Structural breaks in point processes: With an application to reporting delays for trades on the New York stock exchange *

Jonas Andersson and Jan-Magnus Moberg

Norwegian School of Economics and Business Administration
Department of Finance and Management Science

Abstract

In this paper some methods to determine the reporting delays for trades on the New York stock exchange are proposed and compared. The most successful method is based on a simple model of the quote revision process and a bootstrap procedure. In contrast to previous methods it accounts for autocorrelation and for variation originating both from the quote process itself and from estimation errors. This is obtained by the use of prediction intervals. The ability of the methods to determine when a trade has occurred is studied and compared with a previous method by Vergote (2005). This is done by means of a simulation study. An extensive empirical study shows the applicability of the method and that more reasonable results are obtained when accounting for autocorrelation and estimation uncertainty.

1 Introduction

Many studies within the field of market microstructure apply data from the Trades and Quotes (TAQ) database for empirical research. This empirical work is often adaptations of theoretical microstructure models. A critical factor for many of these studies is the ability to identify the quotes in effect at the time of a trade. The identification of this prevailing quote is an important element in determining the information content of trades, the order imbalance and inventory accumulation of liquidity providers, the price impact of large trades, the effective spread, and many other related questions.

The most widely used algorithm to determine the prevailing quote was developed by Lee and Ready (1991). They studied quote revision frequencies around isolated trades in order to identify the prevailing quote. The intuition behind this approach is that although some quote revisions are caused by the arrival

*We would like to thank Jostein Lillestøl and Luc Bauwens for valuable comments.

or cancellation of limit orders, many are triggered by trades. To avoid any contamination from neighboring trades, only isolated trades where studied.[1]. They found, by studying TAQ data from 1988, that a large portion of the quotes were registered ahead of trades. One possible explanation for this could be that if the specialist assistant was faster in recording a quote revision than the floor reporter in recording a trade, the corresponding quote update could be recorded before the trade that triggered it. Lee and Ready (1991) showed that this problem could be mitigated by using a time-delayed quote which, in the case of 1988 data, was the quote in effect 5 seconds before the the trade time stamp. Several studies, among others Ball and Chordia (2001), Busse and Green (2002), Chan et al. (2002), Chordia et al. (2001, 2002), Easley et al. (2001) , Edelen and Gervais (2003), Engle and Patton (2004), Huang and Stoll (2001), Kryzanowski and Zhang (2002), Nyholm (2003), Schultz (2000), Stoll (2000) and Venkataraman (2001) have used the same time-delay as Lee and Ready (1991) even though Lee and Ready explicitly mention that another delay might be appropriate for other time periods.

Henger and Wang (2006) illustrates that the timing specifications of trades and quotes ultimately can change empirical outcome. They find that using a 1-second quote delay is appropriate for their sample of NYSE stocks during 1999, and demonstrate the significance of the timing specifications of economic variables using the Huang and Stoll (1997) spread decomposition model. Using a 5-second quote delay result in severe biases in the estimated parameters. Piwowar and Wei (2006) find that the effective spread estimates are sensitive to trade-quote matching algorithms. In particular, Lee and Ready's 5 second algorithm can overestimate effective spreads for active stocks.

The NYSE records transactions via the Consolidated Trade System (CTS) and quotes via the Consolidated Quote System (CQS). The TAQ database is an extraction of these systems.[2] The way trades and quotes reach CTS and CQS has changed over the years. The actual trade was until 24 July 2001 reported either through a Display Book or by floor reporters.[3] After this the floor reporter position was eliminated and all trading is now done directly through the Display Book. Hasbrouck et al. (1993) studied 144 stocks on the NYSE during the 5 first trading days in 1990 and found that Display Book reported trades have much smaller reporting delay than trades reported by floor reporters.

In Vergote (2005), an important issue is highlighted, namely that the reporting delay of trades on the New York stock exchange (NYSE) varies significantly between stocks and also over time. Strong arguments are given, in the same paper, that the commonly used delay of 5 seconds given by Lee and Ready (1991) is too rigid to apply to all stocks and all periods of time. As in Hasbrouck et al. (1993), Vergote (2005) separates between Display Book reported trades and

[1] They defined an isolated trade as a trade were there are no other trades within a 2-minute window centered on that trade

[2] See Hasbrouck et al. (1993) for more detailed reporting procedures at the NYSE

[3] The Display Book receives and displays orders to specialists and provides a mechanism to execute and report transactions. The floor reporter is an NYSE employee who stands by the specialist on the trading floor.

trades that are not reported via the Display Book.

This paper studies the arrival of quote revisions posted by designated specialists at the New York Stock Exchange (NYSE). We start by assuming that quote revisions arrive according to a homogeneous Poisson process until a trade occurs. When a trade occurs the intensity of quote revisions increase and in this paper we study this structural breakpoint in the assumed data-generating process applying different methods. The first methods we study, PIa and PIb, are based on the calculation of prediction intervals of number of quotes in a given second, and the latter accounts for estimation uncertainty. We also consider two methods based on the index of dispersion (VTa and VTb). In addition to these methods we also propose two methods were we calculate prediction intervals based on the assumption that quotes follow an AR(1)-process. As with PIa and PIb, we present one method which account for estimation uncertainty (PIAR1b) and one that ignores this uncertainty (PIAR1a).

The methods proposed in this paper are first studied and compared with the method presented in Vergote (2005) in a simulation study. This study shows that PIa, PIb, PIAR1a and PIAR1b performs reasonably well in detecting structural breakpoints. This result is robust to even small increases in the intensity after the breakpoint. The method proposed by Vergote, as well as the methods based on the index of dispersion, fails in detecting structural breakpoints as the increase in intensity decrease.

The applicability of the methods are illustrated with data from the New York stock exchange. Applied to the TAQ dataset it seems like the PIAR1b method performs most reasonable. Contrary to a homogeneous Poisson process this method model the quote intensity at each time point as an AR(1)-process, accounting for estimation uncertainty. An irregularity in the quote intensity is observed for some stocks were the quote intensity increase steadily up to the time of the trade. This irregularity is better captured assuming an AR(1)-process. Another peculiar result is that the intensity in quote revisions systematically increase one second prior to the reported trades for most stocks. A likely reason for this observation is the way timestamps are rounded in the TAQ dataset.

The remainder of this paper is organized as follows. In the next section we give a brief description of the data. Section 3 present the methods studied in this paper. In Section 4 we investigate the methods ability to detect the time of the trade under some different data generating processes (DGP's). We also compare this ability with the method of Vergote (2005). Section 5 exemplifies the method on a sample of stocks from the New York stock exchange during April 2002 and 2006 and Section 6 concludes.

2 Data

We study 225 stocks during April 2002 and 2006. The trades and quotes are taken from the TAQ dataset. The reason for only studying one month of data is the size of the dataset. By studying data after the floor position at NYSE was eliminated we do not need to consider the difference between trades that are

Display Book reported and those that are reported by floor reporters. We select three groups according to market capitalization. 75 stocks from each of the three indices; S&P 500, S&P MidCap 400 and S&P SmallCap 600, were studied in this paper. Only trades and quotes that followed certain conditions were selected.[4] Daily descriptive statistics of our dataset are presented in table 3. Trading and quoting activity is as expected largest for the stocks with highest market capitalization and lowest for the stocks with smallest market capitalization. In the empirical study we only consider isolated trades, defined as trades for which there are no other trades within a 40-second window centered around the trade. This time interval is defined as t = [-20, 20] and the isolated trade is reported at t = 0. The reason for only studying isolated trades is to remove any confounding effects between trades that are closely clustered. Lee and Ready (1991) define isolated trades applying a 2-minute window, but this is not suitable for our sample due to the increased trading activity. The number of isolated trades, in percentage of the total number of trades is low, especially during 2006. Less than 2% of all trades in the group with large market capitalization stocks were defined as isolated trades during the sample from 2006. The main reason for this is the increasing trading activity during the last couple of years. The number of quote revisions have also increased significantly, partially due to the introduction of auto quoting. A new quote revision were reported almost every 2 seconds for the 75 stocks in the S&P 500 Index during April 2006.

