

The Information Content of Volume Price Impact for Intraday Liquidity Forecasting

Thibaut Moyaert¹

January 14, 2013

Working Paper

Abstract

This study aims at forecasting intraday liquidity pools in the order book. We compute an adjusted Amihud illiquidity ratio to detect pools of liquidity. Our methodology is built around a fixed volume event chart and presents a threefold advantage. First, we solve the non linearity issue of the widely used Amihud ratio. Second, we consider liquidity as a trade size specific concept. Third, we turn liquidity into a volatility proxy, which allows to transpose the extensive volatility forecasting literature into the liquidity framework. We outline some predictable patterns that statistically outperform a random walk model. We also find that trade-based and order book-based measures are both informative of forthcoming liquidity. These results are robust across several time frames and suggest potential execution improvement for market participants.

JEL Classification: G17, C53, C58

Keywords: Intraday Liquidity, Market Microstructure, Forecasting, GARCH models

¹ Thibaut Moyaert, Louvain School of Management, Université catholique de Louvain, 151 Chaussée de Binche - 7000 Mons (Belgium), e-mail: thibaut.moyaert@uclouvain.be. We are grateful to NYSE Euronext in Paris for providing the data. We also thank Catherine D'Hondt, Rudy De Winne, Sébastien Laurent, Albert J. Menkveld, Marco Pagano and Mikael Petitjean for providing useful comments and suggestions, as well as participants at the 2012 SFM International Conference in Kaohsiung (Taiwan) and at the 2012 AFM International Conference in Auckland (New Zealand). Any remaining errors are the responsibility of the author. The author gratefully acknowledges the support from the ARC grant 09/14-025 and from the CISM.

1 Introduction

Liquidity is a key concept for market efficiency and financial system stability. The expansion of high frequency trading, the 2010 flash crash and the implementation of new regulations, such as Reg NMS in the United States and MiFID in Europe, made liquidity issues a hot topic drawing the attention of both practitioners and academics.

In the theory on market microstructure, an exhaustive definition of market liquidity typically refers to the four dimensions identified in Harris (2003), i.e. tightness, depth, resiliency and immediacy. The literature brings out two broad categories of liquidity proxies, namely trade-based and order book-based measures. As they show little correlation, a bunch of research focuses on determining which type of measure is the best liquidity proxy. Goyenko et al. (2009), among others, show that the choice of liquidity proxy may significantly impact our results.

Given its multidimensional feature, liquidity cannot be captured by a single figure. Hence, practitioners must strike a balance between the different liquidity dimensions to select the best proxy according to their needs. As outlined by Irvine et al. (2000), investors can be patient or impatient to trade. On the one hand, impatient investors, such as arbitrageurs, require an immediate execution of their orders. On the other hand, patient investors, such as institutional investors who rebalance or hedge their portfolio, look for occasional liquidity pools to execute their orders inexpensively.

In this paper, we investigate liquidity from the viewpoint of patient investors. In their view, a liquid market is a market where the uncertainty around the execution price is low. In this framework, liquidity is set through volume price impact as measured by the Amihud ratio.¹ This ex-post liquidity measure neglects the immediacy dimension but allows to capture both potential uncommitted liquidity² and market resiliency³ to assess the level of execution price uncertainty.

Our methodology is based on a fixed volume event chart and has a threefold advantage. First, Kempf and Korn (1997) outline a major pitfall in the widely used Amihud illiquidity

¹The Amihud ratio is defined as follows: $IL_t = \frac{|R_t|}{V_t}$, where $|R_t|$ is the absolute return and V_t is the trading volume at time t .

²Uncommitted liquidity is hidden liquidity that is not displayed in the order book.

³Resiliency is the speed with which prices recover from a random, uninformative shock. Large (2007) and Degryse et al. (2005) outline the particular importance of resiliency in the case of an electronic limit order book, where the price-smoothing agency of a specialist or market maker is absent.

ratio as it implies a linear relationship between volume and volatility, which is typically not the case. Setting a fixed trading volume allows us to bypass the use of a ratio. Second, Irvine et al. (2000) address the idea that a market may be at the same time liquid for small orders and illiquid for large orders. Our event chart allows us to assess liquidity as a trade size specific concept. Furthermore, the fixed trading volume of the event chart can be adjusted to the order size considered by market participants. Third, we turn liquidity into a volatility proxy, which allows us to transpose the extensive literature on volatility forecasting into the liquidity framework.

Using time series models on a sample of Euronext stocks, we document some predictable patterns in volume price impact that statistically outperform a random walk model in- and out-of-sample. Consistent with Cont et al. (2011), we show that trade-based and order book-based liquidity measures are informative of forthcoming price impact. We also find support for Jain et al. (2011) who show that the order book bid side is more informative than the ask side. Among the four tested models, the ARX(3)-GARCH(1,1) model stands out from the crowd while the ARMAX model underperforms the other models. Robustness checks confirm the consistency of our results across stocks and different time frames.

All our results suggest practitioners could benefit from time series models to build execution price improvement strategies. On average, the end of the continuous trading day seems to be the best time for institutionals to hedge or rebalance a portfolio. It is indeed characterized by a slightly -but stock specific- negative seasonality, a higher predictability of volume price impact and a higher trading volume.

The paper is structured as follows. Section 2 provides an overview of the existing literature related to liquidity proxies and liquidity forecasting. Section 3 details the data used for our empirical work. Section 4 presents the applied methodology. Section 5 reports the in-sample and out-of-sample results as well as several robustness checks. Finally, Section 6 sums up the findings and identifies possible extensions.

2 Literature Review

The plethora of liquidity proxies in the literature illustrates that liquidity remains an elusive concept. Market liquidity is an unobserved process with multidimensional characteristics, which makes it impossible to measure by a single figure. Practitioners and academics have to

strike a balance between the different liquidity dimensions to select the best proxy for their purpose.⁴

In this section, we only present the benchmark of ex-post (trade-based) and ex-ante (order book-based) measures.

