

# Measuring Liquidity in the TLX Market

Alessandro Beber, University of Amsterdam

Marco Pagano, University of Naples Federico II, CSEF and CEPR

November 2008

<b>1. Introduction</b> .....	<b>2</b>
<b>2. Alternative measures of liquidity</b> .....	<b>3</b>
2.1. <i>Ex-ante</i> measures of transaction costs .....	4
2.1.1. Quoted bid-ask spread.....	4
2.1.2. Depth.....	5
2.1.3 Liquidity index.....	5
2.1.4 Weighted average spread (or Market Performance Index) .....	6
2.2. Alternative <i>ex-post</i> measures of transaction costs .....	9
2.2.1. Volume weighted effective spread (VWES).....	9
2.2.2. Price impact .....	9
<b>3. Possible developments of the MPI indicator</b> .....	<b>10</b>
3.1. Measurement and release frequency .....	10
3.2. Type of information reported to the market.....	11
3.3. Choice of the reference quantity $K$ .....	14
3.4. Alternatives to choosing a single reference value for $K$ .....	16
3.4.1. Alternative 1: average MPI.....	16
3.4.2. Alternative 2: multiple reference sizes .....	17
3.4.3. Alternative 3: confidence intervals .....	18
3.5. Round-trip versus one-side liquidity costs.....	18
<b>4. Bringing further dimensions into the picture</b> .....	<b>20</b>
4.1. Computing a market-wide TLX MPI indicator .....	20
4.2. Comparing more than two trading venues .....	21
4.3. Accounting for stock market fees .....	22
4.4. Extending the comparison to different types of markets .....	25
<b>5. Conclusions</b> .....	<b>26</b>

# 1. Introduction

Assessing the liquidity of security markets is of essence for any trader who is planning to submit an order. A liquid market is one where trades can be executed with low transaction costs (O'Hara, 1995), and depending on these costs an investor may decide to alter the size of his orders, the timing and mode of their execution and the trading venue where they are routed to, so as to achieve “best execution”. However, measuring these costs is not simple, precisely because they vary depending on the order's size, its timing and the venue where it is placed. As a consequence, a variety of measures are employed to evaluate a market's liquidity.

At least as importantly, the information needed to measure transaction costs is not always promptly available to securities traders. Indeed the availability of such information is a key feature of a market – its transparency. More specifically, the market's “ex-ante transparency” is the availability of information that help market users to estimate the trading costs that they would face when placing an order, whereas the availability of data about past trades defines the market's “ex-post transparency”. Transparency is particularly important to market users when they can choose across alternative trading venues, as it allows them to perform an informed choice between these venues: indeed comparable and timely information about the trading costs of alternative trading venues is at the very basis of the principle of “best execution”.

Currently, the choice between alternative trading venues is no longer limited to wholesale securities trading: retail investors too can choose between alternative platforms to route their orders. In Italy, retail investors can currently trade fixed income securities either in the TLX platform or in the MOT platform managed by Borsa Italiana. Therefore, providing timely, comparable and informative measures of liquidity in these alternative venues is key to allow these investors to achieve best execution.

The present document aims to provide a basis for the design and evaluation of such liquidity measures, in light of both academic literature and market practice. We perform our assessment in two distinct but complementary steps. In Section 2, we assess alternative measures so as to highlight their comparative strengths and limitations, and on

this basis identify some indicators that are particularly informative. In Section 3, we discuss the statistical implementation and dissemination of these indicators, for a variety of financial instruments. In Section 4, we discuss their potential extensions to (i) a setting with multiple instruments being traded on each platform, (ii) a setting with more than two trading platforms for each financial instrument, (iii) the inclusion of explicit trading costs, and (iv) the comparison between market-driven and order-driven market structures. Section 5 concludes by summarizing the main points of the document.

## **2. Alternative measures of liquidity**

In this section we compare alternative *ex-ante* and *ex-post* measures of transaction costs used in the academic literature and in market practice, indicating in each case the data that are required for their computation and assessing their comparative strengths and limitations. In this discussion, we shall classify liquidity measures according to whether they measure *ex-ante* or *ex-post* liquidity.

Measures of *ex-ante* liquidity are based on hypothetical trades at prevailing market conditions. They have the advantage of addressing directly the key problem in order placement strategies: “if I am to place an order of a certain size in this market, at which price will it execute?”. However, *ex-ante* liquidity measures also suffer from a bias, insofar as they do not take into account that orders themselves affect market conditions, chiefly by affecting the market price: for instance, to the extent that an order is regarded as informative by other market participants, it will affect their valuation of the security, and hence both the bid and ask prices that will prevail after the execution of the order.

Therefore, an alternative approach is to consider the liquidity costs experienced by *actual* past trades, or in other words *ex-post* liquidity measures. These measures have themselves an obvious limitation: as market conditions change continuously, the conditions at which recent trades have been executed can offer limited guidance as to which conditions may currently prevail in the market. For instance, the fact that recent transaction prices imply very wide spreads does not necessarily imply that such spreads

will persist, insofar as they may attract more aggressive limit orders, which will enhance market liquidity.

## 2.1. *Ex-ante* measures of transaction costs

### 2.1.1. *Quoted bid-ask spread*

The bid-ask spread is a commonly used measure of market liquidity. It directly measures the roundtrip cost for a small trade, that is, the cost of buying and immediately selling a small quantity of the security. This cost is typically calculated as the difference between the ask price  $P_A$  (at which a small buy order is executable) and the bid price  $P_B$  (at which a sell order of the same size is executable). This difference is scaled by the bid-ask midpoint, or mid-price  $M \equiv (P_A + P_B)/2$  to obtain the percentage bid-ask spread  $s$ , that is, the roundtrip transaction cost per euro (or dollar) traded:

$$s = \frac{P_A - P_B}{M} = \frac{P_A - P_B}{(P_A + P_B)/2}$$

This measure can be calculated quickly and easily with data that are widely available on a real-time basis. Fleming (2003) documents that the bid-ask spread is a useful tool for assessing and tracking liquidity for the U.S. Treasury market.