3 The methods

3.1 Assumed data generating process and notation

Just as in Vergote (2005) we assume that the quote revisions arrive according to a homogeneous Poisson process until a trade occurs. Assume that we have an observation interval $[0, T] \in \mathcal{Z}_+$ for each isolated trade and let Y_t be the number of quote revisions in $[0, t]$ where $t \leq T$. Then, as long as no trade have occurred

$$Y_t \sim Po(\lambda t) \tag{1}$$

where λ is a positive constant. In our methods we will use the number of quotes in a given second before, at or after the registered trade. Because of this we introduce $X_t = Y_t - Y_{t-1}$ with corresponding sample quantities $x_t = y_t - y_{t-1}$. The stochastic variable X_t then has the probability function

$$P(X_t = x_t) = e^{-\lambda} \frac{\lambda^{x_t}}{x_t!} \tag{2}$$

Whether one chooses to work with the momenta-nous or accumulated number of quotes is mainly a matter of preference of one graphical presentation over another.

[4]Same as those specified in Vergote (2005). Note that all trades were Display Book reported in the two periods we study.

Finally, a note on the vocabulary used in the sequel of this section. In order not to mix up the method proposed with the interpretation of the empirical results, we will here use the term *break* instead of *trade* when we refer to a shift in the intensity of the process.

3.2 Method based on prediction intervals of number of quotes in a given second (PIa and PIb)

The perhaps most straightforward approach is to directly consider the frequencies of quotes, X_t, in a given second before or after a registred trade. A histogram of such data can be seen in Figure 1.

Given a sample of such $Po(\lambda)$-distributed data, $x_1, x_2, ..., x_n$, it is well known that the maximum-likelihood estimator is the sample mean

$$\hat{\lambda}_{ML} = \frac{\sum_{t=1}^{n} x_t}{n} \tag{3}$$

A large-sample approximation to the distribution of $\hat{\lambda}_{ML}$ is given by

$$N\left(\lambda, \frac{\lambda}{n}\right) \tag{4}$$

In using the quote revision data to determine whether a break has occurred we need to account for two types of variation. Those are the inherent variation in the quote arrivals themselves, given by (2), and estimation errors, given by (4). The procedure is based on prediction intervals and the estimation uncertainty can be taken into account through a simple resampling procedure. We will later investigate the method both with and without accounting for estimation uncertainty (named PIb and PIa respectively).

Given that no break has occurred at time $t+1$, we would like to construct an interval based on the information in $X_1, X_2, ..., X_t$ which covers X_{t+1} with probability $1-\alpha$. To do this we would, ideally, like to know the exact distribution of the prediction based on $X_1, X_2, ..., X_t$. The functional form of this distribution is unknown. However, such an interval can be obtained by the following parametric bootstrap procedure.

1. Estimate λ by $\hat{\lambda}_{ML}$ with $x_1, x_2, ..., x_t$
2. Generate $\hat{\lambda}_1, \hat{\lambda}_2 ..., \hat{\lambda}_B$ from (4) with $\lambda = \hat{\lambda}_{ML}$.
3. Generate $\hat{X}_{t+1}^{(i)}$ from $Po(\hat{\lambda}_i)$, $i = 1, 2, ..., B$.
4. Estimate the endpoints of the interval by the $[B\frac{\alpha}{2}]$'th and the $[B(1-\frac{\alpha}{2})]$'th order statistics of $\hat{X}_{t+1}^{(1)}, \hat{X}_{t+1}^{(2)}, ..., \hat{X}_{t+1}^{(B)}$.

If the observation X_{t+1} lies above this interval we conclude that a break has occurred at time t. Even though this can be seen as an "estimator" of the delay we choose to present the results graphically. An example of this is given in Figure 1 for the ABM-stock. When it is done for many stocks, however, we automatize the procedure by choosing the second before the first observation falling above the prediction interval.

3.3 Method based on prediction intervals from AR(1) model (PIAR1a and PIAR1b)

There are obviously the possibility that even though no break, as defined by e.g. a trade, has occurred the underlying qoute-process is not necessarily a homogeneous Poisson-process. If $X_1, ..., X_t$ are positively autocorrelated the prediction intervals in the last section will be too narrow, see Section 5 for empirical evidence on this. This can in practice usually be modeled by an AR(1)-specification

$$X_t = \phi_0 + \phi_1 X_{t-1} + \varepsilon_t$$

where $\{\varepsilon_t\}$ is a normally distributed white noise process with standard deviation σ. It is well know that a one-step-ahead predictive distribution based on this model is given by

$$\hat{X}_{t+1}|X_t = x_t \sim N(\phi_0 + \phi_1 x_{t-1}, \sigma^2) \quad (5)$$

In order to account for the parameter uncertainty, which is particularly important since we have small samples, can be done in the same manner as the algorithm described in the previous section. See e.g. Pascual et al. (2001) for more on how to account for parameter uncertainty in prediction intervals for ARMA-models.

1. Estimate ϕ_0 and ϕ_1 by maximum likelihood with $x_1, x_2, ..., x_t$ and extract the residuals $e_t = x_t - \hat{\phi}_0 - \hat{\phi}_1 x_{t-1}$ for $t = 1, ..., t$
2. Create a bootstrap sample, $e_1^*, e_2^*, ..., e_t^*$ from these residuals by drawing t values with replacement.
3. Use these residuals and the estimated parameter values to generate a new artificial time series $x_1^*, x_2^*, ..., x_t^*$.
4. Estimate the parameters of the model based on this artificial time series.
5. Use these parameters and the observation x_t from the original sample to produce a bootstrapped prediction $\hat{x}_{t+1}^*$.
6. Repeat 1 to 5 B times in order to produce B bootstrapped predictions.
7. Estimate the endpoints of the interval by the $[B\frac{\alpha}{2}]$'th and the $[B(1-\frac{\alpha}{2})]$'th order statistics of these bootstrapped predictions.

Also this method can be applied with or without the correction for parameter uncertainty.

3.4 Methods based on the index of dispersion (VTa and VTb)

As shown in Karlis and Xekalaki (2000), if the deviation from the Poisson-distribution is manifested as a difference in the first two central moments, which are equal for the Poisson distribution, one can successfully use the variance test (VT) statistic

$$VT_a = (n-1)\frac{S^2}{\overline{X}} \tag{6}$$

where $\overline{X}$ is the sample mean, $S^2 = \sum_{i=1}^{n}(X-\overline{X})^2/(n-1)$, the sample variance and n the number of observations. VTa can be shown to be asymptotically $\chi^2(n-1)$-distributed. Also here the estimation of the breakpoint is made by sequentially, for $t = 10, 11, ..., 40$, test whether the t'th observation belongs to the same distribution as the first $t-1$.

However, it is known that the $\chi^2(n-1)$-approximation to the sample distribution is not very reliable. Therefore, the alternative statistic

$$VT_b = \sqrt{\frac{n-1}{2}}\left(\frac{S^2}{\overline{X}} - 1\right) \tag{7}$$

will also be studied. VTb is asymptotically $N(0,1)$ under the null hypothesis that $X \sim Poisson$. This is also the test which came out best in Karlis and Xekalaki (2000) for alternatives where the mean and variance was unequal.

These methods are however not working when the ratio of the variance and the mean maintains the value one also after the breakpoint. Also, since we are mainly interested in *increases* in the intensity, these tests, which also reacts to decreases, has a theoretical drawback. We will, nevertheless, consider them in the simulation study of Section 4.