The benchmark of ex-post liquidity measures is the Amihud illiquidity ratio, presented in Amihud (2002). This category of proxies is mainly used when market timing is neglected, such as for brokers or hedgers who rather focus on improving execution price. The Amihud ratio measures liquidity through the relationship between traded volume and absolute return or, in other words, the volume price impact:

$$Amihud_t = \frac{|R_t|}{V_t}$$

where $|R_t|$ is the absolute return and V_t is the trading volume at time t .

The shortcomings of this measure, and more broadly of all ex-post measures, is that they cannot be computed in the absence of trading volume. For illiquid stocks, this measure is thus less relevant as it may be lagging because of the absence of trading volume. Furthermore, trade-based measures do not consider liquidity as a trade size specific concept. Irvine et al. (2000) show that a market may be at a point in time liquid for small orders but illiquid for large orders. Eventually, Kempf and Korn (1997) point out the fact that the use of a ratio implies that the relationship between trading volume and absolute return is linear which is typically not the case. This non-linear relationship entails misleading interpretation and conclusions.

As the availability of high frequency data has increased, attention has shifted towards the informational content of the order book. The benchmark of order book-based (ex-ante) measures is the Cost of Round Trip trade, hereafter CRT. It represents the cost of buying and selling the same quantity of a stock immediately. In other words, it measures the liquidity cost for investors that require immediate execution:

$$CRT_{q,t} = \frac{B_{q,t} - S_{q,t}}{q * MQ_t}$$

⁴For the sake of simplification, two possible biases of liquidity proxies are often neglected in the literature. On the one hand, liquidity is asymmetric as confirmed by Kalay et al. (2004). Buy orders lead to significant higher price impact than sell orders and the price impact is higher for both in market decline. This cannot be accounted for through a unique ratio. On the other hand, Hall and Hautsch (2007) show that financial markets are driven by both information and liquidity. Liquidity proxies do not filter price impact as either liquidity driven or information driven.

where $CRT_{q,t}$ is the cost of round trip for a trade size q at time t , $B_{q,t}$ is the cost of buying q shares, $S_{q,t}$ is the income from selling q shares at time t and MQ_t is the midquote at time t .

The CRT trade captures the maximum implicit transaction cost if immediacy is required. This implies that the measure only makes sense if a high proportion of liquidity is committed. Furthermore, as often only the best quotes of the order book are available, perfect elasticity of supply and demand assumption beyond the best quotes is often taken to compute the CRT for large orders.

Both Trade-based and order book-based measures coexist in the liquidity forecasting literature.

In a quote-driven market environment,⁵ Dufour and Engle (2000) extend the Hasbrouck (1991)'s vector autoregressive model for prices and trades to highlight some correlation between trading frequency and price impact. They find that high trading activity is related to both larger quote revisions and stronger positive autocorrelation of trades.

Later, Engle and Lange (2001) develop a new liquidity proxy, VNET, that captures the depth of the underlying stock corresponding to the net buy and sell volume during a given price movement. Using an event chart based on a given price movement, they also forecast the time to complete a fixed price movement relying on an ACD model variant. This forecasted time to achieve a given price movement could be used to set up market timing strategies.

In an order-driven market environment,⁶ Jain et al. (2011) document that the CRT explains considerably more of the upcoming transaction frequency, trade price and volatility than other measures.

Next, Tombeur and Wuyts (2011) highlight that order flow imbalance is a better predictor of momentum than order book imbalance in the short run. However, order book imbalance remains informative in the long run which is not the case for order flow imbalance.

Finally, Cont et al. (2011) outline order book-based and trade-based measures are both informative of future price impact. They show that highly correlated market orders lead to momentum while limit orders lead to range-bound market. They also find that order flow

⁵A quote driven market is a market that only displays bids and asks of designated market makers and specialists for a specific security.

⁶An order-driven market is a market where all buyers and sellers display the prices at which they wish to buy or sell a particular security, as well as the amounts of the security desired to be bought or sold.

imbalance and momentum exhibit a linear relationship with a slope inversely proportional to depth.

As our liquidity measure shares some stylized features with volatility, we also rely upon the literature on volatility forecasting. Given that this literature is extensive, it is not reviewed entirely here as it is beyond this paper. In a nutshell, there are two strands in this literature. The first stream focuses on assessing models to find some consistent outperformance between various GARCH-type models and autoregressive models. Poon and Granger (2003) and Lunde and Hansen (2005) provide a quick and effective snapshot on this stream. The second stream include researches that aim at defining less noisy volatility estimates. Andersen and Bollerslev (1998) advocate that volatility models do provide accurate volatility forecasts. However, poor performance of those models is due to noisy estimates of volatility. De Vilder and Visser (2008) rank several volatility proxies and finds support for Andersen and Bollerslev (1998) on the noisiness of squared returns to estimate volatility.

3 Data


3.1 Sample

The data pertain to 81 stocks that belong to three national major European NYSE Euronext stock indexes: CAC40, AEX and BEL20. The sample period spans 61 trading days (from February 1, 2006 to April 30, 2006). Building on trades and orders, the database has been rebuilt using a replication of the Euronext algorithm. This allows us to use trade and order book data to compute an extensive set of liquidity measures.

We arrange our database in event-based intervals. Hasbrouck (1996) outlines the advantage of event-based charts over mainstream time-based charts for forecasting purposes as they tend to be less noisy. In this paper, an interval event is defined as a given trading volume, which allows us to consider liquidity as a trade size specific concept. To ensure consistency across stocks, we set a fixed trading volume for each stock that corresponds to the average 15-minute trading volume observed for the stock over the whole sample period.

Our final data consist in 154,633 intervals split across 81 stocks. For each stock, the observations are split into two equal databases, the first half is for in-sample model fitting and the second half is for out-of-sample model forecasting performance.

Figure 1: Volume Seasonality



Distributional properties of trading volume. The chart represents the average trading volume across the 81 stocks. The continuous trading day is discretized in 34 intervals of 15 minutes.

We apply some filters to the data in order to ensure the relevance of our results. We remove the first interval of each continuous trading day to avoid referring to previous day order book data. We also cancel the last observation of the continuous trading day because the event, that is trading volume, on that interval may be uncompleted. We should mention as well that the trigger for a new interval is the order that crosses the fixed volume threshold. The traded volume per interval may thus slightly vary in order to avoid order splitting.

