BUDAPESTI ÉRTÉKTŐZSDE
BUDAPEST STOCK EXCHANGE

The Budapest Liquidity Measure and its Application

Liquidity Risk in VaR measures

Budapest Stock Exchange Working Paper

January 2011

ÁKOS GYARMATI–MARTON MICHALETZKY–KATA VÁRADI

cee stock exchange group

We introduce the Budapest Liquidity Measure (BLM) and one of its possible applications in the field of risk management. BLM is a weighted spread measure, it represents the implicit costs of trading, which arise from the fact that actual trading is not executed at the mid-price.

Traditional VaR measures cover only the risk of the changing mid-price, they ignore the liquidity risk arising from the buying and selling of a position. With the use of BLM we show, how to integrate liquidity risk into the VaR-framework. While our method has already been introduced, it has never been tested on the Hungarian market. We also point out several areas of possible improvement. In our analysis we use the data of the stocks of the Budapest Stock Exchange, and find that even in the case of the most liquid stocks and smallest positions, the daily VaR measures can rise by up to 4% if we take the liquidity risk into account.[1]

1 Introduction of the BLM

In this section we give a short explanation of the BLM: we introduce the concept, the calculation and the interpretation. A more detailed description can be found in Kutas and Végh [2005].

BLM was created in 2005 by the Budapest Stock Exchange (BSE) using the model of the German XLM. The goal was to evaluate numerically one of the most important aspects of liquidity for the market participants, the implicit costs of transacting.

There are basically two groups of transaction cost:

- explicit costs: these are the direct cost of trading (e.g. broker fees, taxes)
- implicit costs: these are the indirect cost of trading (e.g. spreads)

BLM covers the implicit costs. The total implicit costs of a transaction consist of two parts: the bid-ask spread and the adverse price movement. The latter is the effect of the total transaction not being executed at the best level, but at worse levels. In this case the average price the market participant pays is worse than the best price.

BLM measures the implicit costs in percentage of the total transaction value. Consequently, it can only be defined for given order sizes. The standard order sizes used in the BSE are (in EUR thousands): 20, 40, 100, 200, 500.

In the following we take a closer look at the calculation of the BLM. Let a_i be the ith best ask price, b_i the ith best bid price and P_{mid} the mid-price. Then denote:

- $LP = \frac{a_1 - b_1}{2P_{mid}}$, the so-called liquidity premium, the half of the bid-ask spread,
- $b(n) = \frac{\sum b_i \cdot n_i}{n}$, where $\sum n_i = n$, the weighted average bid price at which the total of n shares can be sold,

[1] The authors work at the Department of Finance of Corvinus University of Budapest.

- $a(n)$, similarly the weighted average ask price,
- $APM_{bid}(q) = \frac{b_1 - b(n)}{P_{mid}}$, where $P_{mid} \cdot n = q$ the size of the position in EUR, the adverse price movement for the bid side,
- $APM_{ask}(q) = \frac{a(n) - a_1}{P_{mid}}$, similarly the adverse price movement for the ask side.

With the above notation BLM is calculated in the following way:

$$BLM(q) = \left(2 \cdot LP + APM_{bid}(q) + APM_{ask}(q)\right) \times 100,$$

BLM clearly always depends on the actual state of the order book, thus the calculation can only be done at a given time point. The system of BSE calculates BLM every second on trading days in the time interval of 9:02 am – 4:30 pm. The daily average BLM values are calculated as the time weighted averages of the intraday data.

BLM represents the implicit cost of turning around a position that is both selling and buying a position. E.g. BLM(500) = 60 bps means that the buying and selling of a position of EUR 500 thousand have an implicit cost of 500,000×60bps = EUR 3,000.

The calculation and interpretation of the BLM is illustrated in Figure 6 and Figure 7 of the Appendix.

As we see BLM covers two of the traditional dimensions of liquidity, tightness (the bid-ask spread, LP) and depth (adverse price movement). Therefore it gives a more precise description of the actual state of liquidity than the normal one-dimensional measures. The automated calculation enables the fast and easy collection of the BLM data, thus making its application easier.

A disadvantage of the BLM is that it does not capture another important dimension: time (immediacy). BLM is only a ”snapshot“ of the order book, it can only be defined for immediate transactions, thus it cannot capture the case when the total transaction is not executed immediately (order splitting). Also it cannot deal with the case when it is not possible to execute the total transaction immediately, due to the order book not being deep enough. This latter problem shows up in the system's calculation method, we will discuss this later in detail.

Bearing in mind both the advantages and disadvantages of the BLM we address the practical application in the next section.

2 A possible practical application of BLM

In this section we introduce one of the most promising applications of the BLM in the field of risk management.

One of the most common measures of risk is the Value at Risk (VaR) of an asset. This shows the maximum loss that the asset can suffer in a given period of time and with a given level of confidence. E.g. 10 day 99% VaR = HUF 10 million means that the probability of not losing more than HUF 10 million in the next 10 days is 99%.

The basic idea behind the practical application of BLM is that in VaR measures calculated for stocks only the risk of changing mid-price is taken into account. However with BLM it is possible to represent liquidity risk as well. Considering the liquidity issues related to the recent crisis, this is particularly important; the risk arising from illiquidity must not be ignored.

In the literature there are already several papers that try to capture liquidity risk, a part of these concentrate only on the bid-ask spread, others use different liquidity measures. Giot and Grammig [2005] and Stange and Kaserer [2008] use a similar liquidity measure to ours in the integration of liquidity risk. These studies are the starting point of our research.

In the following we shortly discuss the usual VaR-framework and the notation we use, then we show a possible integration of BLM.