However, a drawback of this measure of transaction costs is that bid and offer quotes are good only for orders of very *small* size. It is unsuited as a measure of the cost of executing a sizeable market order. In a limit order book market (henceforth, LOB), an order of large size may end up executing against limit orders placed at increasingly unfavorable prices, hitting higher ask price levels if it is a buy order or lower bid prices if it is a sell order. Indeed, a sufficiently large market order may even exhaust the entire set of limit orders outstanding on the limit order book. In this case, the order will not be executable, or at least it will not be entirely executed. Therefore, for any order above a minimal size the quoted bid-ask spread fails to provide guidance both to the magnitude of transaction costs and to execution risk.

### 2.1.2. Depth

A simple measure that helps assessing execution risk is the quantity of securities that can be traded at the bid and offer prices, that is, the depth of the market. A simple estimate of this quantity is the total quantity of securities that is explicitly bid for or offered for sale at the posted bid and offer prices, which is the total depth of the market on the sell and buy side, respectively. As already mentioned, this statistic may be useful for traders who wish to place large orders, and for them may complement the information provided by the quoted bid-ask spread. Yet, even when taken together, these two measures fail to provide an estimate of the transaction cost to be expected when placing a large order.

### 2.1.3 Liquidity index

The market liquidity index (MLI) is a measure designed to capture the tradeoff between quoted bid/ask spread and market depth (Bollen and Whaley, 1998). It is the ratio between the average depth at the prevailing bid and ask quotes and the percentage quoted spread, that is:

$$\text{MLI} = \frac{(\text{depth at bid} + \text{depth at ask})/2}{\text{percent spread}}$$

where the depth at bid and the depth at ask are measured in monetary amounts (euros or dollars) rather than nominal amounts, and the percent spread is the quoted bid/ask spread divided by the bid/ask price midpoint.

To illustrate how this market liquidity index works, suppose that on TLX the bid and ask price quotes are 99.875 and 100.125, respectively, with a depth of 1,000 on each side of the market. Then the MLI on TLX is:  $[(1000+1000)/2]/[0.25/(99.875+100.125)/2] = 400,000$ .

Now, suppose that the competing platform has a narrower spread of 0.125 instead of 0.25, but the market depth is now only 500 on each side of the market. The alternative platform's MLI is:  $[(500+500)/2]/[0.125/(99.9375+100.0625)/2] = 400,000$ , indicating no difference between the quality of the two markets. In other words, although the quoted

spread is lower on the alternative platform, the market depth has been reduced commensurately, so that there is no change in market quality.

The liquidity index MLI has the advantage of condensing information about both spreads and depths in a single number. The drawback of the liquidity index is that it does not have an intuitive meaning. Furthermore, the MLI does not use information from the order book other than the depth at the best prices.

#### *2.1.4 Weighted average spread (or Market Performance Index)*

A natural way to overcome the limitations of the quoted bid-ask spread is to extend its logic to an order of any predetermined size  $K$ , where size is defined in terms of nominal value of the instrument, for instance the face value of the bonds to be bought or sold. The idea is to take the difference between the price  $P_A(K)$  at which one can currently buy a quantity  $K$  of securities and the price  $P_B(K)$  at which one could sell the same quantity:  $P_A(K)$  will be the sum of the ask prices multiplied by the corresponding quantities available on the LOB until the total quantity  $K$  is purchased, and  $P_B(K)$  will be the sum of the bid prices multiplied by the corresponding quantities until the quantity  $K$  is sold. In the same way as one would do for the quoted spread, this difference is standardized by the mid-price corresponding to the total quantity  $K$ , which in this case is the arithmetic average of the ask and bid prices just described, i.e.  $M(K) = [P_A(K) + P_B(K)] / 2$ .

Therefore, this measure is the analogue of a percentage bid-ask spread referred to an order of nominal size  $K$ . It is known as the *weighted average bid-ask spread* referred to the quantity  $K$ , since the spread is computed by weighting bid and ask prices by the respective quantities available on the LOB up to the point where the sum of the orders

achieves the predetermined size  $K$ .<sup>1</sup> As this measure is computed by TLX and labeled *Market Performance Index*, we shall equivalently refer to it by the acronym MPI.<sup>2</sup>

Formally, it is defined as:

$$s(K) \equiv \text{MPI}(K) = \frac{P_A(K) - P_B(K)}{M(K)} = \frac{P_A(K) - P_B(K)}{[P_A(K) + P_B(K)]/2}$$

where  $P_A(K) = \sum_{i=1}^n q_{Ai} P_{Ai}$  and  $P_B(K) = \sum_{j=1}^m q_{Bj} P_{Bj}$ ,  $i$  and  $j$  index the level of the LOB on

the bid and ask side respectively,  $P_{Ai}$  and  $P_{Bi}$  are the bid and ask prices at the  $i^{\text{th}}$  LOB level, and  $q_{Ai}$  and  $q_{Bi}$  are the amounts used to weight those prices. The quantity  $q_{Ai}$  is

defined as  $q_{Ai} = \min\left(Q_{Ai}, K - \sum_{r=1}^{i-1} Q_{Ar}\right)$  such that  $\sum_{r=1}^{i-1} Q_{Ar} \leq K$ , where  $Q_{Ai}$  is the

amount that can be bought at the  $i^{\text{th}}$  LOB level. Symmetrically,

$q_{Bi} = \min\left(Q_{Bi}, K - \sum_{r=1}^{i-1} Q_{Br}\right)$  such that  $\sum_{r=1}^{i-1} Q_{Br} \leq K$ , where  $Q_{Bi}$  is the amount that can

be sold at the  $i^{\text{th}}$  LOB level.