Figure 1 reports the distributional properties of the average trading volume across the 81 stocks over our sample period. As expected, the well-known U-shaped volume distribution is observed and our observations are thus outweighed at both the beginning and the end of the continuous trading day.

3.2 Liquidity Proxies

We aim at detecting liquidity pools in the market through volume price impact. For this purpose, we use an adjusted Amihud illiquidity ratio, hereafter AAR:

$$AAR_{i,t} = \frac{HL_{i,t}}{V_{i,t}}$$

where $HL_{i,t}$ is the High-Low and $V_{i,t}$ is the fixed traded volume for stock i during interval t .

Building on Alizadeh et al. (2002) who document the higher efficiency of High-Low compared to absolute returns and squared returns as a volatility proxy, we replace absolute returns by High-Low in the measure. In the framework of a volume event chart methodology, we set a fixed traded volume such that the AAR proxy fines down to the $HL_{i,t}$ measure. This allows us to consider liquidity as related to a given trade size.

Thanks to our rich database, we compute an extensive set of liquidity measures including Kang and Yeo (2008)'s dispersion and Næs and Skjeltorp (2006)'s slope, among others. For the sake of convenience, this section only reviews the explanatory variables that are used in the final model.

The final model includes an exogenous trade-based measure that is the number of trades. This is a state variable that is thus summed over the interval. In the framework of our event chart, the number of trades for a fixed trading volume represents the average trading size and is defined as follows:

$$\#Trades_{i,t} = \sum_{j=1}^N N_{i,j,t}$$

where N is the number of trades for stock i during interval t .

The remaining measures are order book-based measures. They represent a snapshot of the order book at a given point in time and are taken at the end of each interval. Those measures are respectively the depth at the five best bid and ask quotes ($Q5_{i,t}$), the relative spread ($RS_{i,t}$) and both bid and ask slopes ($BidSlope_{i,t}$ and $AskSlope_{i,t}$) which capture order book imbalance. They are defined as follows:

$$Q5_{i,t} = \sum_{j=1}^5 Q_{B,j,i,t} + Q_{A,j,i,t}$$

where $Q_{B,j,i,t}$ is the committed bid quantities and $Q_{A,j,i,t}$ is the committed ask quantities display at the j best limit for stock i at the end of interval t .

$$RS_{i,t} = \frac{PA1_{i,t} - PB1_{i,t}}{MQ_t}$$

where $P_{A1,i,t}$ is the best ask and $P_{B1,i,t}$ is the best bid for stock i at the end of interval t and MQ_t is the midquote at the end of interval t .

$$AskSlope_{i,t} = \frac{P_{A5,i,t} - P_{A1,i,t}}{P_{A1,i,t}}$$

where $P_{A5,i,t}$ is the limit price of the fifth best ask quote and $P_{A1,i,t}$ is the best ask for stock i at the end of interval t .

$$BidSlope_{i,t} = \frac{P_{B5,i,t} - P_{B1,i,t}}{P_{B1,i,t}}$$


where $P_{B5,i,t}$ is the limit price of the fifth best bid quote and $P_{B1,i,t}$ is the best bid for stock i at the end of interval t .

4 Methodology

4.1 Quantitative analysis

As a preliminary step, we point out some characteristics of the Adjusted Amihud Ratio (AAR). The main purpose is to detect any pattern in the AAR to support our model selection. As our liquidity proxy and volatility proxies share some similar stylized features, we pay specific attention to clustering (distribution skewness and kurtosis), heteroscedasticity and autocorrelation. The stationarity of the measure is tested using the Phillips-Perron test. The AAR displays strongly significant stationarity in our sample.

Figure 2: Autocorrelation Function of the adjusted Amihud ratio



The autocorrelation function (ACF) of the AAR is displayed for 20 lags on the left-hand side chart. The right-hand side chart plots the partial autocorrelation function (PACF) of the AAR up to 20 lags.

Figure 2 presents the AAR autocorrelogram, which outlines some autocorrelation in the proxy that nevertheless converges rather quickly to zero. We apply the extended sample autocorrelation function, presented in Tsay and Tiao (1984), to test the relevance of an ARMA dynamics. A long memory ARMA (1,1) process fits two out of three stocks of our sample.


Table 1: Descriptive Statistics of the adjusted Amihud ratio

Skewness	1.5781	Minimum	0
Excess Kurtosis	6.0783	Maximum	2.7231
Mean	0.2720	Normality test	<0.0001
Standard Deviation	0.1692	ARCH test 1-10	<0.0001
Median	0.2430	AR test 1-2	<0.0001

The table presents key descriptive statistics of the AAR across the 81 stocks for an average 15-minute volume observed for the stock over the whole sample period.

As reported in Table 1, the AAR displays highly positive excess kurtosis. We carry out ARCH disturbances test on the residuals that reveals a long-lasting and consistent ARCH effect across stocks. This is a strong support for GARCH (1,1) model. The normality test shows that our measure is not normally distributed. The positive skewness suggests that we should consider extensions of the GARCH model that account for asymmetry as well.

Figure 3: Adjusted Amihud ratio Seasonality



The chart plots the average AAR across the 81 stocks for a given trading volume, that is set to the average 15-minute volume observed for the stock over the whole sample period to ensure consistency across stocks. The continuous trading day is discretized in 34 intervals of 15 minutes.

The AAR exhibits seasonality as plotted in Figure 3. At the beginning of the continuous trading day, volume price impact is significantly higher. This is consistent with Cont et al. (2011) who show that there is lower depth in the morning because of information asymmetry. The very end of the continuous trading day displays a slightly lower price impact as well. We should point out that, while trading volume is lower at midday, there is no observed effect on price impact.

4.2 Model presentation

We use the general-to-specific model selection process, presented in Hendry and Krolzig (2001), to select the optimal set of explanatory variables. This procedure simplifies a general unrestricted model (GUM) to a congruent model. The selection encompasses test for model mis-specifications and diagnostic tests at each simplification step.