- We calculate returns in continuous time: $r_t^{\Delta t} = \ln\left(\frac{P_{mid}^{t+\Delta t}}{P_{mid}^{t}}\right)$
- Traditional VaR for returns can be calculated using the following general formula:

$$VaR_{return}^{\alpha,\Delta t} = r_t^{\alpha,\Delta t} = \mu_{t+\Delta t} + \sigma_{t+\Delta t} q_{1-\alpha},$$

- where $\mu_{t+\Delta t}$ is the predicted expected return in $t+\Delta t$, $\sigma_{t+\Delta t}$ is the standard deviation of the prediction and $q_{1-\alpha}$ is the $1-\alpha$th quantile of a chosen distribution.
- To be able to compare the results with the above mentioned studies we calculate traditional VaR for the price as well:

$$VaR^{\alpha,\Delta t} = \frac{P_{mid}^{t} - P_{mid}^{t}\exp\left(r_t^{\alpha,\Delta t}\right)}{P_{mid}^{t}} = 1 - \exp\left(r_t^{\alpha,\Delta t}\right),$$

- which, considering that the connection between the prices is given by $P_{mid}^{t+\Delta t} = P_{mid}^{t}\exp\left(r_t^{\Delta t}\right)$, shows the maximum percentage loss in Δt period of time and with α confidence level. E.g. $VaR^{95\%,1day} = 5\%$ means that *due to the change of the mid-price* the probability of not losing more than 5% in the following day is 95%.

Now we will show how can to incorporate liquidity risk. The idea is to adjust returns using the BLM.

First we must discuss an issue considering BLM: originally BLM is an effective type measure, we calculate returns on continuous time however. To be able to match them we have to transform BLM to continuous time as well. This can be done the following way: let V be the gross value of a position, then only by buying and selling this position, its value is decreased because of the implied costs:

$$V_{net} = V \cdot \left(1 - BLM\left(V\right)\right) \text{ (effective form)},$$

on continuous time this can be written as:

$$V_{net} = V \cdot \exp\left(BLM_{cont}\left(V\right)\right) \text{ (continuous form)},$$

where BLM_{cont} is the notation used for the continuously calculated BLM. Both forms measure the same implicit costs, so we get:

$$BLM_{cont}\left(V\right) = \ln\left(1 - BLM\left(V\right)\right).$$

Now we focus on the integration of liquidity risk:

- Consider the following transformed returns, *net returns* from now on:

  $$r_{net,t}^{\Delta t}(q) = r_t^{\Delta t}(q) + \ln\left(1 - \frac{BLM_t(q)}{2}\right).$$

- Notice as a consequence of the definition of BLM, we have $r_{net,t}^{\Delta t}(q) < r_t^{\Delta t}(q)$.

Net returns are the normal returns decreased by the continuously calculated liquidity measure. We use the half of the BLM because - as was already mentioned - BLM contains the implicit cost of both selling and buying, however we only need one of them now. Naturally, by using half of the figures we implicitly assume that bid and ask sides are symmetric. We make a proposition to relax this assumption later in our paper.

- We can use the general formula to calculate VaR for net returns, then with these we calculate "total" price VaR, which now includes both mid-price and liquidity risk. We shall call this *Liquidity adjusted VaR* and use the following notation:

  $$L - VaR^{\alpha,\Delta t}(q) = 1 - \exp\left(r_{net,t}^{\alpha,\Delta t}(q)\right).$$

The interpretation of this measure is similar to the traditional VaR except that now not only the risk of changing mid-price, but also the liquidity risk is taken into account.

In the measure defined above liquidity risk is integrated, it does not show liquidity risk alone, but together with mid-price risk. However, using the traditional VaR measure it can easily be extracted. Let us consider the following measure:

$$\lambda(q) = \frac{L-VaR^{\alpha,\Delta t}(q) - VaR^{\alpha,\Delta t}}{VaR^{\alpha,\Delta t}}$$

We call $\lambda(q)$ relative liquidity impact or relative liquidity measure and it shows the maximum percentage loss due to illiquidity on a given time horizon and confidence level. In other words, it is the error made when liquidity risk is ignored.

Another approach of integrating liquidity risk into a VaR-framework is when VaR is calculated for the liquidity measure (like the traditional VaR) and then it is added to the traditional price VaR (see Bangia et al. [1999]). The problem with this is that this approach does not take into account that the correlation between mid-price and liquidity may not be perfect, thus the VaR values calculated for them may not be simply added (because the correlation may be imperfect, there may be some effect of diversification). Our approach however covers this issue as well.

3 Practical modeling

3. 1. About the data

The Budapest Stock Exchange provided us the data set from January 1, 2007. The data contain BLM, APM and LP time series for all the traded securities at BSE, in the case of the first two for all standard order sizes. In our analysis we use daily data, the used sample is from the 01. 01. 2007 – 16. 07. 2010 period.

We already mentioned that BLM theoretically cannot deal with the case when the complete order cannot be executed immediately. In practice, to avoid missing BLM values, the system calculates as if the order could be executed at the last non-empty price level, so practically it assumes infinite number of orders at this level. Thus this liquidity measure actually underestimates the real state of liquidity and can produce misleading values particularly for large order sizes (EUR 200 and 500 thousand) and for further calculations relying on them. As we will see, this will lead to some unrealistic results.

3. 2. Modeling

In this section we discuss the technical tools used for applying the above method. These are not necessary for understanding the results, they are presented for the sake of completeness and traceability.