This measure of transaction costs is more appealing and flexible than the quoted bid-ask spread, because it overcomes its limitation of applying only to very small orders, as it takes explicitly into account the quantities that are available for trading on the order book when a specific quantity  $K$  needs to be traded. Hence,  $K$  must be chosen so as to be representative of the costs that are relevant for the typical customer. (Section 3.3 below specifically addresses the issue of representativeness of  $K$ ).

By the same token, this measure also captures the depth dimension, for two reasons. First, if the market features little depth, then the limit ask (bid) prices available on the LOB will rapidly increase (decrease) as a function of the traded quantity  $K$ , so that the MPI will

---

<sup>1</sup> In the French equity market, this statistic is labeled *fourchette moyenne ponderée* and is commonly used by practitioners as a measure of liquidity.

<sup>2</sup> Currently TLX compares the MPI in its own market and in the competing MOT market for each financial instrument to determine which trading venue offers the lower spread, and provides to its customers the proportion of times in which the MPI of the TLX market is lower than or equal to that of the MOT market.

also increase rapidly as a function of  $K$ . In other words, observing a large MPI even for a low quantity  $K$  indicates that the market lacks depth. Second, it may even happen that the entire quantity  $K$  cannot be bought or sold on the market, so that the MPI cannot be computed for that quantity. Therefore, also the impossibility to compute the MPI signals lack of depth.

Compared with the MLI, the MPI has the advantage of being a more intuitive measure of transaction costs. The MPI is an *ex-ante* measure of immediate liquidity for a given trade size, which shows the investor the upper bound of transaction cost for immediate execution of a trade of that size. It is well understood in the academic literature that liquidity is a transaction-size-specific concept. For example, Jones and Lipson (2001) report that liquidity for institutional trades worsened following NYSE's reduction of the tick size, while liquidity for small trades improved. However, it must be recognized that the MPI is more arbitrary than the MLI, in the sense that it requires specifying a traded quantity  $K$ , whereas the MLI does not.

While conceptually MPI is appropriate for any market, it is most informative in markets where market orders are mostly executed against a limit order book, and least accurate in markets where most trades are negotiated. This is because the precision with which MPI describes the market rises with an increase in the proportion of committed liquidity relative to the total liquidity supply. In a pure limit order market without intermediaries and hidden limit orders, all immediate liquidity is committed. In this situation, MPI is the most informative.

Moreover, the MPI has a number of advantages in dealing with common market imperfections. For example, it is not very sensitive to liquidity holes (prices on the grid where limit orders could be placed, but are not) in the limit order book.<sup>3</sup> This issue is particularly relevant when market makers temporarily leave the market, generating holes at the best quotes. The MPI in this market will be higher than in the other market, but this event will simply count as one of the observations recorded from the LOB. In contrast,

---

<sup>3</sup> Biais, Hillion and Spatt (1995) first observed this phenomenon when they have found holes in the limit order book of the Paris Bourse.

other indicators, such as the bid-ask spread, could be severely affected by these market imperfections.

## **2.2. Alternative *ex-post* measures of transaction costs**

There are several *ex-post* liquidity indicators. Those that seem more appropriate within the current framework are the volume-weighted effective spread and the price impact.

### *2.2.1. Volume weighted effective spread (VWES)*

The effective cost for a security on a specific day is computed as a trade-weighted average for all trades relative to the prevailing quote midpoint. Unlike the MPI, the VWES does not require to specify a traded quantity  $K$ . The main drawback is that it can be computationally intensive, because it requires using both all trades and all quotes. It can also be misleading if TLX is compared with a market where average traded quantities are very different. In that case, the VWES can be computed for trade-buckets to insure appropriate comparisons.

### *2.2.2. Price impact*

A popular measure of liquidity, suggested by Kyle (1985), considers the rise (fall) in price that typically occurs with a buyer-initiated (seller-initiated) trade. If data on price changes and signed trades are available for each security  $i$ , one can estimate a representative price impact coefficient as the  $\lambda_i$  coefficient in the regression:

$$\Delta p_t = \lambda_i \left( \text{signed} \sqrt{\text{dollar volume}} \right)_{it} + \varepsilon_{it}.$$

The specification is estimated using price changes and signed volume aggregated over a certain time intervals, typically five-minutes. The main drawback is again that it is computationally intensive, requires econometric estimation, and is probably less intuitive than the other measures. Again, the advantage is that the specification of a quantity  $K$  is not required.

### 3. Possible developments of the MPI indicator

The previous section indicates that the MPI indicator has several advantages as a compact and flexible measure of transaction costs. However, in principle the MPI can be computed any number of times per day for each financial instrument, thereby generating a time series of MPI values for each trading day or week. This naturally raises several issues, which will be addressed in this section:

- (i) Sampling and computation frequency: How often should one collect the data needed to compute the MPI from the LOB and calculate it?
- (ii) Choice of statistics: Which statistic should one compute to aggregate the MPI over time for a given financial instrument? And how should these statistics be aggregated across different financial instruments?
- (iii) Defining the reference quantity  $K$ : Which criteria should one use to establish this quantity, and should the MPI be computed for one or more such reference quantities?
- (iv) Dealing with asymmetries in the bid-ask spread: Can one define separate measures for the liquidity that one can expect to find on the buy and the sell sides of the market?

#### 3.1. Measurement and release frequency

If the MPI is to be released on a daily basis to market users, the sample size for each day should be of at least 30 observations,<sup>4</sup> given that market conditions – and especially trading costs – can (and do) change considerably during the trading day, as highlighted by many studies in market microstructure based on high-frequency, intra-daily observations. The measurement frequency can be lowered if the release of statistics to market users is less frequent (for instance, weekly): conversely, it should be raised if one wishes to release infra-daily estimates of transaction costs. For instance, if one wishes to release statistics three times per day, one would need to make sure that each release is preceded by at least 30 observations.