We consider three extensions of the model relying on the quantitative analysis of the adjusted Amihud ratio in Section 4.1, the models are presented in details hereafter:

- The ARX(3) model

$$\begin{aligned} AAR_{i,t} = & \beta_{i,0} + \sum_{j=1}^3 \gamma_{i,j} AR_{i,t-j} + \beta_{i,1} RS_{i,t-1} + \beta_{i,2} Q5_{i,t-1} \\ & + \beta_{i,3} BidSlope_{i,t-1} + \beta_{i,4} AskSlope_{i,t-1} + \beta_{i,5} \#Trades_{i,t-1} \\ & + \sum_{j=1}^5 \alpha_{i,j} SE_{i,t-j} + \alpha_{i,6} SE_{i,34} + \varepsilon_{i,t}, \end{aligned}$$

where $AAR_{i,t}$ is the High-Low for stock i at interval t , $SE_{i,t-j}$ is a dummy for a 15-minute time-based interval for stock i that captures the beginning and end of day seasonality in the dependent variable⁷ and $SE_{i,34}$ is the dummy for the last 15-minute time-based interval for stock i . Other explanatory variables were defined previously in section 3.2.

- The ARMAX (1,1) model

$$\begin{aligned} AAR_{i,t} = & \beta_{i,0} + \gamma_{i,1} AR_{i,t-j} + \delta_{i,1} \varepsilon_{i,t-1} + \beta_{i,1} RS_{i,t-1} + \beta_{i,2} Q5_{i,t-1} \\ & + \beta_{i,3} BidSlope_{i,t-1} + \beta_{i,4} AskSlope_{i,t-1} + \beta_{i,5} \#Trades_{i,t-1} \\ & + \sum_{j=1}^5 \alpha_{i,j} SE_{i,t-j} + \alpha_{i,6} SE_{i,34} + \varepsilon_{i,t}, \end{aligned}$$

where $AAR_{i,t}$ is the High-Low for stock i at interval t , $SE_{i,t-j}$ is a dummy for a 15-minute time-based interval for stock i that captures the beginning and end of day seasonality in the dependent variable and $SE_{i,34}$ is the dummy for the last 15-minute time-based interval for stock i . Other explanatory variables were defined previously in section 3.2.

⁷That is $SE_{i,1}$ is the seasonality dummy for the first 15 minutes of the trading day, $SE_{i,2}$ is the seasonality dummy for the next 15 minutes of the trading day, etc.

- The ARX (3)-GARCH (1,1) model

$$\begin{aligned}
AAR_{i,t} &= ARX(3)model, \\
\varepsilon_{i,t} &= z_{i,t}\sqrt{h_{i,t}}, \\
z_{i,t} &= N(0, 1), \\
h_{i,t} &= \omega_{0,i} + \lambda_{i,1}\varepsilon_{i,t-1}^2 + \theta_{i,1}h_{i,t-1},
\end{aligned}$$

where $h_{i,t}$ is the variance process for stock i at interval t .

- The ARX (3)-EGARCH (1,1) model

$$\begin{aligned}
AAR_{i,t} &= ARX(3)model, \\
\varepsilon_{i,t} &= z_{i,t}\sqrt{h_{i,t}}, \\
z_{i,t} &= N(0, 1), \\
\log(h_{i,t}) &= \omega_{0,i} + \lambda_{i,1}\log(h_{i,t-1}) + \theta_{i,1} \left(\frac{|\varepsilon_{i,t-1}|}{\sqrt{h_{i,t-1}}} - \sqrt{\frac{2}{\pi}} \right) + \zeta_{i,1} \left(\frac{\varepsilon_{i,t-1}}{\sqrt{h_{i,t-1}}} \right),
\end{aligned}$$

where $h_{i,t}$ is the variance process for stock i at interval t .

The model calibration is made through maximum likelihood function using Levenberg-Marquart optimization algorithm. The initial values are set to the parameter values of the pooling regression over the whole sample period. Autoregressive models do not impose a strictly positive value for the adjusted Amihud ratio. In few exceptional cases where the models return negative value, we set the value of the dependent variable to zero.

As mentioned earlier, we divide the sample period into an estimation period (first half of data) that is used for fitting the model in-sample and an evaluation period (second half of data) used for out-of-sample forecasting. For the out-of-sample data, we perform a rolling regression. Each model is estimated over an initial window that corresponds to the estimation period and a 1-period ahead liquidity forecast is performed. The window is then rolled forward one period to obtain the second forecast and so on and so forth for all the out-of-sample data.

We use two mainstream loss functions to measure the discrepancies of the model compared to the realized value: the Mean Squared Error (MSE) and the Mean Absolute Error (MAE). As a remainder, the MSE loss function gives some extra penalty to models that display large errors as the observed discrepancy is squared. The MSE is defined as:

$$MSE_{i,t} = \frac{1}{N} \sum_{j=1}^N \left(\widehat{AAR}_{i,j,t} - AAR_{i,j,t} \right)^2$$

where $\widehat{AAR}_{i,t}$ is the estimated High-Low for stock i at interval t and $AAR_{i,t}$ is the realized High-Low for stock i at interval t . The MAE is defined as:

$$MAE_{i,t} = \frac{1}{N} \sum_{j=1}^N | \widehat{AAR}_{i,j,t} - AAR_{i,j,t} |$$

where $\widehat{AAR}_{i,t}$ is the estimated High-Low for stock i at interval t and $AAR_{i,t}$ is the realized High-Low for stock i at interval t .

We test the statistical outperformance of our models against a random walk model through the superior predictability ability test (SPA) presented in Hansen (2005). Hansen's SPA test is less prone to data mining issues than the Diebold-Mariano test. The SPA test addresses the null hypothesis that any alternative forecasting models are better than a benchmark model, here a random walk model.

5 Results

5.1 In-sample analysis

We report the fit summary statistics of each model in Table 2. For the mean modeling, the ARX model obviously outperforms the ARMAX model with a better information criteria⁸ and R^2 statistics. The modeling of the variance, through both a GARCH and an EGARCH process, significantly improves the information criteria. The standard deviations across stocks of the fit summary are lower for the GARCH-type models, which suggests a higher consistency across stocks than the other models.