In order to account for the clustering volatility of returns and net returns, we fit the following AR(1)-GARCH(1,1) model for the time series:

$$r_t = c + \phi r_{t-1} + \varepsilon_t$$
$$\varepsilon_t = \sigma_t \cdot \eta_t$$
$$\sigma_t^2 = a_0 + a_1 \varepsilon_{t-1}^2 + b_1 \sigma_{t-1}^2$$

where η ~ IID(0,1). In particular, cases following Giot and Grammig [2005] and Stange and Kaserer [2008] we will use t- and empirical distributions. In the results of the next section we used t-distribution. The sample used to estimate the model was the first 2.5 years, while the last year was used as a control period. We calculated the daily 95% and 99% VaR using forecasts from the GARCH model, as described in Section 2. We used a rolling window of 2.5 years to continuously re-estimate the GARCH model, i.e. we estimated a GARCH model for the first 2.5 years and made a forecast for the next year, and then we repeated the procedure while rolling the sample period with one day.

The test of the correction of the risk forecasts was done the following way: the predicted VaR values for both net and normal returns were compared to the corresponding values of the control period and empirical exceedance frequencies were calculated. Then the significance of deviation from the theoretical frequencies was determined statistically using the LR-test of Kupiec [1995].

The test is the following. Let Nu denote the number of days when the (net) returns exceeded the forecasted (L-)VaR values, and N the number of days in the sample. Then the empirical exceedance frequency is Nu/N, and let α denote the theoretical frequency. The test statistic using this notation is the following:

$$LR = -2\ln\left(\left(1-\alpha\right)^{N-N_u} \cdot \alpha^{N_u}\right) + 2\ln\left(\left(1-\frac{N}{N_u}\right)^{N-N_u} \cdot \left(\frac{N}{N_u}\right)^{N_u}\right).$$

Under the null-hypothesis of H0: α = Nu/N the test statistic is chi-squared distributed with one degree of freedom. We used the test uniformly on confidence level of 95%, thus H0 was accepted if LR ≤ 3.84. This test will reject the model if the empirical exceedance frequency is significantly below the theoretical value (model overestimates risk) or significantly above (model underestimates risk).

3. 3. Results

We illustrate the above method by using the daily data of the four major Hungarian stocks (OTP, MOL, Richter, MTelekom). Our goal is to analyze the effect of taking liquidity into account.

In the following figures we plotted the VaR forecasts and normal returns for the final year and the different stocks. We plotted both the L-VaR and the traditional VaR values in order to be able to make comparison and to see the difference between the two. In the figures we used uniformly order sizes of EUR 20 thousand and EUR 200 thousand (1 denotes the former, 4 denotes the latter) and 95% VaR. The numbers on the horizontal axis show the time of the forecast (e.g. 650 means the forecast for the 650th trading day from 01.01.2007.), while the numbers on the vertical axis are the percentage values.

Figure 1: Liquidity adjusted (VaR_1, VaR_4) and traditional VaR (VaR_n) forecasts compared with actual returns for OTP (Major Hungarian Bank)

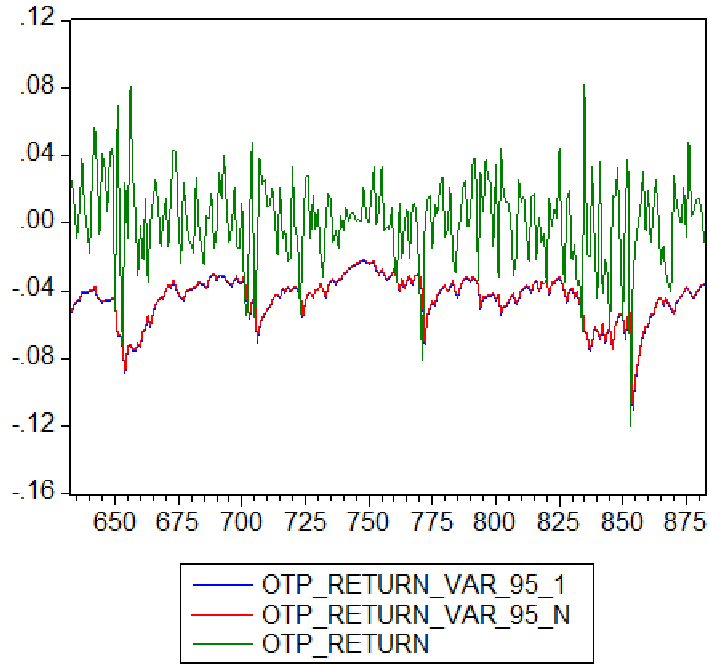

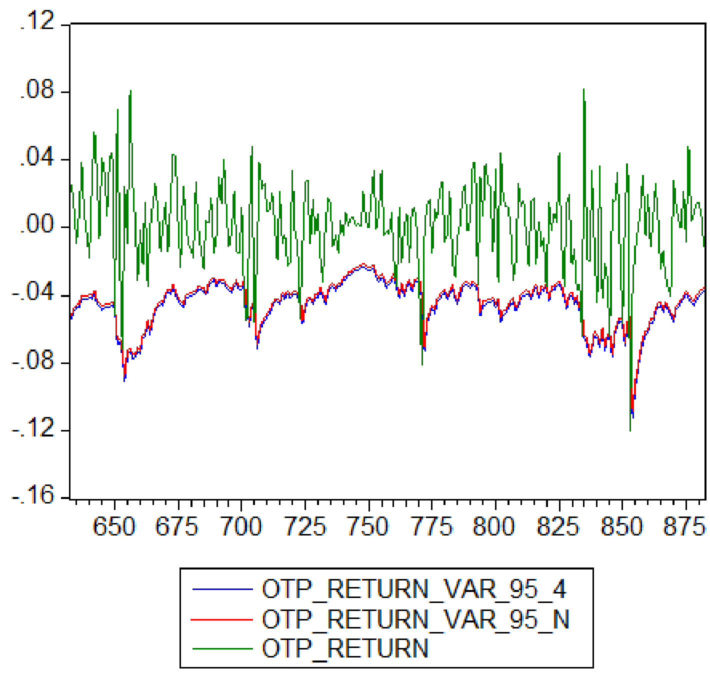

As we see in Figure 1, in the case of OTP (Major Hungarian Bank) there is no significant difference in the traditional and L-VaR values, which exactly indicates that OTP is a very liquid stock, its liquidity risk is low. Things are different if we take another stock, e.g. Richter (Hungarian Pharmaceutical Company).