---

<sup>4</sup> This is often considered by statisticians as the typical sample size beyond which one can apply asymptotic theory.

The sampling frequency is important if one is also to provide a measure of the statistical significance of the MPI mean difference between the two markets, as suggested in the subsequent section. Of course, in deciding the magnitude of the sample to be collected one should also take into account the occurrence of outliers: the more likely the presence of such influential observations, the larger the sample to be collected for the sample to be representative of the underlying population.

### **3.2. Type of information reported to the market**

In deciding which statistics about the MPI is to be released to the market, one has to balance the richness of informational detail with the need to provide synthetic indicators that are easy to grasp and compare across markets. We consider three types of statistics:

- (i) *average* MPI, and test statistic for its equality across platforms;
- (ii) *median* MPI, and test statistic for its equality across platforms;
- (iii) *frequency* of cases in which the MPI is lower in one market, and test statistic for its significance.

The *average* MPI in each of the two markets (and the difference between the two) has a very intuitive meaning for investors: it tells them on which platform they can expect to find greater liquidity. However, it is important to complement this measure with an indication of the confidence level that investors can place in this difference. For instance, if the MPI has a very large variance either in one or in both markets, what might look like a considerable liquidity advantage of one of the two markets may not be statistically significant at conventional confidence levels. To inform customers about the confidence that they can place on the mean MPI difference, it is important to also publish the confidence level with which this difference is different from zero, or simply mark the mean difference with 1, 2 or 3 stars if it is significantly different from zero at the 10, 5 or 1 percent confidence level (as routinely done in econometric studies). It should be noticed that this confidence measure (based on the usual Student's *t*-test of the null hypothesis that the means of two normally distributed populations of MPIs are equal) requires a normal distribution of MPI.

The *median* MPI in each the two markets (and the difference between the two) can also be a helpful statistic. Indeed, this statistic is more reliable than the mean if the distribution of MPI is highly asymmetric, so that the non-normality assumption fails. Clearly, the normality assumption can be tested on a sample of MPI observations to identify the more appropriate statistical approach. The median is also more reliable in the presence of outliers (or influential observations), possibly generated by anomalous market events. For example, sometimes during the day for some financial instruments, market makers could temporarily leave the market generating anomalous spreads.

If the observations used in the estimation include such abnormally large spreads, they could distort the average MPI. However, this is not necessarily a good reason to delete these outliers from the sample: if the customer trades indeed during these anomalous market conditions, then the MPI indicator should not ignore those observations because they reflect potential trading costs that the customer would have to pay in certain cases. It would be more advisable to use the median, which is not as sensitive to their presence as the average. By the same token, in these cases it is preferable to rely on a non-parametric test such as the Wilcoxon signed-rank test than on a *t*-test to establish the statistical significance of the MPI difference.

Beside the mean and the median, another potentially useful statistic to compare the liquidity of TLX with that of another trading venue is the *observed frequency* of cases in which the MPI is lower on TLX. Just as the median, this statistic is not very sensitive to the presence of outliers due to liquidity holes. Moreover, it does not require the distribution of MPI to be normal. Therefore, like the median, it is a more robust statistic than the average difference in the MPI. Another advantage of this statistic is its intuitive meaning for market users.

However, reporting *only* the proportion of cases in which the MPI is lower on TLX may provide a very incomplete picture of the comparative liquidity on the two markets, since this statistic provides no indication as to the *magnitude* of the cost advantage or disadvantage from trading on TLX relative to trading on MOT.

To illustrate why the magnitude may be important in assessing relative market quality, consider for instance a situation where in 50 percent of the occurrences the MPI is lower on TLX than on MOT, but only by a narrow margin, but it is considerably larger in the 50 percent of occurrences when it is dominated by MOT. For concreteness, suppose that on average  $\text{MPI}_{\text{TLX}} - \text{MPI}_{\text{MOT}} = -5$  basis points when liquidity is better on TLX, whereas it is 20 basis points on average when liquidity is worse on TLX. In this case, the average MPI difference between the two markets is  $-5 \times 0.50 + 20 \times 0.50 = 7.5$  basis points, indicating that on average transaction costs are larger on TLX than on MOT. Note that in this example the median MPI difference does not distinguish between the two platforms, since the median MPI is identical in the two markets.

This highlights the point that for a customer it may be important to know not only how frequently transaction costs are lower on TLX but also how much lower they are on average. For instance, in the example just reported, the mean MPI difference and the frequency with which the MPI is lower on TLX would yield opposite indications, and this would be informative for a potential customer.

Finally, even for the frequency of successes, it is advisable to estimate and disseminate a statistic to assess its reliability. To this purpose, one could report the result of a binomial test, rather than simply the observed frequency.

This frequency statistic would also help to cope with the cases in which executing a trade of size  $K$  is impossible, because the quantities available on the LOB are not large enough. In these cases the MPI cannot be computed, unless an arbitrarily large transaction cost is assumed in case of non-execution. Suppose a trade cannot be executed in market A but it can in market B. It is reasonable to conclude that the transaction cost is higher on market A than on market B, even though the difference in medians or means between the MPIs in the two markets cannot be determined. The occurrence of these cases is obviously more likely the larger is the reference amount  $K$  for which the comparison is being made. This highlights a general point regarding the MPI computation. Suppose that  $K$  is chosen to capture the most representative size of customer trades is key. If it is often the case that this typical order of size  $K$  cannot be fully executed, then the frequency with which it can be executed on one market but not on the other is important information that


customers should be aware of. If instead this is a rare occurrence, then this becomes a marginal concern. So reporting information about the frequency with which one market dominates the other is more important the larger is the typical order size  $K$  compared to the typical order book depth. This naturally leads to the next point, that is, the choice of  $K$ .