Table 2: Fit Summary Statistics

	Log likelihood		AIC		R-square	
	mean	std	mean	std	mean	std
ARX	584.5609	224.7625	-995.40	446.73	0.257648	0.057231
ARMAX	580.2106	225.2216	-810.97	439.86	0.246212	0.057267
ARX-GARCH	604.0356	218.7943	-1103.65	435.25	0.254267	0.053659
ARX-EGARCH	591.5351	219.4758	-1108.50	436.59	0.252971	0.053093

The table reports the mean and standard deviation of key fit statistics across the 81 stocks for the four forecasting models. AIC is the acronym for Akaike Information Criteria.

As the ARX model better captures the mean process than the ARMAX model, we only report the ARX model parameter estimates in Table 3. The other models parameter estimates are provided in Table 9, 10 and 11 (in Appendix).

⁸The AIC is defined as follows: $AIC = -2\ln(ML) + 2k$, where k is the number of parameters in the statistical model and ML is the maximum likelihood function

Table 3: Parameter Estimates of the ARX model

Variable	Estimate	StdErr	tValue	**
AR1	0.191043	0.0348	-5.48	99%
AR2	0.090239	0.0334	-2.71	68%
AR3	0.069899	0.0329	-2.13	57%
Ask_slope	0.111076	0.0565	2.03	48%
Bid_slope	-0.133657	0.0558	-2.48	56%
N	0.000414	0.0002	1.73	41%
Q5	-0.000003	0.0000	-5.46	96%
RS	0.304829	0.0929	3.73	80%
SA_1	0.093568	0.0372	2.63	68%
SA_2	0.070035	0.0295	2.40	58%
SA_3	0.062984	0.0282	2.29	58%
SA_4	0.035948	0.0283	1.28	30%
SA_5	0.027848	0.0274	1.02	19%
SA_34	-0.04276	0.0357	-1.25	20%

This table reports the parameter estimates of the ARX(3) model. Exogenous variables are those presented in Section 3.2. SA_1 parameter is the dummy for seasonality of the first 15 minutes of the trading day; SA_2 the next 15 minutes of the trading day; SA_3 the third 15 minutes of the trading day, etc. SA_34 parameter is the last 15 minutes of the trading day seasonality. The column ** reports the percentage of stocks for which the parameter is significant at a 95% confidence interval.

The set of explanatory variables consists in both trade-based and order book-based measures. This finding is consistent with Cont et al. (2011) who show that both measures are informative of future volume price impact. Table 3 points out some parameters that are significant across almost all the stocks: the first autoregressive lag, depth and, to a lesser extent, relative spread. The bid slope displays a somewhat higher coefficient and is more often significant than the ask slope. This particular result brings support for Jain et al. (2011) who find that the bid side of the order book carries more information than the ask side. The average trading size is also informative of future liquidity for almost half of the 81 stocks (41%). The first hour of the continuous trading day displays a positive seasonality that is pronounced especially at the very beginning of the day. The end of day seasonality seems to be more stock

specific with only one out of five stocks that exhibits a significant seasonality during the last 15 minutes of the trading hours.

The GARCH parameters are reported in Table 10 available in Appendix. They all exhibit strong consistency across stocks. The EGARCH model parameters, in Table 11 available in the Appendix, are also significant across almost all the stocks, except for the asymmetry parameter Θ . Only one out of four stocks displays a significantly positive asymmetry in the dependent variable.

In a second step, we compare the performance of the forecasting models against a random walk model for both loss function in Table 4. All the models outperform the random walk with a 99% confidence level. Nevertheless, the choice of a loss function affects the ranking of the forecasting models. The performance of the GARCH-type models are very close to the ARX model using the MAE loss function. The performance of the ARX and ARMAX model improves when using the MSE loss function.

Table 4: Average Loss Function

MAE	MEAN	STD	SPA	MSE	MEAN	STD	SPA
RW	0.128715	0.132757		RW	0.034192	0.100329	
ARX	0.101001	0.094717	***	ARX	0.019172	0.055678	***
ARMAX	0.101365	0.095058	***	ARMAX	0.019310	0.056293	***
ARX-GARCH	0.100910	0.095689	***	ARX-GARCH	0.019339	0.057158	***
ARX-EGARCH	0.101034	0.095785	***	ARX-EGARCH	0.019382	0.057832	***

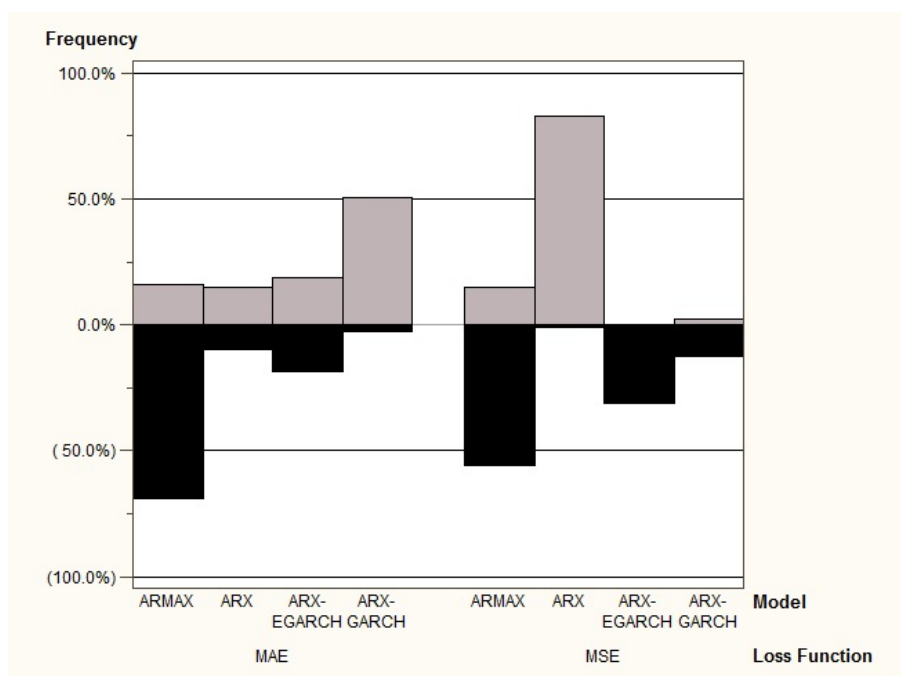
The table exhibits the mean loss function value and the standard deviation across the 81 stocks for the four forecasting models and the random walk model (benchmark). The column SPA reports the statistical outperformance of each forecasting model compared to the random walk model (RW). *** means the model outperforms the benchmark with a 99% confidence level.