Figure 2: Liquidity adjusted (VaR_1, VaR_4) and traditional VaR (VaR_n) forecasts compared with actual returns for Richter (Hungarian Pharmaceutical Company)

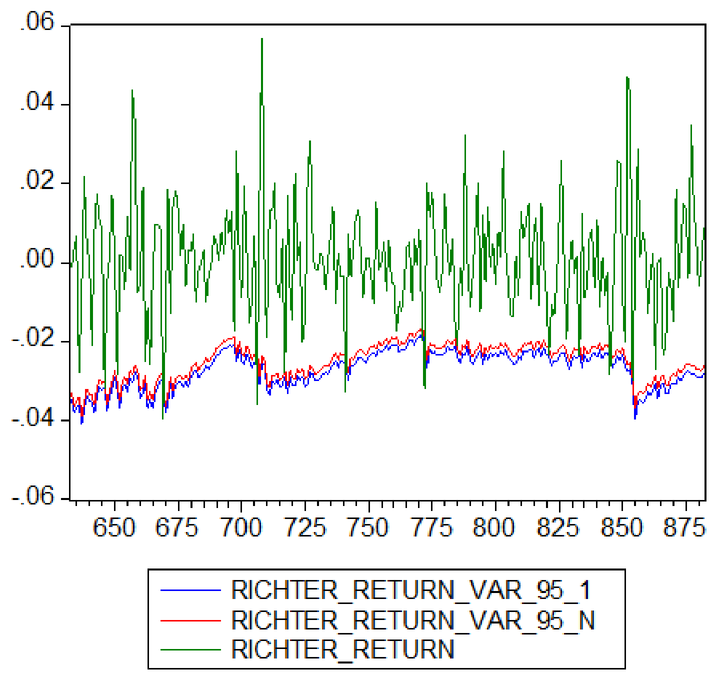

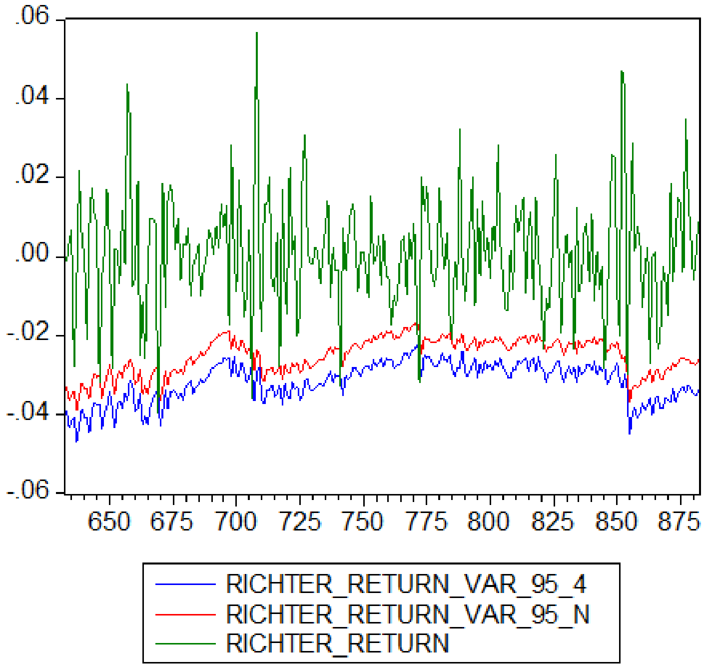

In this case even for the smallest order size, there is clearly a visible difference between the two forecasts, and this increases drastically if we move to the larger order sizes. These results show that Richter is much less liquid than OTP, its liquidity risk is high.

The other two stocks show similar issues, their figures can be found in the Appendix (Figure 8 and Figure 9).

During the test of exceedances for OTP and MTelekom both the traditional and L-VaR forecasts work properly, the empirical values do not differ significantly from the theoretical 5%. This means that in the case of OTP and MTelekom by taking liquidity into account we do not worsen the accuracy of forecasts. For Richter the situation is similar, only 99% L-VaRs for 100 and 200 order size are inaccurate, this is probably due to the above mentioned calculation problem of BLM. In the case of MOL both the 99% traditional VaR and L-VaR values are inaccurate, we get too strict forecasts - instead of the expected 1% exceedance there are in fact no exceedances at all. This is probably due to the used sample as it contains the entire period of the 2008 crisis.

To summarize, we conclude that by taking liquidity into account the accuracy of risk forecasts do not worsen.

To illustrate the difference between the traditional and L-VaR better, we look at the time series of the above defined $\lambda(q)$ relative liquidity measure for the different stocks. In the figures, the $\lambda(20)$ and $\lambda(200)$ measures are plotted simultaneously. These figures show the percentage difference between the forecasts of the previous figures (the horizontal axis shows the time of the forecast, the vertical axis shows the value of the measure).