3.5 Vergote's method (VER)

Since we will include the method by Vergote (2005), a short review of it is appropriate here. In contrast to the method of this paper it exploits only the variation in the parameter estimator of λ when the number of observations to estimate it increases. If n is the total number of isolated trades he defines $\hat{\lambda}_s = Y_s/n$ as an estimator of λ. Furthermore, a partial mean of λ's for all observations from time 0 to s (in the notational convention of this paper) $\hat{\lambda}_s^p = \sum_{i=0}^{s}\hat{\lambda}_s/(s+1)$ is used as a reference. The conclusion that a trade has occurred is made if the quantity $\lambda_s - 1.2\lambda_{s-1}^p$ is significantly larger than zero. λ_s and λ_{s-1}^p are the population values corresponding to $\hat{\lambda_s}$ and $\hat{\lambda}_{s-1}^p$. The test is performed by using the asymptotically standard normal quantity

$$\frac{\hat{\lambda}_s - 1.2\hat{\lambda}_{s-1}^p}{se}$$

where

$$se^2 = \left(\frac{\hat{\lambda}_s}{n} + \frac{1.44}{ns^2} \sum_{i=0}^{s-1} \hat{\lambda}_i \right)$$ [5]

Even though the assumed data generating process of this method is the same as in Vergote (2005) there is a significant difference between the methods. Our approach, based on prediction intervals, accounts for the natural variation in Y_{t+1} and not only the variation originated from parameter estimation which is the case of the former method. Another notable difference that can be mentioned is that the tuning parameter 1.2 in the latter method is arbitrary and is motivated in Vergote (2005) by "Judging from the quote revision distributions, λ_s appears to vary within the interval before the trade". In the proposed method no such parameter exist.

In Section 4 we will study the performance of the different methods.

3.6 Summary of the methods

We have in the previous subsections presented 4 methods with a few subclasses. These are

1. Prediction interval based on the assumption of quotes being a sequence of independent Poisson distributed stochastic variables.
   (a) Ignoring estimation uncertainty (PIa).
   (b) Accounting for estimation uncertainty by a bootstrap procedure. (PIb)
2. Prediction interval based on the assumption of quotes being an AR(1)-process. Also here we have two alternatives.
   (a) Ignoring estimation uncertainty. (PIAR1a)
   (b) Accounting for estimation uncertainty by a bootstrap procedure. (PIAR1b)
3. Using the index of dispersion,$VT = S^2/\overline{X}$, and evaluate whether this deviates from one, i.e. from the property of the Poisson distribution.
   (a) Using the asymptotic distribution of $(n-1)VT$ which is $\chi^2(1)$ under the Poisson-distribution. (VTa)
   (b) Using the asymptotic distribution of $\sqrt{(n-1)/2}(VT-1)$ which is $N(0,1)$ under the Poisson-distribution. (VTb)
4. The method by Vergote (2005).

[5] The formulas look slightly different from Vergote (2005) since our observation interval runs between 1 and 40 while his runs from -10 to 10.

4 Performance of the methods

Some simulation experiments are performed in order to study the power of the method to detect the correct time of the breaking point given two different data generating processes. Simulations as well as the empirical analysis in the next section is performed by the statistical package R (R Development Core Team, 2005).

In the first simulation the process of the accumulated number of quote revisions in $[0, t]$, Y_t, is presumed to be a Poisson process with intensity 10 until the break occurs at time point 19 when it increases to λ_2 and decreases back to 10 at time point 24. λ_2 is varied between 20 and 100. We specify the time-varying λ as

$$\lambda_t = \begin{cases} \lambda_2 \text{ if } t = 19, 20, 21, 22, 23 \\ 10 \text{ otherwise} \end{cases}$$

Figure 2 illustrates the λ_t and a typical frequency histogram of events (trades).

Figure 3 presents the results in the form of boxplots of the resulted estimated breakpoints for the seven investigated methods. As can be seen from them, the Monte Carlo medians of the estimated break point is very close to the true value 19 when the intensity jumps to 30. This is true for all seven methods. However, when the jump in the intensity decreases the performance of the VT-methods as well as Vergote's method deteriorates while the proposed methods PIa, PIb, PIARa and PIARb stays reasonably on track detecting the break point. It should also be noted that the larger variation in the PIAR1-methods is explained by the fact that more parameters are estimated and that it in this simulation study is not coming fully to its right since there is no autocorrelation in the simulated DGP's. The sometimes slight deviation from a median of 19 for the proposed methods is explained by the fact that occasionally the prediction interval based on observations $Y_1, ..., Y_t$ will cover Y_{t+1} even though the intensity has jumped at $t + 1$ because of a trade. On the other hand sometimes the prediction interval will **not** cover Y_{t+1} even if the intensity has stayed the same. The sign of this deviation depends on the relationship of the probabilities for these events to occur.

The reason for failure of the index of dispersion based methods in this simulation exercise is explained by the fact that the first and second order sample moments are changing close to proportionally to each other. The success of these methods for λ_2 equal to 30 is actually somewhat surprising.

The next simulation is supposed to be more realistic and the choice of the lambda-function is made based on the guidance of the real data that we will consider in Section 5. λ_t is now defined to make a jump at a certain time point (time of trade), make another jump at the next time point and then decay exponentially back to its original level. Specifically

$$\lambda_t = \begin{cases} 10 \text{ if } t \leq 18 \\ 10 + (\lambda_2 - 10)/2 \text{ if } t = 19 \\ \exp(0.1(19 - t) + \ln(\lambda_2 - 10)) + 10 \text{ if } t > 19 \end{cases}$$

Figure 4 shows the variation of λ and a typical frequency histogram of events (trades).

The results shown in Figure 5 and in Table 2 indicate that the shape of λ_t after the breakpoint might not be of great importance for this method. This is good, as an indication of robustness. However, it also indicates that information after the breakpoint is not very much exploited. This is obviously the usual trade-off between fully and partially parametrized models. A possible future line of research could be to parametrize the shape of λ_t after the breakpoint in order to exploit this information better in the detection of it. The breakpoint will then itself be a parameter to estimate in a more classical setting than with the method proposed in this paper.

The behaviour of the index of dispersion based methods is again explained by the close-to-constant ratio of the sample variance and sample mean.

5 Empirical example

In this, empirical, section we study structural breaks in the quote intensity around trades using the methods introduced in Section 3. We study NYSE trades and quotes obtained from the TAQ dataset during April 2002 and 2006. For each of the two years we divide the sample stocks into three groups according to market capitalization. The large, medium and small groups each include 75 stocks from S&P 500, S&P 400 MidCap and S&P 600 SmallCap respectively. By doing this we can study whether there are any differences in breakpoints not only across time but also with respect to market capitalization. This section is only based on isolated trades and the quote revisions around these trades. In other words, we only account for a part of the total amount of trades and quotes. This is, as already mentioned, to avoid any confounding effects from neighboring trades. The percentage of isolated trades can be increased by reducing the 40-window used to define an isolated trade. Preliminary results indicate that the quote intensity does not alter significantly when we decrease this to a 20-second window. We therefore assume that the quote intensity is the same for all trades.

The reported breakpoint for each stock is presented in Table 4 and Table 5 (for 2002 and 2006 respectively). Breakpoints are reported for each of the methods introduced in Section 3, and the number of quotes are accumulated for each time point and reported in the intervals $[< 16, 17, ..., 22, > 22]$. Trades are reported at $t = 21$. Compared to the results in the simulation study the empirical results are more similar to those where $\lambda_2 = 30$ in the sense that no breakpoints are reported at $t > 21$ in the empirical study. The reason for this is that all methods detect the increase in quote intensity at (or before) $t = 21$ in the real data. The jump in intensity at this time point is quite large, often more than 100% of average intensity up to $t = 20$. As in the simulation study, this indicates that the shape of λ_2 is of little importance to the result. The breakpoints for PIa and PIb are distributed in the interval from t=[10,21] throughout the sample with a large part of the breakpoints reported at $t < 16$. This result might indicate that the quote intensity does not follow an homogeneous Poisson process for all stocks, independent of whether we account for estimation uncertainty or not. If the quote intensity at each time point, $X_1, ..., X_t$, are positively autocorrelated the prediction intervals will be too narrow. The two methods based on prediction intervals from an AR(1)-model accounts for this positive autocorrelation. PIAR1a, which does not account for estimation uncertainty, produce similar results to PIa and PIb. PIAR1b most frequently report breakpoints at one second before the reported trade, particularly during 2006 and for stocks with high market capitalization. Under PIAR1b, 99% of all stocks in the large group were reported with a breakpoint at $t = 20$ in 2006. For medium and small market cap stocks this number was 92% and 87% respectively. These numbers were somewhat lower during April 2002.