### **3.3. Choice of the reference quantity $K$**

Computing the MPI with a single quantity  $K$  has the advantage of being simple to communicate and understand. However, it might be inaccurate whenever a customer needs to trade a different quantity. It becomes important then to define the quantity  $K$  in the most objective manner and, possibly, in a way that is significant for most customers. The choice of  $K$  should be made on the basis of data about the typical order size, which may well differ across market segments – government bonds, corporate bonds, supranational bonds and sovereign bonds.

To exemplify, we report an analysis of detailed transaction data for one recent week of trading. The analysis reveals that in a sample week the average trade size was about € 47.000 for Italian/EU government bonds and € 20,000 for the remainder. The graph in Figure 1 illustrates a frequency histogram of trade sizes for Italian/EU government bond category for our sample week, ignoring sporadic trades above € 300.000.

The statistics computed from our sample week reveal that the most frequent trade for the Italian/EU government bond category is € 10.000, the 25<sup>th</sup> percentile trade is € 6.000, the median trade is € 15.000, the 75<sup>th</sup> percentile trade is € 40.000 (implying that the distribution has strong positive skewness). The average is not the best central tendency measure in this case, since the largest outliers push up the average. The mode is also not appropriate in this context, given the extremely wide range of trade sizes that investors can submit. The preferred measure of central tendency is then the median trade size, which is less affected by outliers than the average trade size. It would be advisable to complement the median with the 25<sup>th</sup> and 75<sup>th</sup> percentile. The reference median on which the MPI indicator is computed should be updated over time, though at low frequency: for example, in each month it could be computed from data for the previous month.



**Figure 1. Size distribution of trade sizes on TLX**

The MPI needs to be meaningful for the average TLX customer. It is possible that TLX quotes larger bid-ask spread than MOT for small quantities and tighter ones for greater levels into the LOB. If this is the case, setting a very large  $K$  would make TLX look better than it is for most of the customers, who place considerably smaller orders. Of course, also the opposite may be the case: it may be that TLX offers better liquidity than MOT for smaller quantities. In either case, the picture offered may be quite sensitive to the choice of  $K$ .

The data also suggest that even though the number of trades is very different across segments (for instance, Italian government bonds trade 5 times more frequently than other EU bonds), the category “Italian/EU government bond” is relatively homogeneous in terms of trade sizes. Still, the segment T02 (BOT) has median (average) trade sizes that are 33% (16%) larger than the segment T01 (Italian government no BOT). It is likely that within T01, bonds with shorter maturities attract larger sizes. Therefore, it would be sensible to adopt a different  $K$  for the different segments of Italian government bonds.

The other category “corporate, sovereign and supranational” is much more heterogeneous than Italian/EU government bond”. If we consider only the most active segments (at least 400 trades per day), the average trading size goes from €18.600 (segment E06: structured corporate bonds) to €26.300 (segment E04: floating rate corporate bonds). Therefore, this represents another example where it would be sensible to adopt a different  $K$  for the different segments of the corporate bond category.

A final note of caution concerns the use of the *nominal* value for the choice of  $K$ . This can be very misleading for zero-coupon bonds with long maturities and for securities issued in non-euro currencies. For these securities, it is more appropriate to use the euro amount of the transaction or to set specific values of  $K$  for each of these categories.

### **3.4. Alternatives to choosing a single reference value for $K$**

To a certain extent, choosing a single reference trade size  $K$  is bound to make information about liquidity inadequate for many customers: a “one size fits all” criterion will be inevitably unsatisfactory if customers have very different trading needs. In this subsection we explore various alternatives that would remove the dependence of the indicator on the choice of a single value of  $K$ , and hence would overcome this limitation of the indicator.

#### *3.4.1. Alternative 1: average MPI*

Like the single- $K$  approach described above, this methodology preserves the simplicity of providing customers a single MPI for instrument/market, but it does not require choosing somewhat arbitrarily the traded quantity  $K$ . It requires to average MPI for all  $J$  quantities  $K$  between a minimum and maximum and then calculating the average MPI. Formally:

$$\text{MPI}^* = \frac{1}{J} \sum_{j=1}^J \text{MPI}(K_j).$$

This approach coincides with the approach of computing MPI for the average  $K$  *only if* MPI is a linear function of  $K$ . This is unlikely to be the case, since the empirical evidence on different security markets suggests that liquidity is not a linear function of trade size.

This approach can be specialized using a *weighting scheme*, which reintroduces some arbitrariness but may increase representativeness. Weights by quantity are an easy approach, but they would assign great importance to large  $K$ 's and therefore may be unsuited to a mainly retail-oriented customer base.

Possibly, a more sensible approach is to assign weights based on the observed number of trades with a specific size  $K$ . The weights could refer to the previous day, previous week, or previous month, just to keep them in line with the actual distribution of trades in the TLX customer base. Formally,

$$\text{MPI}_W^* = \frac{1}{\sum_{j=1}^J n_j} \sum_{j=1}^J \text{MPI}(K_j) n_j ,$$

where  $n_j$  can be the number or euro value of trades with size  $K_j$  .

#### 3.4.2. *Alternative 2: multiple reference sizes*

In the tradeoff between simplicity and representativeness, trying to produce a single statistic favors simplicity. The alternative is to sacrifice this simplicity somewhat, and compute and communicate MPI for more than a single quantity  $K$ , so that customers with different trading needs may better evaluate which platform offers the more liquid venue for his/her intended trade.

Our recommendation here would be to compute MPI for three or five quantities, reflecting very small, small, medium, large, and very large trades, where these size categories are to be based on the 10<sup>th</sup> percentile, 25<sup>th</sup> percentile, the median, the 75<sup>th</sup> percentile, and the 90<sup>th</sup> percentile of typical trades to the platform. The set of typical trades for which these summary statistics would be computed should be rolling and defined on the previous week, or two weeks, or month of trading, so that the evolving features of the market can be taken into account.