Figure 3: $\lambda(q)$ for OTP, it shows the percentage error made, when ignoring liquidity

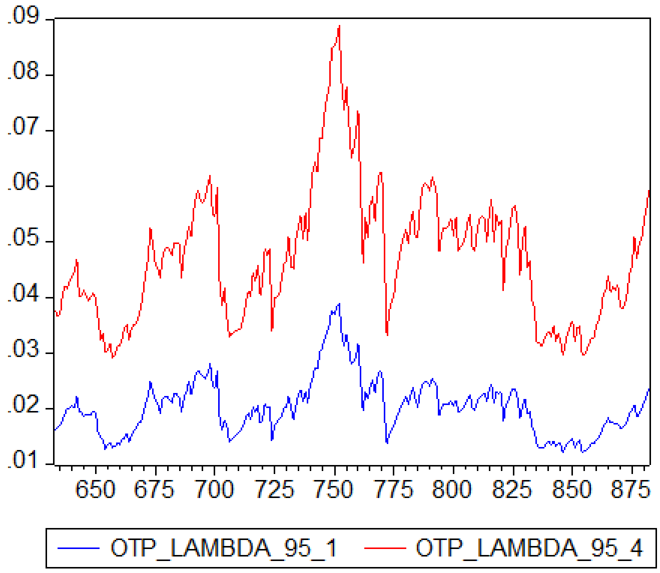

From the relative liquidity measures of OTP we can conclude on one hand that by increasing the order size liquidity risk increases significantly; this is expected intuitively as the implicit costs of a larger order are clearly greater. On the other hand if we examine the values of $\lambda(q)$ we see that for the smallest order size liquidity risk is always above 1%, but can be up to 4%, while for the order size of EUR 200 thousand liquidity risk is always above 3% and can go up to as high as 9%. This is the added risk we ignore if we concentrate only on mid-price risk. While these values may not be very large, we should bear in mind that OTP is (one of) the most liquid stocks at BSE.

For Richter we have the following figure:

Figure 4: $\lambda(q)$ for OTP, it shows the percentage error made, when ignoring liquidity

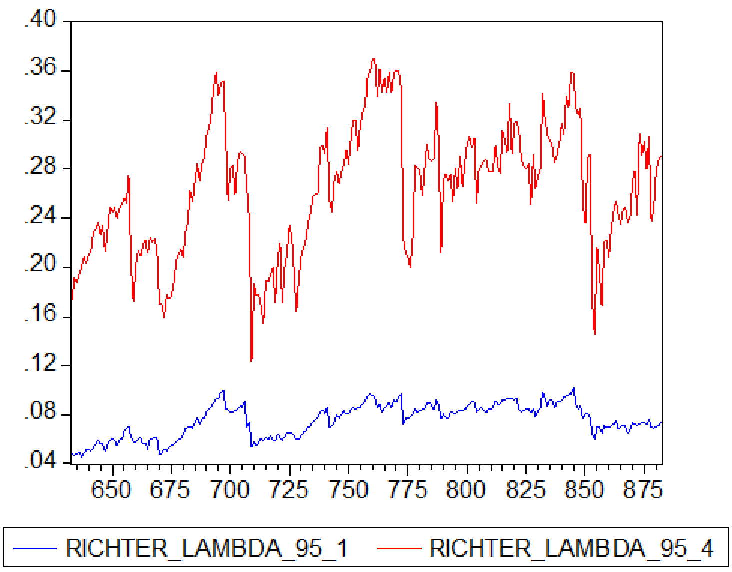

As we see, for Richter liquidity risk is significantly greater, even for the smallest order size it is always above 4%, but often reaches 8%, while for the order size of EUR 200 thousand it is stably above 20%. This backs up numerically our previous conclusion from Figure 1 and Figure 2 that Richter is much less liquid than OTP and it has significant liquidity risk.

The relative liquidity measures of the remaining stocks can be found in the Appendix (Figure 10 and Figure 11).

In Figure 5 the relative liquidity measures of the major Hungarian stocks are compared for the smallest order size. The liquidity order of OTP-MOL-Richter is clearly visible, as expected. The significant difference among them, however, shows that only OTP is really liquid at BSE.

Figure 5: From liquidity to illiquidity: $\lambda(q)$ for OTP, MOL and Richter respectively

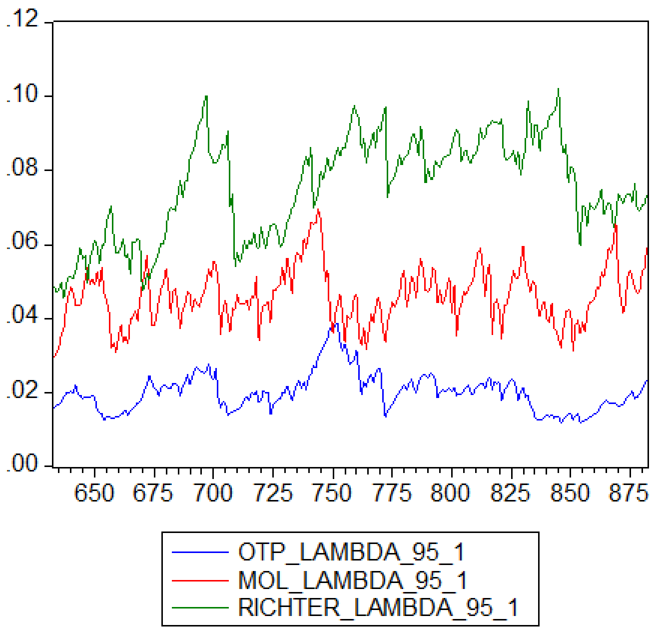

It is worth looking at the average values of the above relative liquidity measures for the different stocks and order sizes. Table 1 summarizes these values.