The methods based on the index of dispersion (VTa and VTb) reports structural breaks that are more diversified up until the trade is reported at $t = 21$. These results supports the findings from PIa and PIb that the quote intensity does not follow a homogeneous Poisson process for all stocks. It does not seem to be any systematic difference between the two asymptotic distributions applied in these methods. As already mentioned in subsection 3.4 these models does not only react to increase in intensity, but also decreases. Natural variation can therefore cause these methods to report a breakpoint when there is a drop in the quote intensity. This makes them less reliable when applied here since the main intuition in this empirical study is that the quote intensity should increase at the time of a trade.

The methods presented so far are based on prediction intervals and accounts for natural variation in the quote intensity. Vergote's method, identifies a breakpoint by comparing the quote intensity, defined as $\hat{\lambda}_s$ and the partial mean of the quote intensity $\hat{\lambda}_s^p$. A parameter value, which is set arbitrary, is used to determine breakpoints. This method is efficient in capturing the significant increase in quote frequency that takes place at the same time and in the second before the trade is reported, but does not capture variation in the Poisson process beyond the natural variation very well. As a result of this, Vergote's method report breakpoints close up to the trade at $t = 20$ and $t = 21$. In 2006 it reports similar results as the PIAR1b method for large market cap stocks. 96% of the breakpoints in this group are reported at $t = 20$. Contrary to the method proposed by Vergote (2005) PIAR1b does not assume a homogeneous Poisson process.

Figure 6 and Figure 7 present the results graphically in a boxplot. The median for the stocks in the large group is close to $t = 20$. This result is consistent across the methods we study. For PIa, PIb and PIAR1a the interquartile range is considerably smaller for stocks in the large market cap stocks compared to that of the stocks in the medium and small group. Similar results is obtained when studying the earliest non-outlier observation. The results obtained from PIAR1b are somewhat similar to that of Vergote's method, in the sense that the breakpoints are centered around the time point of the trade. The main difference is that VER more frequently report breakpoints at $t = 21$.

The structural breakpoints are most frequently reported one second ahead of the isolated trades. These results are not consistent with the intuition of a constant intensity up to the reported trades at $t = 21$ for all stocks. Our results indicate however that most trades have a relatively constant intensity up to one second prior to the trade. At $t = 20$ all methods report a structural breakpoint for a large part of the stocks in the sample. This result is consistent with the findings in Moberg (2007). They find that the quote intensity increase one second before the trade is reported, i.e. at $t = 20$. This result is most likely due to a rounding in the reported timestamps in the TAQ dataset.

If all quote revisions resulting from the isolated trades were reported no more than one second ahead of the trades we would expect a fairly constant quote intensity up to $t = 20$. This is not the case for some stocks, especially those with small market cap. Interestingly, the intensity for some stocks increase steadily up to the second ahead of the reported trade. Figure 8 illustrates this effect. This abnormality is best captured in the methods accounting for autocorrelation in the quote intensity. This effect is likely due to an imperfect trade reporting system. If the timestamps in the Consolidated Trade System and the Consolidated Quote System, which stores all trades and quotes respectively, are not fully synchronized this might result in trades reported with a lag. Quotes that are triggered by these trades could then be reported ahead of the trades that triggered them. The effect can also occur if reporting is synchronized but trades and quotes are reported out of proper sequence. This might occur to a different extent for different market makers.

6 Conclusion

In this paper we propose some methods to determine reporting delays for trades at the New York stock exchange. We start by assuming that quote revisions arrive according to a homogeneous Poisson process until a trade occurs. The proposed methods under this assumption are PIa and PIb, which are based on prediction intervals of quotes in a given second. PIb accounts for estimation uncertainty. Two methods based on the index of dispersion are also presented, namely VTa and VTb. To account for the possibility that the underlying quote process is not a homogeneous Poisson process, we present two methods which account for autocorrelation in the quote revision process; PIAR1a and PIAR1b. PIAR1b accounts for estimation uncertainty through a parametric bootstrap procedure. The performance of these methods is first compared with a method by Vergote (2005) in a simulation study. This study shows that as the jump in the quote intensity decrease the performance of VTa, VTb and Vergote's method deteriorate. The proposed methods PIa, PIb, PIAR1a and PIAR1b performs reasonably well detecting the breakpoint. The methods based on the index of dispersion fail due to the fact that the first and second order sample moments are changing close to proportionally to each other. Vergote's method performs poorly as λ_2 decrease due to a tuning parameter designed to react to more significant increases in intensity. The simulation study also suggest that the shape of the quote intensity in λ_2 has little impact on the reported breakpoint.

Our empirical study of stocks at the New York stock exchange shows that more reasonable results are obtained when applying the PIAR1b method, which account for both autocorrelation and estimation uncertainty. Similar results are also obtained from the method proposed by Vergote. As the quote intensity contemporaneously with and one second prior to a trade increase significantly the tuning parameter applied in Vergote (2005) effectively detects these breakpoints, but it is not as sensitive to less identifiable breaks in the quote intensity. This is an issue for stocks that report a steady increase in quote intensity up to the time of the reported trade. PIAR1b is better at capturing this effect. Our proposed methods based on a homogeneous Poisson process produce narrower prediction intervals and hence report earlier breakpoints than the two methods that account for autocorrelation. The two methods based on the index of dispersion also have a drawback in that they report at breakpoint when there is a drop in the quote intensity. This makes them less applicable to the data considered in this empirical study.

Contrary to what one might expect the intensity in quote revisions increase one second prior to the reported trades for most stocks in our sample. This result may occur due to a rounding in the TAQ dataset. While our proposed methods based on the assumption of a homogeneous Poisson process tend to report structural breakpoints prior to this time point, accounting for autocorrelation and estimation uncertainty seems to provide more reasonable empirical results.

References

Ball, C., and Chordia, T. (2001), "True Spreads and Equilibrium Prices," *Journal of Finance*, 56(5), 1801–1835.

Busse, J., and Green, T. (2002), "Market Efficiency in Real Time," *Journal of Financial Economics*, 65(3), 415–437.

Chan, K., Chung, Y., and Fong, W. (2002), "The Informational Role of Stock and Option Volume," *Review of Financial Studies*, 14(4), 1049–1075.

Chordia, T., Roll, R., and Subrahmanyam, A. (2001), "Market Liquidity and Trading Activity," *Journal of Finance*, 56(2), 501–530.

Chordia, T., Roll, R., and Subrahmanyam, A. (2002), "Order Imbalance, Liquidity and Market Returns," *Journal Of Financial Economics*, 65(1), 111–130.

Easley, D., O'Hara, M., and Saar, G. (2001), "How Stock Splits Affect Trading: A Microstructure Approach," *Journal of Financial and Quantitative Analysis*, 36(1), 25–51.

Edelen, R., and Gervais, S. (2003), "The Role of Trading Halts in Monitoring a Specialist Market," *Review of Financial Studies*, 16(1), 263–300.

Engle, R., and Patton, A. (2004), "Impacts od Trades in an Error-Correction Model of Quote Prices," *Journal of Financial Markets*, 7(1), 1–25.

Hasbrouck, J., Sofianos, G., and Sosebee, D. (1993), "New York Stock Exchange Systems and Trading Procedures," , . NYSE Working Paper 93-01.

Henger, T., and Wang, J. (2006), "On the importance of timing specifications in market microstructure research," *Journal of Financial Markets*, Forthcoming.

Huang, R., and Stoll, H. (1997), "The Components of the Bid-ask Spread: A General Approach," *Review of Financial Studies*, 10, 955–1034.

Huang, R., and Stoll, H. (2001), "Exchange Rates and Firm's Liquidity: Evicence from ADRs," *Journal of International Money and Finance*, 20(3), 297–325.

Karlis, D., and Xekalaki, E. (2000), "A Simulation Comparison of Several Procedures for Testing the Poisson Assumption," *The Statistician*, 43(3), 355–382.

Kryzanowski, L., and Zhang, H. (2002), "Intraday Market Price Integration for Shares Cross-Listed Internationally," *Journal of Financial Quantitative Analysis*, 37(2), 243–369.