### 3.4.3. *Alternative 3: confidence intervals*

The general intuition for this approach is to come up with a quantity threshold  $K^*$  below which TLX is better than the alternative platform in a certain proportion  $x$  of occurrences (where the proportion  $x$  can be chosen to be arbitrarily high, say 90%). The implementation of this approach requires that the dominance of TLX over an alternative platform is (or tends to be) monotonically decreasing for larger trade sizes, i.e. TLX is less and less liquid with larger quantities compared to the alternative platform.

This approach changes the nature of the problem: instead of choosing  $K$  to compute MPI, this approach requires the choice of  $x$ , which is far less subjective given that it must be set to a representative value like 90%. In terms of communication, this seems like a very simple way of describing liquidity, because it is intuitive and it is going to be relevant for a large number of TLX clients (all customers trading less than or equal to  $K^*$ ). In this sense, the construction of this indicator resembles the KiloVaR indicator for risk developed by Tradinglab.

There are two drawbacks of this approach. The first is that before being implemented it requires an empirical investigation to understand to what extent TLX dominates for smaller quantities. This is important also to understand empirically whether the threshold  $K^*$  is large enough to be a strong marketing tool for TLX. The second drawback is that the indicator is blind about liquidity for traded quantities that are larger than  $K^*$ . In this sense, it is important to test this alternative carefully on the available data.

## **3.5. Round-trip versus one-side liquidity costs**

When submitting a buy or sell order of size  $K$ , an investor may be more interested in the cost of that specific trade on one side of the market rather than in the MPI, which is the cost of a roundtrip trade. This may be particularly important in markets where the depth of the LOB is asymmetric on the two sides of the market. It may even happen that, in comparing the liquidity of two trading platforms, each of them may have greater liquidity on one of the two sides of the LOB.

In order to measure separately differential buy-side and sell-side costs, one can resort to two simple additional statistics: (i) on the buy side, the difference between the amount

that a customer pays on one market (say, MOT) and the amount he would pay on the other (say, TLX); (ii) on the sell side, the (negative of the) difference between the amount that a customer receives on one market and the amount he would receive on the other. If the difference is scaled by the price paid on TLX, in both cases it measures the percentage gain conferred by trading on TLX rather than on the competing market. For instance, on the buy side, this measure would give the “per-euro discount” obtained in the purchase of the amount  $K$ , if the figure is positive. Of course, if the figure is negative, it becomes a “per-euro premium” paid in the purchase of the amount  $K$ .

Suppose that we denote these two measures as  $\Delta_A(K)$  on the buy side and  $\Delta_B(K)$  on the sell side. Then, if we denote by an asterisk the variables referring to the competitor of TLX (say, MOT), the two measures are defined as:

$$\Delta_A(K) = \frac{P_A^*(K) - P_A(K)}{P_A(K)} \quad \text{and} \quad \Delta_B(K) = -\frac{P_B^*(K) - P_B(K)}{P_B(K)},$$

where  $P_A \equiv \sum_{i=1}^n q_{Ai} P_{Ai}$  and  $P_A^* \equiv \sum_{j=1}^{n^*} q_{Aj}^* P_{Aj}^*$ , and all the notation is analogous to that used for the MPI indicator.

Let us show a simple example of calculations of the two measures  $\Delta_A(K)$  on the buy side and  $\Delta_B(K)$  on the sell side. Suppose that we observe a 99-101 bid-ask price for a specific bond on TLX and 97-102 bid-ask price on MOT, and assume for simplicity that both the depth and the order size  $K$  is equal to one. The two measures are computed as follows:

$$\Delta_A(1) = \frac{102 - 101}{101} = 0.01 \quad \text{and} \quad \Delta_B(1) = -\frac{97 - 99}{99} = 0.02,$$

denoting about a 1% discount for buy orders on TLX versus MOT and about a 2% gain for sell orders on TLX versus MOT.

Also for each of these two indicators, of course, one can choose between reporting the average differential, the median differential or the frequency with which the differential

is positive. The considerations are exactly the same as those expounded above with reference to the MPI, and therefore they will not be repeated here.

## **4. Bringing further dimensions into the picture**

All the discussion so far has focused on the comparison between the liquidity of two markets for one instrument. In practice, customers may be interested in a richer type of comparison. First, they may want to have an estimate of liquidity of the TLX market as a whole, and compare it to its analogue for a competing market, say MOT. Second, they may want to compare TLX with two or more alternative trading venues. Third, customers may be interested in comparing not just the “implicit transaction costs” arising from a wide bid-ask spread, but also the “explicit transaction costs” specific to each market, an important component of which are the stock exchange fees to be paid for a given order. Finally, in some cases the alternative trading venue available for trading may be qualitatively different from TLX, that is, be a dealer (or quote-driven) market rather than another LOB (or order-driven) market. Each of these more complex comparisons would require an appropriate methodology and additional information. In this section, we explore which additional information might be required, and in which format this additional information may be usefully disseminated to the customer base.

### **4.1. Computing a market-wide TLX MPI indicator**

How do you aggregate the MPI indicators computed for individual financial instruments into a market-wide measure? Since the MPI is in percentage terms, a market-wide MPI can be computed as a weighted average of the single instrument averages, where the weights should represent the importance of each instrument  $j$  in the market.

$$\text{MPI}_W^* = \frac{1}{\sum_{j=1}^J n_j} \sum_{j=1}^J \text{MPI}(K_j) n_j$$

The weights, which must be relatively stable and common between markets being compared, can be obtained from the average transacted volume over the previous month on TLX and MOT.