Table 1: Average $\lambda(q)$s for different stocks, order sizes and 95% - 99% forecasts

95%	OTP	MOL	Richter	MTelekom	99%	OTP	MOL	Richter	MTelekom
20	2,03%	4,61%	7,54%	8,46%	20	1,25%	3,07%	4,65%	4,78%
40	2,41%	5,76%	9,57%	11,29%	40	1,47%	3,90%	6,40%	6,38%
100	3,36%	8,91%	15,71%	18,78%	100	2,03%	6,25%	11,33%	10,71%
200	4,72%	13,86%	26,29%	31,98%	200	2,83%	10,03%	18,00%	18,49%
500	8,40%	29,75%	91,74%	133,52%	500	5,04%	22,24%	60,73%	97,43%

From the average values clearly show the liquidity ranking of the stocks. OTP proves to be the most liquid again (has the smallest liquidity risk by far).

In the table above the methodological feature of the BLM appears; we get unrealistic values for Richter and MTelekom (even above 100%!) for the largest order size. This phenomenon is the consequence of the order book not being deep enough, i.e. total limit orders in the book do not reach EUR 500 thousand on average, thus transactions of this size could not be executed in reality.

As a conclusion we can say that the above results show that liquidity risk is not irrelevant, it is highly advised to take it into account when calculating VaR measures.

4 Conclusion and outlooks

The analysis of this paper shows that by taking liquidity into account risk increases significantly even in the case of the largest and most liquid stocks. It must not be ignored.

BLM and the method presented above offer an easy and rapid way to incorporate liquidity in capital requirement calculation. Bearing in mind the deficiencies and calculation methodology of the BLM the results should be treated with care. Nevertheless, the presented model is able to reproduce main empirical observation like OTP is by far the most liquid stock at BSE, therefore we advise its integration into risk management systems.

However, prior to introducing L-VaR is the use of smaller order sizes for the BLM should be considered.

Furthermore in this study we only focused our analysis on single stocks, in practice the focus is usually on portfolios however. Due to the standard (and fixed) order sizes used in the calculation of the BLM and to the correlation among the liquidity of particular stocks, it is not evident how to extend the above method. A possible solution of the fixed order sizes is to fit a properly chosen function to the given BLM values, this way estimating the liquidity measure for order sizes other than the standard ones.

It should be also emphasized that although BLM covers two important aspects of liquidity (tightness, depth, and breadth to some extent) it does not capture the dynamic dimensions, i.e. resiliency and immediacy. From the BLM figures it is not obvious whether the total transaction can be executed immediately or not and if not then how much does liquidity worsen because of the delay. As far as we know, such a liquidity measure that covers this aspect effectively is yet to be proposed.

Beside these issues further interesting areas of research are the following:

- At the calculation of net returns we assumed the symmetry of the bid and ask sides. This assumption can be relaxed using the provided data. We can calculate net returns for both the bid and ask side by defining a bid-side BLM:
- $BLM_{bid}(q) = LP + APM_{bid}(q)$.
- The calculated VaR measures may be made even more accurate, with a more proper way of choosing the quantiles, e.g. using results from Extreme Value Theory (EVT).
- The provided intraday data is yet to be used.

Appendix

Figure 6: The principle of the calculation of BLM (XLM); Source: Deutsche Börse AG

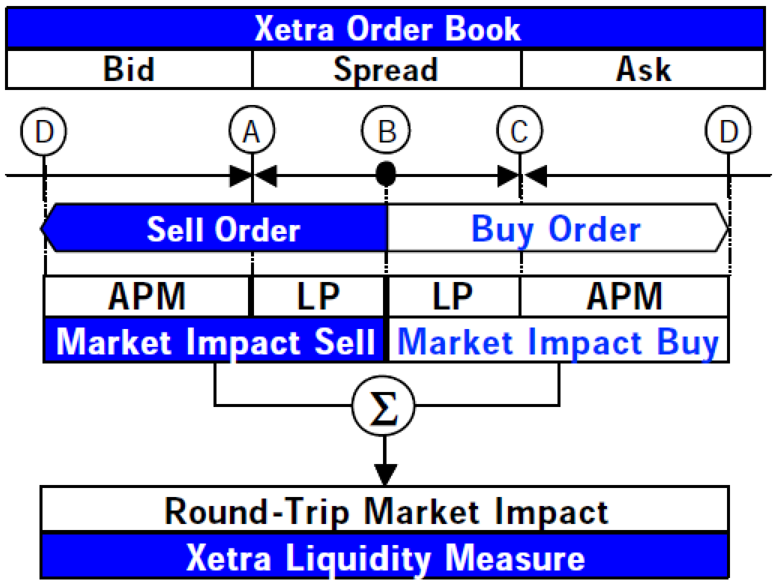

Figure 7: Visual representation of the calculation of BLM. The grey area measures the total implicit cost, BLM is this divided by the size of the transaction; Source: Stange and Kaserer [2008]