Lee, C., and Ready, M. (1991), "Inferring Trade Direction From Intraday Data," *Journal Of Finance*, 46, 733–747.

Moberg, J.-M. (2007), "Trade-quote matching algorithm for NYSE stocks," *Working Paper*, .

Nyholm, K. (2003), "Inferring the Private Information Content of Trades: A Regime Switching Approach," *Journal of Applied Econometrics*, 18(4), 457–470.

Pascual, L., Romo, J., and Ruiz, E. (2001), "Effects of parameter estimation on prediction densities: a bootstrap approach," *International Journal Of Forecasting*, 17, 83–103.

Piwowar, M., and Wei, L. (2006), "The Sensitivity of Effective Spread Estimates to Trade-Quote Matching Algorithms," *Electronic Markets*, 16(2), 112–129.

R Development Core Team (2005), *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0.
URL: *http://www.R-project.org*

Schultz, P. (2000), "Stock Splits, Tick Size and Sponsorships," *Journal of Finance*, 55(1), 429–450.

Stoll, H. (2000), "Friction," *Journal of Finance*, 55(4), 1479–1514.

Venkataraman, K. (2001), "Automated versus Floor Trading: An Analysis of Excecution Cost on the Paris and New York Exchange," *Journal of Finance*, 56(4), 1445–1485.

Vergote, O. (2005), "How to Match Trades and Quotes for NYSE Stocks," , . Working paper, Katholeike Universiteit Leuven.

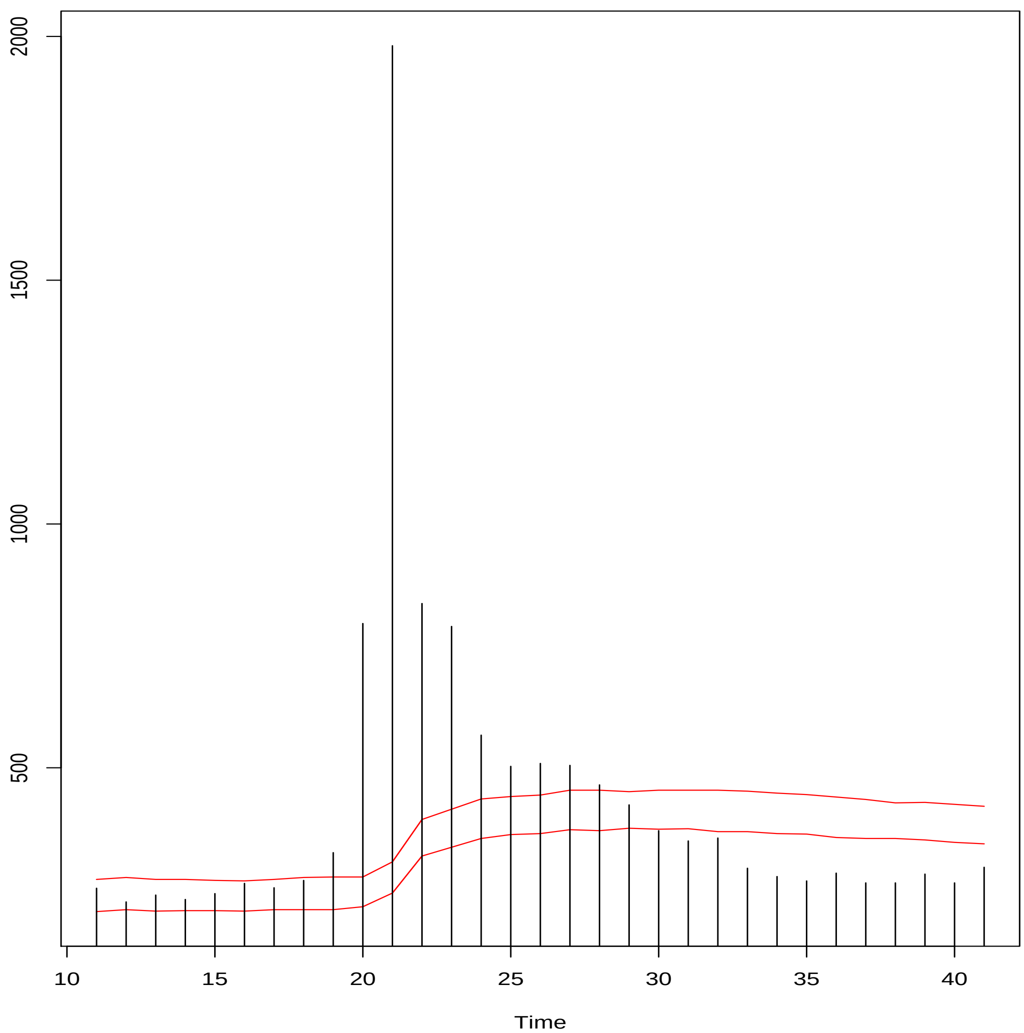

Figure 1: Frequency of quotes around registred trades with corresponding 95%-prediction intervals for the ABM-stock.

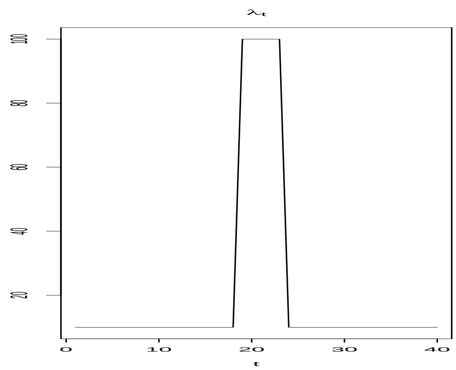

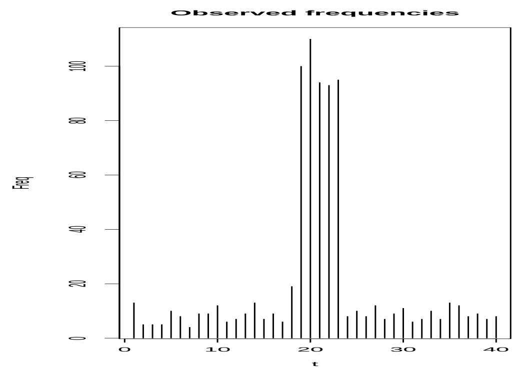

Figure 2: λ-function and a typical frequency histogram of observed events (trades) for DGP 1.

λ_2	PIa	PIb	PIARa	VTa	VTb	VER
30	19	19	19	19	19	19
20	19	19	18	21	20	19
15	19	19	18	40	40	22
12	20	20	19	40	40	40

Table 1: Monte Carlo medians of estimated breakpoint for DGP 1 with different values of λ_2. The true value is 19

λ_2	PIa	PIb	PIARa	VTa	VTb	VER
30	19	19	19	20	20	19
20	19	19	19	23	21	20
15	20	20	19	40	40	26
12	20	20	19	40	40	40

Table 2: Monte Carlo mean of estimated breakpoint for DGP 2 with different values of λ_2. The true value is 19