The competition summary plotted in Figure 4 confirms that the performance of the GARCH-type models is strongly affected by the choice of the loss function. Based on the mean average error, the GARCH-type models seem to outperform the ARX model while their performances strongly deteriorate when the mean squared error is considered at the benefit of the ARX model. This result is in line with our expectations for two reasons. On the one hand, the

loss functions measure discrepancy in the mean process. Hence, we do not expect a strong improvement from the GARCH-type models. On the other hand, the MAE loss function is robust to the presence of outliers, which is not the case of the MSE loss function. The presence of outliers in the data explains the deterioration of the GARCH models performance when using the MSE loss function.

The ARMAX model underperforms the ARX model for the mean modeling across both loss functions but it still strongly outperforms the random walk model.

Figure 4: Summary of the forecast competition



Summary of the forecast competition between the four forecasting models for the 81 stocks. The figure plots the frequency of best (gray) and worst (black) in-sample loss for both loss functions.

5.2 Out-of-sample analysis

The out-of-sample rolling regression results, reported in Table 5, are consistent with the in-sample findings. The random walk model strongly underperforms all the forecasting models at a 99% confidence level. Overall, the ranking of the models is unchanged. The ARX model

outperforms the ARMAX model. The GARCH-type models are still affected by the choice of the loss function but their out-of-sample performances are better on average than in-sample.


Table 5: Average Loss Function

MAE	MEAN	STD	SPA	MSE	MEAN	STD	SPA
RW	0.123411	0.127283		RW	0.031431	0.099755	
ARX	0.097024	0.088965	***	ARX	0.017328	0.043806	***
ARMAX	0.097467	0.089252	***	ARMAX	0.017466	0.044152	***
ARX-GARCH	0.096694	0.089595	***	ARX-GARCH	0.017377	0.045133	***
ARX-EGARCH	0.096773	0.089571	***	ARX-EGARCH	0.017388	0.045338	***

The table exhibits the mean loss function value and the standard deviation across the 81 stocks for the four forecasting models and the random walk model (benchmark). The column SPA reports the statistical outperformance of the four forecasting models compared to the random walk model (RW), *** means the model outperforms the benchmark with a 99% confidence level.

The same conclusions are drawn from Figure 5. The ARX model dominates the ARMAX model for the mean modeling across both loss functions. The forecast competition also confirms the better out-of-sample behavior of the GARCH-type models.


Figure 5: Summary of the forecast competition



Summary of the forecast competition between the four forecasting models for the 81 stocks. The figure plots the frequency of best (gray) and worst (black) out-of-sample loss for both loss functions.

To gauge the uncertainty surrounding our mean forecast, we plot the evolution of the GARCH variance process throughout the trading day in Figure 6. The average confidence interval around our liquidity forecast exhibits a downward trend. This suggests that the uncertainty around our forecast decreases along with the trading day. The variance process of the EGARCH model in Figure 7, available in Appendix, displays the same pattern.

Figure 6: Evolution of the confidence interval throughout the trading day for the GARCH model



The figure plots the average confidence interval of the forecasted mean across the 81 stocks throughout the trading day. The variance is modeled by a GARCH process. The continuous trading day is discretized in 34 intervals of 15 minutes.

Finally, we run SPA test between the four forecasting models to determine whether there is some statistical outperformance that exists among the forecasting models. The results are summarized in Table 6.

Table 6: Statistical underperformance frequency of the forecasting models

	ARX	ARMAX	ARX-EGARCH	ARX-GARCH
MAE	33.32%	35.80%	19.75%	8.64%
MSE	8.64%	18.52%	17.28%	7.41%

The table presents the underperformance frequency at a 99% confidence level for each forecasting model across the 81 stocks and for both loss functions.

The ARX-GARCH model stands out from the crowd. It statistically underperforms other forecasting models less than 10% of the cases for both loss functions. The other forecasting models display less consistent results. The statistical performance of the ARX and the ARMAX

model strongly depends on the loss function. Nevertheless, the ARMAX model remains the worst model. At last, the ARX-EGARCH model performance is quite consistent across loss functions but it underperforms other models for about 20% of the stocks. As a conclusion, the ARX(3)-GARCH(1,1) model seems to better capture the adjusted Amihud ratio dynamics.

All our results highlight some predictable patterns in the adjusted Amihud ratio against a random walk model. In the absence of more sophisticated strategies of execution price improvement, the end of the continuous trading day seems to be the best time for institutional to hedge or rebalance a portfolio. Indeed, it is characterized by a slightly -but stock specific- negative seasonality, a higher adjusted Amihud ratio predictability and a higher trading volume.

5.3 Robustness checks on other time frames

To test the robustness of our results, we consider different time frames that correspond to different trading sizes, specifically the 10-minute average trading volume and the 30-minute average trading volume that are reported in Table 7 and 8.

In both time frames, the results support our initial findings. The random walk model underperforms at a 99% confidence level all the forecasting models. The ranking of the model remains unchanged. As mentioned earlier, the loss function choice impacts the performance of the models. Nevertheless, the mean modeling is better captured by the ARX process than the ARMAX process across both loss functions.

The average loss function reduction is in the same scope as in our initial results. Both time frames display very similar results, which suggests our findings are robust across several trading sizes.