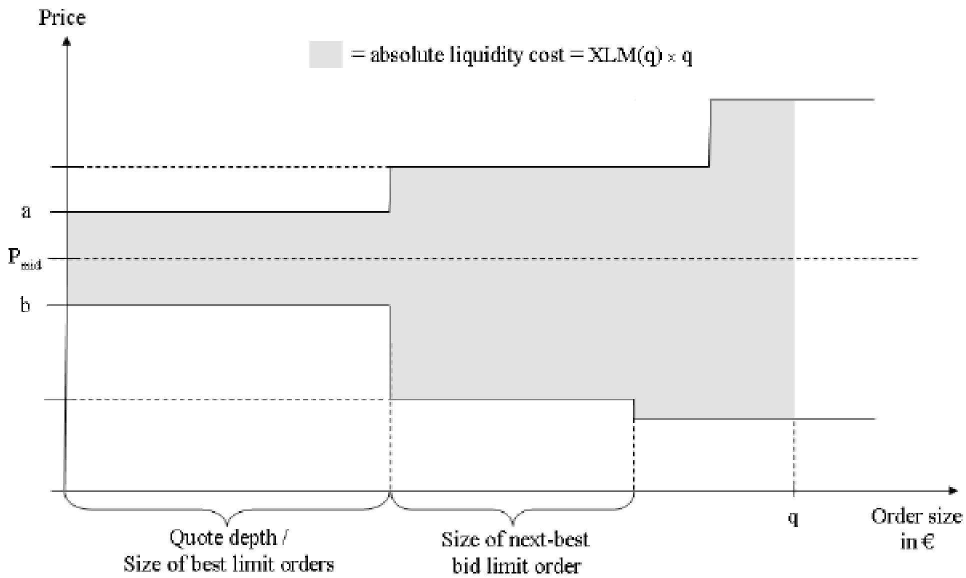

Figure 8: Liquidity adjusted (VaR_1, VaR_4) and traditional VaR (VaR_n) forecasts compared with actual returns for MOL (Hungarian Oil Company)

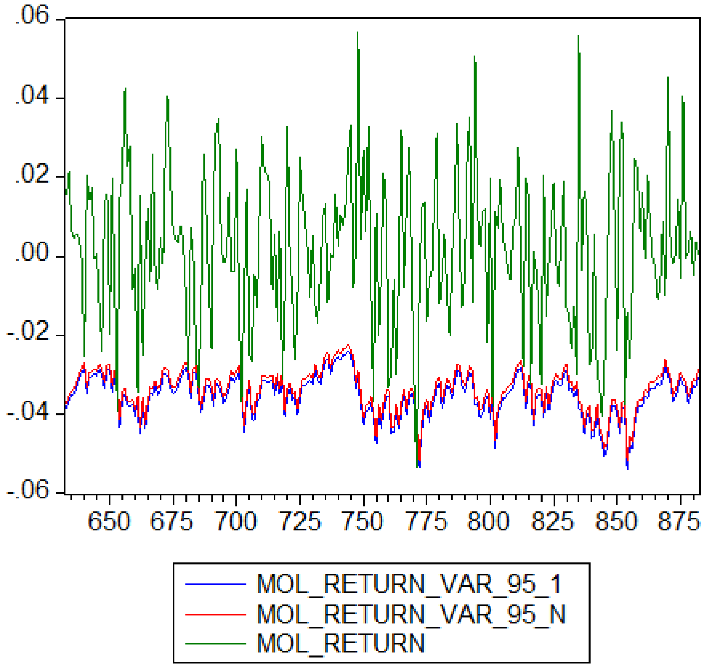

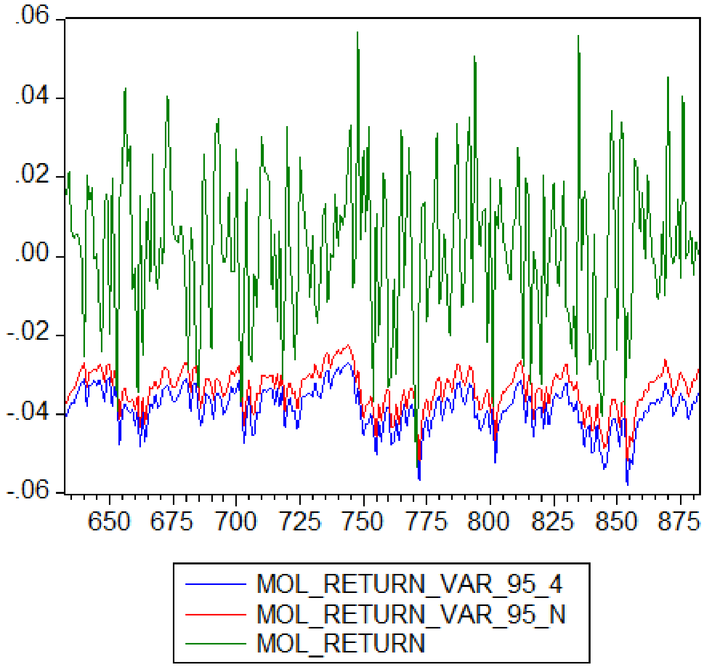

Figure 9: Liquidity adjusted (VaR_1, VaR_4) and traditional VaR (VaR_n) forecasts compared with actual returns for MTelekom (Hungarian Telecom Company)