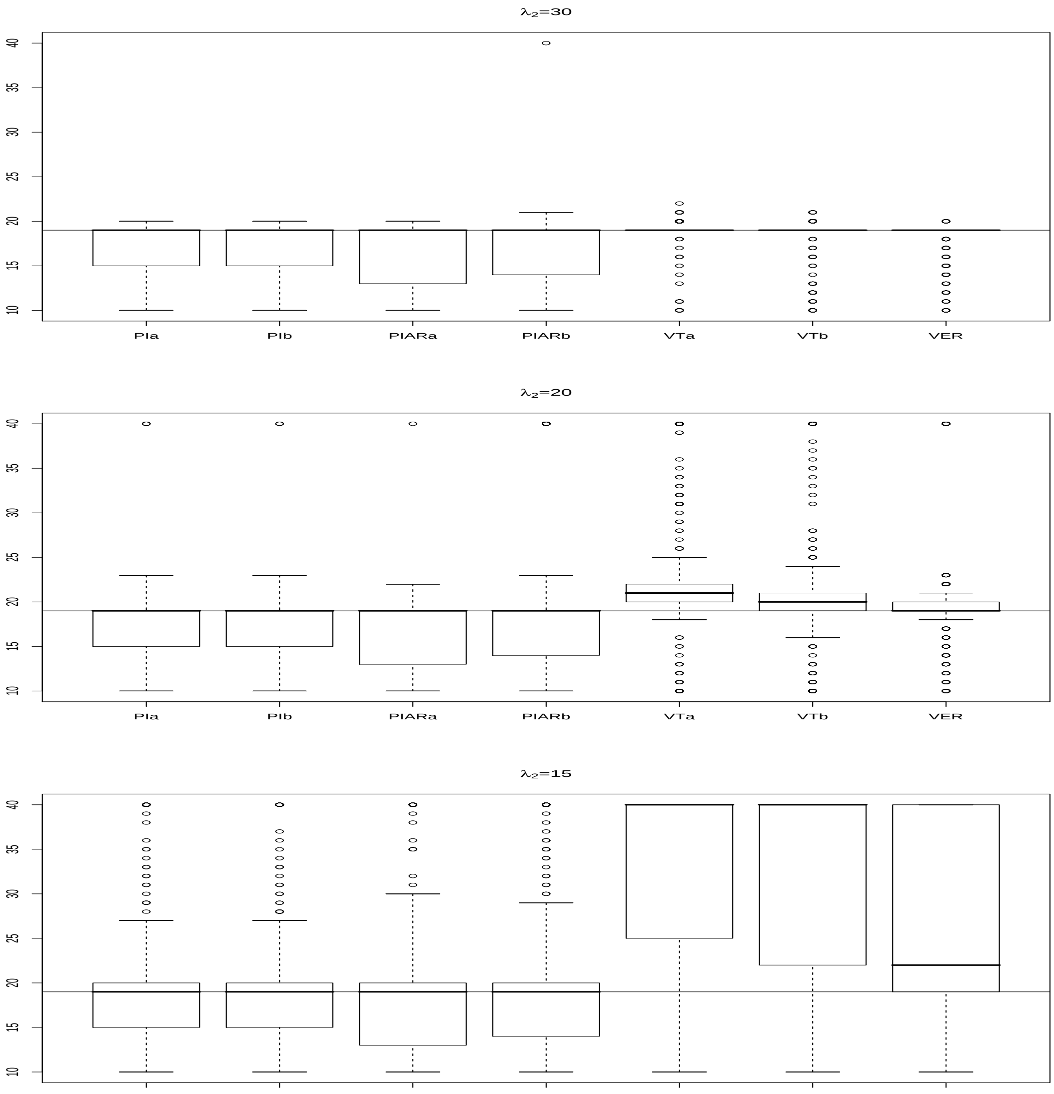

Figure 3: Boxplots of the simulation results under DGP 1. The horizontal line corresponds to the true breaking point 19.

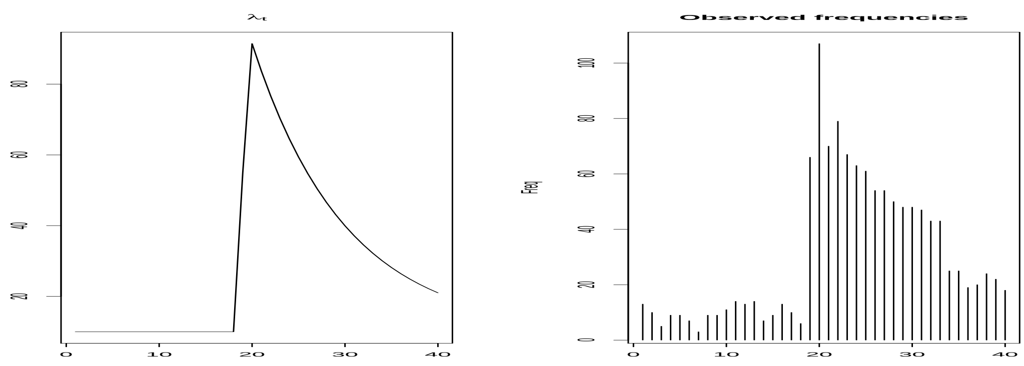

Figure 4: λ-function and a typical frequency histogram of observed events (trades) for DGP 2.

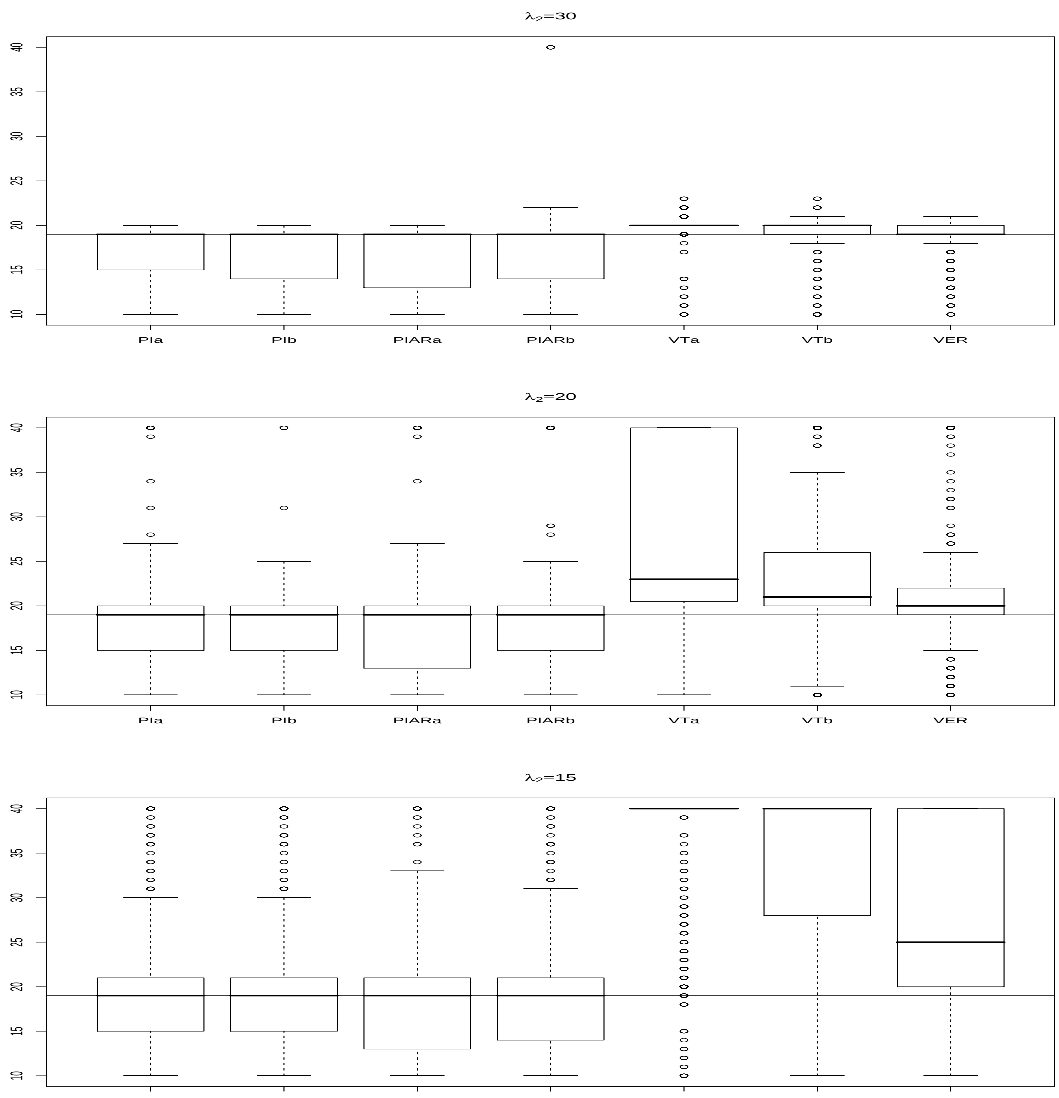

Figure 5: Boxplots of the simulation results under DGP 2. The horizontal line corresponds to the true breaking point 19.