Table 7: Average Loss Function on 10-minute volume time frame

MAE	MEAN	STD	SPA	MSE	MEAN	STD	SPA
RW	0.104520	0.110008		RW	0.023026	0.073184	
ARX	0.083176	0.077016	***	ARX	0.012850	0.033804	***
ARMAX	0.083506	0.077407	***	ARMAX	0.012965	0.034174	***
ARX-GARCH	0.082817	0.077714	***	ARX-GARCH	0.012899	0.034913	***
ARX-EGARCH	0.082897	0.077722	***	ARX-EGARCH	0.012913	0.035030	***

The table exhibits the mean loss function value for a 10-minute volume time frame event chart and the standard deviation across the 81 stocks for the four forecasting models and the random walk model (benchmark). The column SPA reports the statistical outperformance of the four forecasting models compared to the random walk model (RW), *** means the model outperforms the benchmark with a 99% confidence level.

Table 8: Average Loss Function on 30-minute volume time frame

MAE	MEAN	STD	SPA	MSE	MEAN	STD	SPA
RW	0.169705	0.168880		RW	0.057319	0.148279	
ARX	0.129158	0.116900	***	ARX	0.030347	0.077660	***
ARMAX	0.129453	0.117385	***	ARMAX	0.030537	0.078827	***
ARX-GARCH	0.128805	0.117587	***	ARX-GARCH	0.030417	0.080388	***
ARX-EGARCH	0.128856	0.117891	***	ARX-EGARCH	0.030502	0.081049	***

The table exhibits the mean loss function value for a 30-minute volume time frame event chart and the standard deviation across the 81 stocks for the four forecasting models and the random walk model (benchmark). The column SPA reports the statistical outperformance of the four forecasting models compared to the random walk model (RW), *** means the model outperforms the benchmark with a 99% confidence level.

6 Conclusion

In this paper, we investigate the predictability of intraday liquidity pools. We detect liquidity pools through an adjusted Amihud ratio. In this framework, a liquid market is a market where the uncertainty around the execution price is low.

Our forecasting methodology relies on an event chart as advocated by Hasbrouck (1996). This approach has two main advantages. First, Kempf and Korn (1997) outline a major pitfall in the widely used Amihud illiquidity ratio. This ratio implies a linear relationship between the variables, that is volume and volatility, which is typically not the case. Setting a fixed trading volume allows us to bypass the use of a ratio. Second, Irvine et al. (2000) address the idea that a market may be at the same time liquid for small orders and illiquid for large orders. Our event chart allows us to assess liquidity as a trade size specific concept. Indeed, the fixed trading volume of the event chart can be adjusted to the order size considered by market participants.

We select the optimal forecasting model through the general-to-specific process presented in Hendry and Krolzig (2001). We use the Hansen (2005)'s SPA test on two mainstream loss functions to document the performance of our forecasting models against a benchmark, both in- and out-of-sample.

Our results outline some predictable patterns in the adjusted Amihud ratio that statistically outperform a random walk model. The ARX(3)-GARCH(1,1) model stands out from the crowd among the four tested forecasting models. Trade-based and order book-based liquidity measures are informative of forthcoming price impact as confirmed by Cont et al. (2011). We also find support for Jain et al. (2011), who show that order book bid side is more informative than the ask side.

We carry out robustness checks of our results on other time frames. The results are similar to our initial findings, which suggest our models are robust across several order sizes.

All our results suggest that practitioners could take advantage of such models to minimize the uncertainty around execution price. On average, the closing of the continuous trading day seems to be the best time for institutionals to hedge or rebalance a portfolio. Indeed, it is characterized by a slightly -but stock specific- negative seasonality, a higher adjusted Amihud ratio predictability and a higher trading volume. The economic significance and the

implementation of execution price improvement strategies are beyond the scope of this analysis and left for further research.

This paper may be extended in various directions. First, we can consider market timing issues within the same methodological framework. Second, we can transpose volatility literature extensions into the liquidity framework such as realized kernels and intraday jumps.

References

- Alizadeh, S., M. W. Brandt, and F. X. Diebold (2002). Range-based estimation of stochastic volatility models. *Journal of Finance* 57(3), 1047–1091.
- Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets* 5, 31–56.
- Andersen, T. and T. Bollerslev (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39(4), 885–905.
- Cont, R., A. Kukanov, and S. Stoikov (2011). The Price Impact of Order Book Events. *Working paper*.
- De Vilder, R. and M. Visser (2008). Ranking and combining volatility proxies for garch and stochastic volatility models. *MPRA paper*.
- Degryse, H., F. Jong, M. Ravenswaaij, and G. Wuyts (2005). Aggressive Orders and the Resiliency of a limit Order Market. *Review of Finance* 9(2), 201–242.
- Dufour, A. and R. Engle (2000). Time and the price impact of a trade. *Journal of Finance* 55(6), 2467–2498.
- Engle, R. and J. Lange (2001). Predicting VNET: A model of the dynamics of market depth. *Journal of Financial Markets* 4, 113–142.
- Goyenko, R., C. Holden, and C. Trzcinka (2009). Do liquidity measures measure liquidity. *Journal of Financial Economics* 92(2), 153–181.
- Hall, A. and N. Hautsch (2007). Modeling the buy and sell intensity in a limit order book market. *Journal of Financial Markets* 10, 249–286.
- Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business & Economic Statistics* 23(4), 365–380.
- Harris, L. (2003). Trading and Exchanges: Market Microstructure for Practitioners. *Financial Management Association survey and synthesis series*. Oxford University Press.
- Hasbrouck, J. (1991). Measuring the Information Content of Stock Trades. *Journal of Finance* 46, 179–207.