The weighted average has two potential drawbacks. First, it might be tricky to choose common weights across alternative trading platforms if financial instruments that are prevalent in one venue are relative less important in the other. Second, the choice of weights is necessarily arbitrary and, as such, subject to potential criticism.

For these two reasons, it might be sensible to use just a *simple* average of the MPIs computed on each financial instrument. In this case, there is a natural adjustment that increases the significance of liquidity measures on bonds that are more frequently traded by the customer base. In general, the more important categories (e.g., Italian government bonds) feature a larger number of traded instruments than the less important categories (e.g., structured corporate bonds). As a result, even in a simple average, the liquidity of the more relevant financial instruments gets implicitly weighted more.

The  $MPI_W^*$  or the  $MPI^*$  in its simple average version will give an indication of liquidity for a market as a whole. The comparison to its analogue for the competing market should be performed according to the same guidelines offered in the previous section for the MPI computed at the single instrument level. For example, the customer base might be interested in the median liquidity difference between TLX and MOT and in the statistical significance of this difference.

## **4.2. Comparing more than two trading venues**

Allowing customers to compare TLX with other markets beside MOT may become quite important as the enactment of the Markets in Financial Instruments Directive (MiFID) may spawn a proliferation of alternative trading venues for the same financial instrument, especially considering that the technical setup costs of a new trading platform nowadays are not very high.

As always when the set of possible options becomes larger, the danger is to overwhelm customers with a deluge of information that they would not be able to process at the speed that is typically required in the design of an optimal order placement strategy. For instance, one could provide the difference of MPI in TLX and each of the other trading platforms concerned, together with statistics to evaluate the statistical significance of each difference.

But this may be too much information for customers to be able to process it, especially if MPI is provided for several values of  $K$ , and/or it is complemented by other measures of liquidity. As an alternative, one may provide only information on how MPI (and/or any other chosen liquidity indicator) in TLX compares with (i) the trading platform offering best execution at that point in time and (ii) the trading platform offering the mean (or median) MPI at that point in time.

### **4.3. Accounting for stock market fees**

In practice, total trading costs may be heavily affected by the “explicit trading costs”, which are formed by the sum of post-trading fees for clearing and settlement services, stock exchange fees, transaction taxes and intermediaries’ commissions. In some cases these “explicit trading costs” could exceed the “implicit trading costs” typically measured by the bid-ask spread or the price impact of trades. Moreover, they can vary greatly across markets and financial instruments. For example, based on the estimates of Domowitz et al. (2001), the total costs for Italian equities amounted to 34.8 basis points between 1996 and 1998, of which 75% are explicit costs (including broker commission costs and fees).

Since best execution refers to the total trading costs borne by customers to carry out a given transaction, TLX should consider the possibility of providing information also about its comparative “explicit trading costs”. This would typically reduce to the comparison between the post-trading and stock market fees applying to a given order. Indeed, in the choice between two trading venues placed in the same fiscal jurisdiction, transaction taxes do not differ between trading venues. Intermediaries’ commissions for the same type of order should also not differ systematically across venues, and in any

event providing such information to customers is best left to the intermediaries themselves, not least because information about commissions released by a trading venue may easily be abused to perturb competition and might come under the purview of competition authorities.

Effecting a comparison of stock market fees for a given order is not complicated if the fee schedules of the two competing exchanges are relatively simple, for instance if they both charge a fee proportional to the value traded or to the number of units of the financial instrument being traded (for instance, number of shares). It can become very complicated if in one of the two markets (or in both) the fee schedule is non-linear and depends on several qualitative characteristic of the order (for instance, whether it is a market order or a limit order). Then attempting to compare stock market fees for a “typical order” runs the risk of becoming a deceptive oversimplification: for instance, if in one of the two trading venues limit orders and market orders are charged very different fees, which one would qualify as the “typical” order to implement the comparison?

Insofar as the specific comparison between TLX and MOT is concerned, the comparison can be performed only subject to making quite specific assumptions, due to the non-linearity of the fee schedule used by the two exchanges: as we shall see TLX charges a fee per contract, and MOT a membership fee per year plus a fee per contract, so that in neither case the fee is simply proportional to the trade size. Ignoring transaction taxes and intermediary fees (for the reasons mentioned above) leaves us with a comparison of the post-trading costs of the TLX platform with competing trading venues.

It is useful to understand the magnitude of these costs in the same vein as Domowitz et al. (2001) compare total explicit and implicit costs. The post-trading cost per contract on the TLX platform is merely €0.20. Remarkably, TLX does not charge any other fee beside those required to cover post-trading costs. To put this figure in perspective, the implicit trading cost for TLX Italian government bonds, as measured by the average cost of a round trip transaction for the most frequent trade of €10.000 on a sample day is €3.50.<sup>5</sup>

---

<sup>5</sup> The implicit cost is obtained as the average MPI computed on all Italian government bonds traded on the 13<sup>th</sup> of February 2008 on the TLX platform.

It turns out that the explicit cost in this case is only about 6% of the implicit costs. This is an upper bound, since for trades exceeding € 10.000 and securities that are less liquid than Italian government bonds, implicit costs are going to be much larger, and therefore post-trading costs are going to be an even tinier fraction of implicit costs. Therefore, accounting for post-trading costs to obtain the total trading costs on TLX is likely to produce a figure very close to the MPI itself.

Let us describe nevertheless the methodology to adjust the MPI indicator to include an explicit cost per contract:

$$\text{MPI}^{\text{adj}}(K,c) = \frac{P_A(K) - P_B(K) + c \times 100}{M(K)},$$

where all the notation is defined as above and  $c$  is the explicit cost per contract. For the TLX platform,  $c$  is equal to €0.20 and there are no annual membership costs, so that the computation of the adjusted MPI is trivial. Suppose that the ask price is 101, the bid price is 99, the traded size is € 1,000 and the depths at the best prices are large enough to accommodate this traded quantity. The unadjusted MPI indicator would be equal to 2%, corresponding to an implicit trading cost of €20. The adjusted MPI would be equal to 2.02%, corresponding to a total trading cost of €20.2.