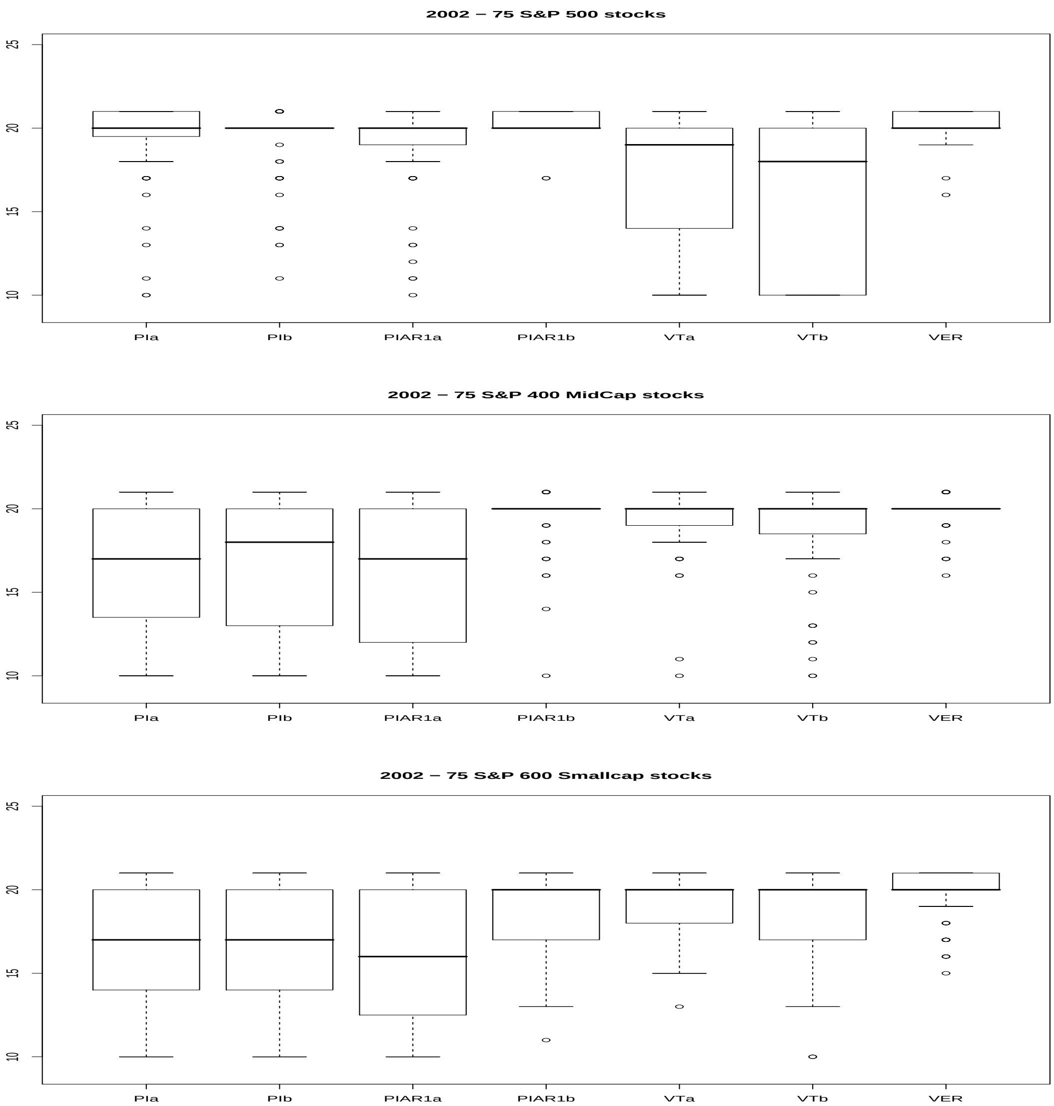

Figure 6: Boxplots of the empirical results during April 2002. The sample in each group is based on 75 stocks.

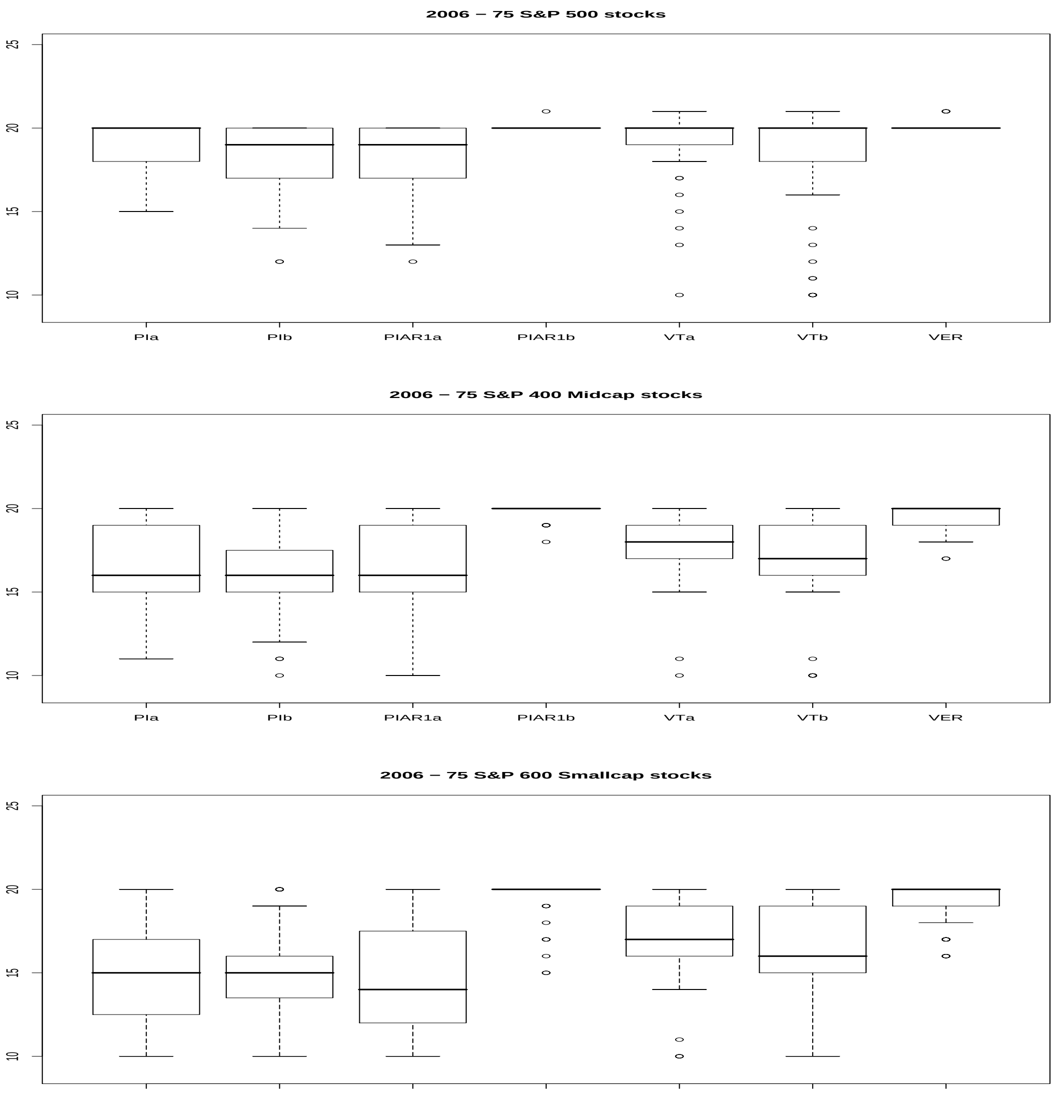

Figure 7: Boxplots of the empirical results during April 2006. The sample in each group is based on 75 stocks.