- Hasbrouck, J. (1996). Modeling Microstructure Time Series. *Handbook of Statistics 14*, 647–692.
- Hendry, D. and H.-M. Krolzig (2001). Computer automation of general-to-specific model selection procedures. *Journal of Economic Dynamics and Control 25*, 831–866.
- Irvine, P., G. Benston, and E. Kandel (2000). Liquidity beyond the inside spread: Measuring and using information in the limit order book. *Unpublished Working paper*.
- Jain, P., P. Jain, and T. McNish (2011). The predictive power of limit order book for future volatility, trade price, and speed of trading. *Working paper*.
- Kalay, A., O. Sade, and A. Wohl (2004). Measuring stock illiquidity: An investigation of the demand and supply schedules at the tase. *Journal of Financial Economics 74*(3), 461–486.
- Kang, W. and W. Yeo (2008). Liquidity beyond the best quote: A study of the nyse limit order book. *Working Paper Series*.
- Kempf, A. and O. Korn (1997). Market depth and order size. *Journal of Financial Market 2*(1), 29–48.
- Large, J. (2007). Measuring the resiliency of an electronic limit order book. *Journal of Financial Markets 10*(1), 1–25.
- Lunde, A. and P. Hansen (2005). A forecast comparison of volatility models: Does anything beat a garch(1,1)? *Journal of Applied Econometrics 20*(7), 873–889.
- Næs, R. and J. Skjeltorp (2006). Order book characteristics and the volume-volatility relation: Empirical evidence from a limit order market. *Journal of Financial Markets 9*(4), 408 – 432.
- Poon, S.-H. and C. Granger (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature XLI*, 478–539.
- Tombeur, G. and G. Wuyts (2011). Does the Information in the Limit Order Book Help to Predict Returns? *Working paper*.
- Tsay, R. S. and G. Tiao (1984). Consistent Estimates of Autoregressive Parameters and Extended Sample Autocorrelation Function for Stationary and Nonstationary ARMA Models. *Journal of the American Statistical Association 79*(385), 84–96.

7 Appendix

Table 9: Parameter Estimates of ARMAX model

Variable	Estimate	StdErr	tValue	**
AR1	0.364015	0.0545	7.23	99%
MA1	0.177301	0.0561	3.31	57%
Ask_slope	0.109189	0.0466	2.48	49%
Bid_slope	-0.12801	0.0445	-3.12	64%
N	0.000441	0.0001	2.03	47%
Q5	-0.000002	0.0000	-4.93	96%
RS	0.309721	0.0812	4.71	80%
SA_1	0.073911	0.0226	3.52	72%
SA_2	0.047836	0.0209	2.41	56%
SA_3	0.040007	0.0211	2.05	41%
SA_4	0.015703	0.0234	0.71	19%
SA_5	0.014723	0.0239	0.63	16%
SA_34	-0.018551	0.0413	-0.37	2%

This table reports the parameter estimates of the ARMAX model. Exogenous variables are those presented in Section 3.2. SA_1 parameter is the dummy for seasonality of the first 15 minutes of the trading day; SA_2 the next 15 minutes of the trading day; SA_3 the third 15 minutes of the trading day, etc. SA_34 parameter is the last 15 minutes of the trading day seasonality. The column ** reports the percentage of stocks for which the parameter is significant at a 95% confidence interval.

Table 10: Parameter Estimates of ARX-GARCH model

Variable	Estimate	StdErr	tValue	**
AR1	0.195058	0.0396	4.95	100%
AR2	0.096319	0.0370	2.63	68%
AR3	0.072862	0.0358	2.06	52%
Ask_slope	0.109938	0.0516	2.35	53%
Bid_slope	-0.130508	0.0495	-3.01	58%
N	0.000410	0.0002	1.85	46%
Q5	-0.000003	0.0000	-5.34	96%
RS	0.309035	0.0865	4.40	80%
SA_1	0.104549	0.0274	4.08	78%
SA_2	0.072160	0.0241	3.17	68%
SA_3	0.063488	0.0246	2.78	68%
SA_4	0.035369	0.0269	1.47	38%
SA_5	0.025258	0.0267	1.02	21%
SA_34	-0.033337	0.0367	-0.81	28%
ARCH0	0.004970	0.0022	4.21	82%
ARCH1	0.115539	0.0270	4.11	90%
GARCH1	0.626170	0.0771	20.40	92%


This table reports the parameter estimates of the ARX(3)-GARCH(1,1) model. Exogenous variables are those presented in Section 3.2. SA_1 parameter is the dummy for seasonality of the first 15 minutes of the trading day; SA_2 the next 15 minutes of the trading day; SA_3 the third 15 minutes of the trading day, etc. SA_34 parameter is the last 15 minutes of the trading day seasonality. The column ** reports the percentage of stocks for which the parameter is significant at a 95% confidence interval.

Table 11: Parameter Estimates of ARX-EGARCH model

Variable	Estimate	StdErr	tValue	**
AR1	-0.200212	0.0398	-5.07	98%
AR2	-0.098229	0.0365	-2.74	67%
AR3	-0.074189	0.0349	-2.14	53%
Ask_slope	0.100753	0.0512	2.18	49%
Bid_slope	-0.123747	0.0493	-2.87	58%
N	0.000399	0.0002	1.86	42%
Q5	-0.000003	0.0000	-5.15	95%
RS	0.305397	0.0866	4.36	81%
SA_1	0.095921	0.0278	3.70	73%
SA_2	0.069866	0.0243	3.06	65%
SA_3	0.059834	0.0244	2.65	63%
SA_4	0.033833	0.0263	1.43	37%
SA_5	0.023732	0.0259	0.97	22%
SA_34	-0.042692	0.0356	-1.22	33%
EARCH0	-1.041414	0.4146	-3.53	88%
EARCH1	0.168538	0.0419	4.01	75%
EGARCH1	0.749040	0.0989	30.89	93%
THETA	0.946708	1.1687	0.76	27%

This table reports the parameter estimates of the ARX(3)-EGARCH(1,1) model. Exogenous variables are those presented in Section 3.2. SA_1 parameter is the dummy for seasonality of the first 15 minutes of the trading day; SA_2 the next 15 minutes of the trading day; SA_3 the third 15 minutes of the trading day, etc. SA_34 parameter is the last 15 minutes of the trading day seasonality. The column ** reports the percentage of stocks for which the parameter is significant at a 95% confidence interval.

Figure 7: Evolution of the confidence interval throughout the trading day for the EGARCH model



The figure plots the average confidence interval of the forecasted mean across the 81 stocks throughout the trading day. The variance is modeled by an EGARCH process. The continuous trading day is discretized in 34 intervals of 15 minutes.