M2	M1	M1
IV	M1	M1	M1	M1
V	M1, M2	M1	M1	M1
VI	M1	M1	M1	M1

Notes: This table reports the Model Confidence Set for each data set and for each horizon r, based on the MCS p-values. M1 stands for the SCD-ARMA-OR-SVLM model, M2 stands for the SCD-ARMA-SVLM model, and M3 stands for the ARMA-OR-GARCHLM model.

Table 11: Summary table for the Geweke-Amisano results.

Data	$r = 1$	$r = 5$	$r = 10$	$r = 20$
I	M1(0.92), M2	M1(0.95), M2	M1(0.94), M2	M1(0.91), M2
II	M1(0.90), M3	M1(0.98), M3	M1(0.98), M2	M1(0.99), M2
III	M1(0.97) , M3	M1(0.91), M2	M1(0.97), M2	M1(0.97), M2
IV	M1(0.90), M3	M1(0.94), M2	M1(0.96), M2	M1(0.92), M2
V	M1(0.93), M3	M1(0.96), M2	M1(0.92), M2	M1(0.92), M2
VI	M1(0.95), M3	M1(0.92), M2	M1(0.95), M2	M1(0.96), M2

Notes: This table reports the model that received most of the weight for each data set and for each horizon r. The numbers in parentheses are the weights for the pool of models optimized over all the last 500 observations. M1 stands for the SCD-ARMA-OR-SVLM model, M2 stands for the SCD-ARMA-SVLM model, and M3 stands for the ARMA-OR-GARCHLM model. The model in red is the model that received the second largest weight.

Table 12: Summary table for the trading rule (all trades).

Data	$r = 1$	$r = 5$	$r = 10$	$r = 20$
I	2.71%	3.44%	3.12%	3.00%
II	1.89%	2.52%	2.97%	3.31%
III	2.35%	2.44%	3.05%	3.20%
IV	2.82%	3.12%	3.55%	3.50%
V	3.27%	3.61%	3.84%	4.13%
VI	3.82%	4.10%	4.10%	4.22%

Notes: This table reports the percentage by which the average excess return from the trading rule for the optimal pool of models exceeds that for the SCD-ARMA-SVLM model.

Table 13: Summary table for the trading rule (restricted trades).

Data	$r = 1$	$r = 5$	$r = 10$	$r = 20$
I	0.44%	0.72%	1.33%	1.52%
II	0.81%	1.05%	1.28%	1.32%
III	1.23%	1.35%	1.44%	1.51%
IV	0.78%	1.06%	1.12%	1.48%
V	1.52%	1.90%	2.31%	2.40%
VI	1.44%	1.60%	1.99%	1.90%

Notes: This table reports the percentage by which the average excess return from the trading rule for the optimal pool of models exceeds that for the SCD-ARMA-SVLM model, when we restrict to three trades per day.