For the MOT platform,  $c$  depends on the number of contracts per year traded by the intermediary and there is an annual membership fee. Since it is not feasible to take the actual situation of every intermediary into account, one must choose as a benchmark an intermediary trading a specific amount of contracts. This requirement inevitably introduces another subjective decision besides the choice of the quantity  $K$  in the computation of the adjusted MPI. If the benchmark intermediary trades at least 1,000,000 contracts per year, the average cost per contract arising from the membership fee would be €0.0396, besides which one would have to consider the €0.40 cost per contract charged by MOT. In this case, the appropriate cost  $c$  to include in the adjusted MPI calculation for MOT would be 0.4396, over twice the TLX corresponding figure. So TLX appears to have a competitive advantage over MOT regarding explicit transaction costs, although not a large one (about 20 basis points) due to the tiny magnitude of these

costs in both markets. This advantage would be slightly larger if the assumed benchmark intermediary were to trade less than 1,000,000 contracts per year, since this would raise the average cost per contract on MOT.

#### **4.4. Extending the comparison to different types of markets**

Some of the financial instruments traded on the TLX market – such as corporate bonds – are also traded in Over-The-Counter, dealer markets. This raises one of the thorniest questions in the area of market microstructure, *viz.*, how to compare the trading costs of a LOB market, where order execution is centralized and based on strict price and time priority rules and price discovery is order-driven, with those of a dealer market, where order execution is decentralized and price discovery is quote-driven. The comparison involves several difficulties.

First, a dealer market does not feature a single bid-ask spread applying to a given order size but as many bid-ask spreads as there are active dealers. However, it is not even clear that one should focus on any particular dealer (for instance, the one offering the tightest quoted spread), because a customer may be able to split his/her order among several dealers, which implies that the trading cost faced by the customer would be some kind of weighted average of the set of dealers patronized by that customer.

Second, the bid and ask prices quoted by dealers may differ from those actually offered to their customers, since dealers may negotiate a “price improvement” to their customers relative to the quoted prices. This implies that the “effective spread” of a dealer market may well be tighter than the quoted spread of any dealer in that market. Insofar as these improved bid and ask prices are generally negotiated over the phone and are not publicized *ex ante*, it is impossible to compare the true *ex-ante* liquidity offered by a dealer market with that offered by a LOB, where instead such a measure is easy to construct – the MPI being precisely one such measure.

Of course, comparing the best quoted bid-ask spread provided by a dealer market with the MPI for the corresponding order size in a market such as TLX may still provide customers with some vague *ex-ante* information about the comparative liquidity of the

two markets. But, almost by definition, it will be a potentially very inaccurate comparison. A more reliable form of comparison between a LOB auction market and an OTC dealer market can instead be effected based on *ex-post* measures of liquidity, and therefore on transaction prices – as it is done by the VWES and the price impact measures described in Section 2 above.

## 5. Conclusions

In this study, we have examined various indicators that can be used to measure the *ex-ante* and *ex-post* liquidity of a security trading platform and compare it with that of another platform where the same security is traded.

Among the *ex-ante* measures of liquidity, we have singled out the weighted average percentage spread (or Market Performance Indicator - MPI) as being particularly informative and flexible. We have then discussed the statistical implementation of this indicator, that is: the sampling and computation frequency, the choice of statistics, and the choice of the reference trade size (or sizes) for which the spread is to be computed.

The most important conclusions of the analysis are that it is advisable to: (i) compute and disseminate both the average (or median) value of the spread in the two competing markets as well as a measure of the frequency with which it is lower on TLX than on the competing market; (ii) provide also measures of the statistical significance of difference between the mean (or median) spreads in the two markets and of the frequency with which it is lower on TLX; (iii) pay great attention to the choice of a reference trade size that is truly representative of the trading needs of the market users, and consider the possibility of multiple reference sizes.

Our analysis also offers some indications as to how one can extend the comparison of liquidity to (i) a setting with multiple instruments being traded on each platform, (ii) a setting with more than two trading platforms for each financial instrument, (iii) the inclusion of explicit trading costs, and (iv) the comparison between market-driven and order-driven market structures.

## References

Benston, G., P. Irvine, and E. Kandel, 2000, Liquidity Beyond the Inside Spread: Measuring and Using Information in the Limit Order Book, Working Paper, Emory University.

Biais, B., Hillion, P., Spatt, C., 1995. An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse. *Journal of Finance* 50, 1655--1689.

Bollen, N. and R.E. Whaley, 1998, Are “Teenies” Better?, *Journal of Portfolio Management* 25, 10-24.

Cohen, K., S. Maier, R. Schwartz, and D. Whitcomb, 1981, Transaction Costs, Order Placements Strategy and Existence of the Bid-Ask Spread, *Journal of Political Economy* 89, 287–305.

Domowitz I., Glen J. and Madhavan A., 2001, Liquidity, Volatility and Equity Trading Costs Across Countries and Over Time, *International Finance* 4, 221-255.

Fleming, M. J., 2003; Measuring Treasury Market Liquidity, *Federal Reserve Bank of New York Policy Review*, 83—108.

Hasbrouck, J., 2006, Trading Costs and Returns for US Equities: Estimating Effective Costs from Daily Data, Working Paper, New York University.

Kyle, A. S. 1985, Continuous Auctions and Insider Trading, *Econometrica* 53, 1315-35.

Jones, C., Lipson, M., 2001, Sixteenths: Direct Evidence on Institutional Execution Costs, *Journal of Financial Economics* 59, 253-278.

O’Hara, M.. 1995. *Market Microstructure Theory*. Cambridge: Blackwell.