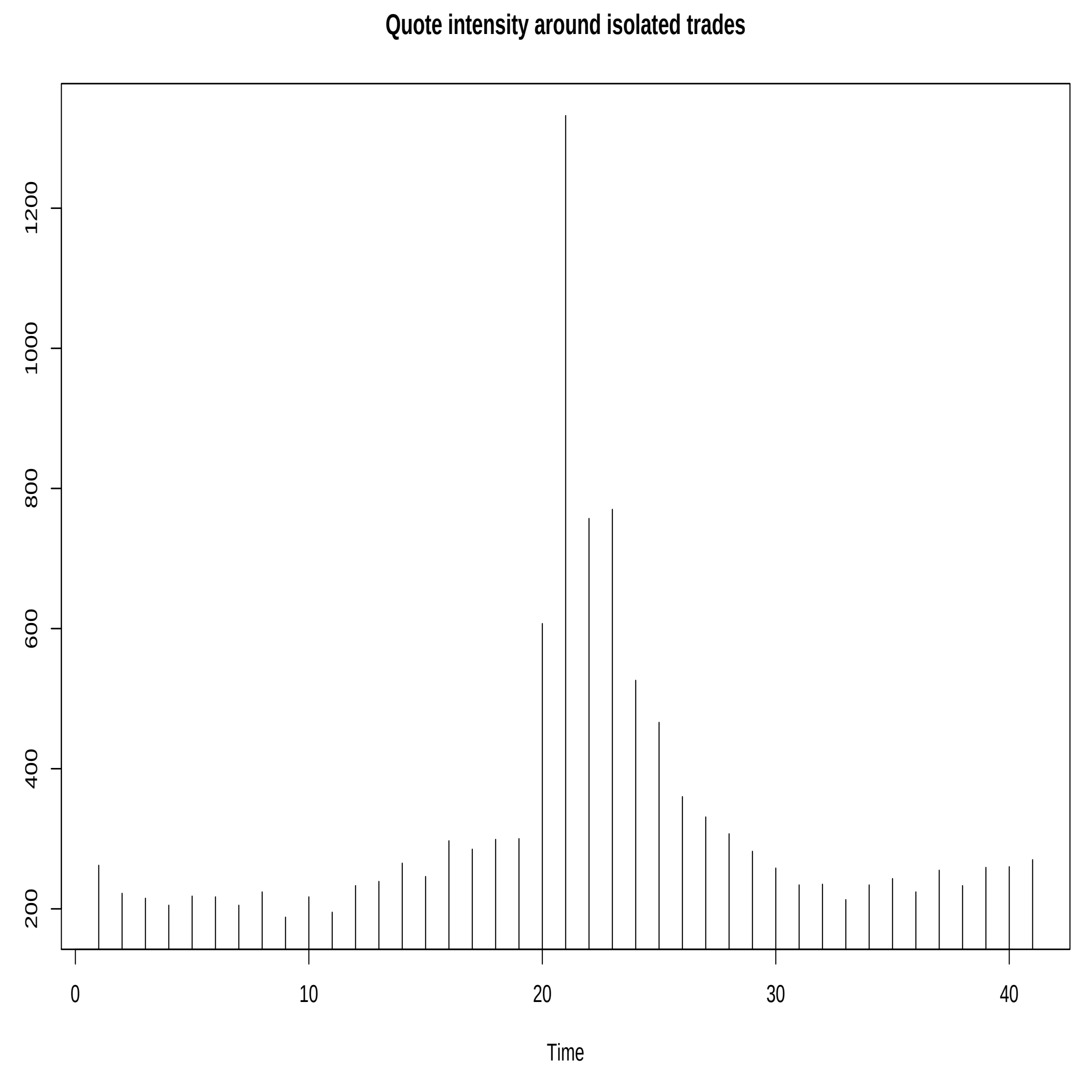

Figure 8: Frequency of quotes around isolated trades with corresponding 95%-prediction intervals for the AOS-stock during April 2006. A trade is defined as isolated if there are no other trades in a 40 second window centered around the trade under consideration. The quote intensity are plotted for each time point t=[11,...,41]. Trades are reported at $t = 21$. The y-axis present the number of quotes reported at each time point.

Table 3: Descriptive statistics

	Market Capitalization		
	Large	Medium	Small
2002:			
Trades	1257	401	178
Isolated trades	108	105	70
Quotes	2513	1227	781
2006:			
Trades	3165	1518	789
Isolated trades	55	95	104
Quotes	11622	6277	3951

This table shows average numbers per trading day during April 2002 and 2006.

Table 4: Distribution of quotes around isolated trades: 2002

Large Market Capitalization

t_{rel}	Pia	Pib	PIAR1a	PIAR1b	Ver	VT1	VT2
< 15	7	8	9	0	0	28	36
15	0	0	0	0	0	1	1
16	1	1	0	0	1	0	3
17	7	9	9	3	1	4	4
18	8	4	5	0	0	9	9
19	3	1	3	0	3	13	16
20	47	57	57	69	57	41	29
21	28	19	16	28	37	3	1
22	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0
> 23	0	0	0	0	0	0	0

Medium Market Capitalization

t_{rel}	Pia	Pib	PIAR1a	PIAR1b	Ver	VT1	VT2
< 15	28	28	36	5	0	1	8
15	9	12	11	7	1	4	5
16	7	9	7	7	4	5	4
17	11	12	11	17	8	8	15
18	5	3	3	5	4	12	11
19	7	4	4	5	4	5	5
20	27	29	27	45	52	44	43
21	7	3	3	8	27	20	9
22	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0
> 23	0	0	0	0	0	0	0

Small Market Capitalization

t_{rel}	Pia	Pib	PIAR1a	PIAR1b	Ver	VT1	VT2
< 15	28	28	36	5	0	1	8
15	9	12	11	7	1	4	5
16	7	9	7	7	4	5	4
17	11	12	11	17	8	8	15
18	5	3	3	5	4	12	11
19	7	4	4	5	4	5	5
20	27	29	27	45	52	44	43
21	7	3	3	8	27	20	9
22	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0
> 23	0	0	0	0	0	0	0

This table shows the distribution of quotes around isolated trades for 225 stocks during April 2002. The results for these stocks are grouped together according to market capitalization. The three tables show the results for 75 stocks in the S&P500 Index, the S&P MidCap 400 Index and the S&P SmallCap 600 Index respectively. The first column (t_{rel}) indicates the time of the quote revisions relative to the isolated trades. $t_{rel} = 20$ indicates the percentage of quotes that were registered 1 second before the trade.

Table 5: Distribution of quotes around isolated trades: 2006

Large Market Capitalization

t_{rel}	Pia	Pib	PIAR1a	PIAR1b	Ver	VT1	VT2
< 15	0	4	3	0	0	4	13
15	4	5	5	0	0	1	0
16	8	11	12	0	0	1	1
17	5	9	8	0	0	3	5
18	16	13	12	0	0	4	12
19	12	12	11	0	0	20	12
20	55	45	49	99	96	63	53
21	0	0	0	1	4	4	3
22	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0
> 23	0	0	0	0	0	0	0

Medium Market Capitalization

t_{rel}	Pia	Pib	PIAR1a	PIAR1b	Ver	VT1	VT2
< 15	13	15	16	0	0	3	9
15	16	21	17	0	0	1	5
16	21	21	21	0	0	15	16
17	15	17	13	0	3	23	24
18	5	5	4	1	8	12	8
19	15	8	12	7	17	24	21
20	15	12	16	92	72	23	16
21	0	0	0	0	0	0	0
22	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0
> 23	0	0	0	0	0	0	0

Small Market Capitalization

t_{rel}	Pia	Pib	PIAR1a	PIAR1b	Ver	VT1	VT2
< 15	44	39	51	0	0	7	15
15	19	15	15	3	0	13	20
16	8	23	5	1	5	24	17
17	7	8	4	4	9	11	13
18	4	1	3	1	5	12	7
19	7	7	7	4	15	12	12
20	12	8	16	87	65	21	16
21	0	0	0	0	0	0	0
22	0	0	0	0	0	0	0
23	0	0	0	0	0	0	0
> 23	0	0	0	0	0	0	0

This table shows the distribution of quotes around isolated trades for 225 stocks during April 2006. The results for these stocks are grouped together according to market capitalization. The three tables show the results for 75 stocks in the S&P500 Index, the S&P MidCap 400 Index and the S&P SmallCap 600 Index respectively. The first column (t_{rel}) indicates the time of the quote revisions relative to the isolated trades. $t_{rel} = 20$ indicates the percentage of quotes that were registered 1 second before the trade.