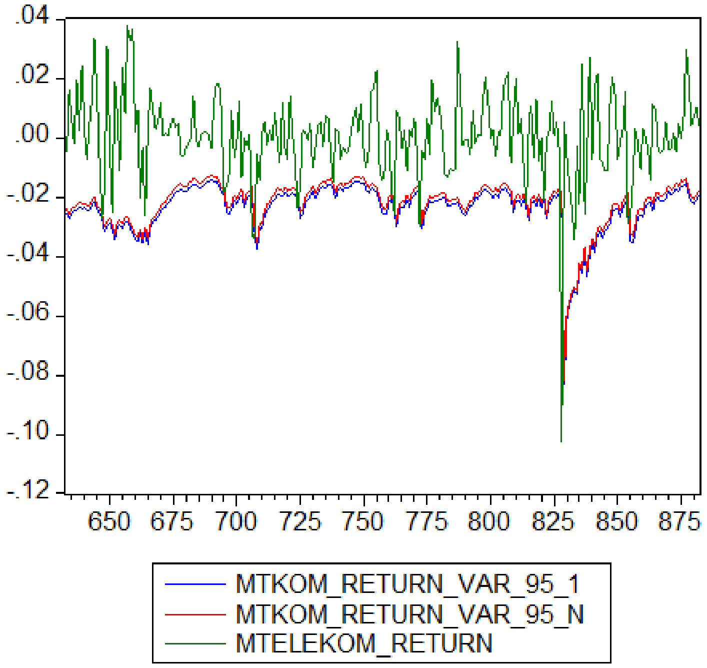

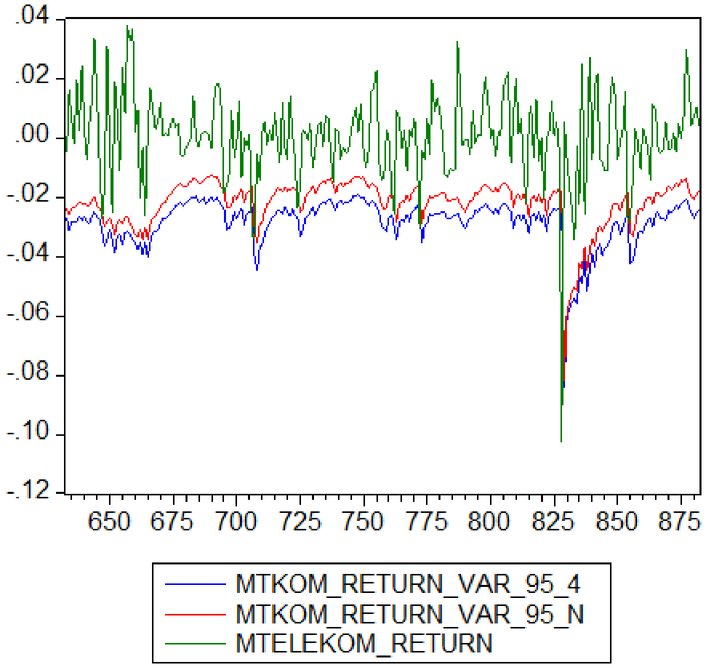

Figure 10: $\lambda(q)$ for MOL, it shows the percentage error made, when ignoring liquidity

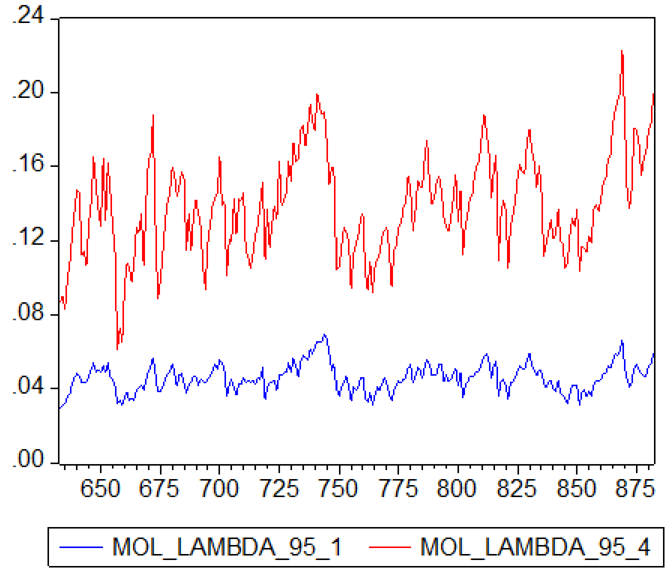

Figure 11: $\lambda(q)$ for MTelekom, it shows the percentage error made, when ignoring liquidity

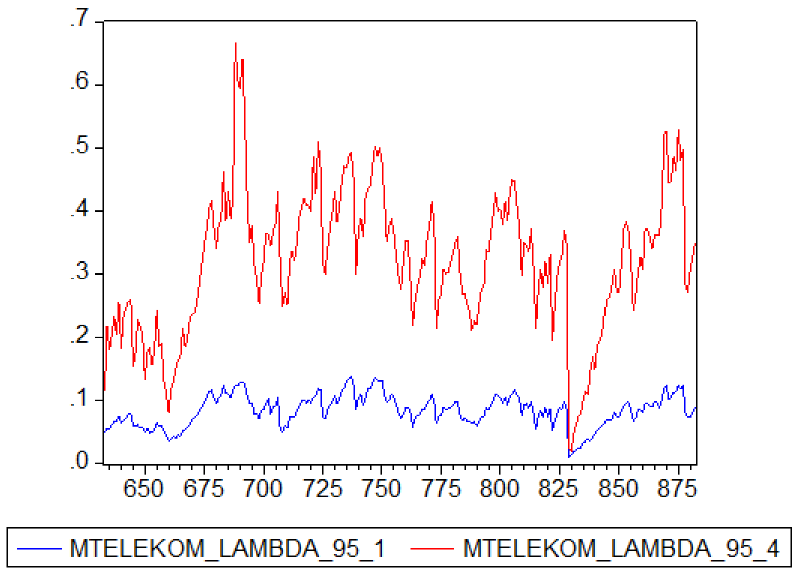

References

BANGIA, A., F. X. DIEBOLD, T. SCHUERMANN ÉS J. D. STROUGHAIR [1999]: Liquidity on the Ouside, Risk, 12, pp. 68-73

GIOT, P. AND J. GRAMMIG [2005]: How large is liquidity risk in an automated auction market? Empirical Economics, 30(4), pp. 867-887

KUPIEC, P. [1995]: Techniques for verifying the accuracy of risk management models. The Journal of Derivatives, 3, pp. 73-84

KUTAS, G. AND R. VÉGH [2005]: A Budapesti Likviditási Mérték bevezetéséről. Közgazdasági szemle vol. LII. July-August 2005, pp. 686-711

STANGE, S. AND C. KASERER [2008]: Why and How to Integrate Liquidity Risk into a VaR-Framework. CEFS working paper 2008 No. 10, available: http://ssrn.com/abstract=1292289, downloaded on 2010. 11. 08.