An alternative approach to measuring the liquidity risk of public investors' investment assets

David Doran, Steve Kilkenny, Šarūnas Ramanauskas and Alex Shablov[1]

Abstract

Public investors – particularly central banks – often apply different criteria, compared with private investors, when deciding how best to allocate their reserves across the range of assets eligible for investment in their own portfolios. While return is undoubtedly an important factor, other criteria, in particular liquidity, often take on greater importance when managing public funds. Furthermore, where the public investor's mandate extends to managing reserves in order to facilitate effective management of exchange rates, the investment criteria can take an even more specific focus, with liquidity risk of particular importance. In recent times, the measurement of liquidity risk has attracted far more attention given the lessons of the Great Financial Crisis; both in terms of understanding actual liquidity when a crisis occurs, and in terms of perceived liquidity in the context of the subsequent expansionary central bank policies and large-scale asset purchase programmes. Given these developments, and evolving market conditions, this paper considers an alternative approach to measuring the liquidity risk of public investors' investment assets. The approach set out in the paper allows for a consideration of liquidity risk in the context of risk assessments and asset allocation decisions that are specific to the mandate and policy objectives of public investors, and central banks in particular. To demonstrate the approach, the paper applies the methodology to the euro area sovereign market and, in doing so, appears to track the deterioration in liquidity during a period of political or fundamental uncertainty while also indicating an increase in liquidity (or a reduction in liquidity risk) following the announcement of the ECB's asset purchase programme. Other potential applications of the methodology, from a central bank risk management perspective, are also considered.

[1] The authors are Head of Financial Risk Management and Senior Risk Analysts, respectively, in the Central Bank of Ireland. The views expressed in this article are solely the views of the authors and are not necessarily those held by the Central Bank of Ireland or the European System of Central Banks. The authors would like to thank Glenn Calverley and Ruth Gleeson for helpful comments as well as Steve Flanagan and Naoise Metadjer for their assistance with the data. Any remaining errors or omissions are our own.

1. Introduction

Public investors – particularly central banks – often use a different set of criteria, compared with private investors, when deciding how best to allocate their funds across the range of assets eligible for investment in their own portfolios. While return is undoubtedly an important factor, other concerns, such as liquidity and capital preservation, often take on greater importance when managing public funds. Furthermore, where the public investor's mandate extends to managing reserves in order to facilitate effective management of exchange rates, the investment criteria can take an even more specific focus, with liquidity risk being of particular importance.

For the purposes of this paper, a security can be considered as liquid if its acquisition or disposal can be executed relatively quickly and with a minimal effect on the market price. There are a number of factors that may determine the liquidity of a particular security; for example, fixed income securities are generally less liquid compared with equities, since the latter are normally traded on open exchanges and the former largely in less transparent over-the-counter (OTC) markets (Laganá et al (2006)). Given that public institutions primarily invest in fixed income securities, liquidity risk should therefore be of particular importance from a risk management perspective. Within the fixed income asset class, bond liquidity is further determined by the subsector to which the issued security is assigned; for example, sovereign bonds are generally considered more liquid than corporate, covered or other non-government bonds (Galliani et al (2014)).

In addition, specific security features such as the amount issued, rating, duration and time to maturity (Galliani et al (2014)), as well as whether they were issued as part of a private placement (Amihud and Mendelson (1988)), further distinguishes between the actual and perceived liquidity of fixed income securities. It therefore follows that different bonds issued by the same issuer may have significantly different liquidity risk profiles. Additionally, the concept of "on-the-run" or benchmark bonds also determine a bond's relative liquidity; bonds that have most recently been issued should be relatively more liquid (Pasquariello and Vega (2009)). Securities that exhibit characteristics of illiquidity will usually experience a direct and visible impact on their trading patterns. Illiquid securities generally have lower traded volumes and display wider bid-ask spreads (Favero et al (2010)).

Heightened liquidity risk, in the form of an inability to dispose of a security within a required time frame and with limited impact on the price, is a risk to which public investors may have differing degrees of sensitivity. Firstly, there might be a negative impact on the security's price, resulting in potential loss of capital. Secondly, an inability to dispose of a security within the required timeframe might result in an inability to meet liquidity cash-flow requirements. However, it is important to note that one cannot consider risk in isolation of potential investment returns, and illiquid securities should command a liquidity premium. Therefore, there is some potential to earn higher returns if there is an appetite to take on more liquidity risk in a manner consistent with an investor's risk appetite. Private placements and hold-to-maturity (HTM) portfolios are examples of this.

However, the understanding of liquidity risk measurement has arguably changed as a result of the Great Financial Crisis (GFC); both in terms of assessing actual liquidity when a crisis occurs, and also in terms of perceived liquidity following the subsequent emergence of a benign environment driven by expansionary central bank policies and large scale asset purchase programmes. Regarding the latter, measures of liquidity risk can be clouded by fewer market participants, in some respects, yet offset by the presence of large-scale central bank purchases, which has provided a guaranteed market purchaser for many fixed income assets. It is noteworthy to recall that conditions prior to crisis, wherein issuer ratings and an overly benign outlook proved not to be representative of the underlying risks and subsequent liquidity risks when market conditions deteriorated. As such, in considering the development of market conditions, public investors should question the accuracy of current measures of liquidity risks, particularly in relation to the universe of assets in which they typically invest. Given the lessons of the crisis, and the effect of non-standard central bank monetary policies on market conditions, this paper considers an alternative approach to measuring the liquidity risk of public investors' investment assets. The outlined approach advances the assessment of liquidity risk in the context of asset allocation decisions that are specific to the mandate and policy objectives of public investors, and central banks in particular.

The paper is set out as follows: Section 2 discusses liquidity risk in the context of the aftermath of the GFC and the liquidity risk considerations for public investors that inform the approach proposed in this paper. Section 3 discusses the importance of the data in this approach and the algorithmic methods that can be employed to ensure its robustness. Section 4 introduces a scoring model and how it is constructed, while Section 5 discusses some of the potential uses of a scoring approach, including both analytical and risk management applications. Section 6 concludes.

2. Liquidity risk considerations for public investors

2.1 Liquidity risk since the crisis

Since the GFC, there has been an increasing focus on liquidity[2] risk as a risk factor in its own right. For example, a number of new regulatory measures were introduced in response to the crisis, such as the *Liquidity Coverage Ratio* and the *Net Stable Funding Ratio* (Basel Committee (2010)), which aim to ensure that commercial banks can better withstand stressed liquidity events. Separately, other regulations have impinged on the ability of global investment banks to act as market-makers and provide two-way markets in fixed income securities (ESRB (2016)). This is more apparent in the non-sovereign segments of bond markets and has affected liquidity in corporate bonds (see eg CGFS (2016)). It is in the context of this increased recognition of liquidity risk that all investors, including public investors, have become more interested in the liquidity risk of discretionary investment asset portfolios.

Two additional developments affecting liquidity risk since the GFC are noteworthy. Firstly, in recent years, global financial markets have experienced a historically low interest rate environment – particularly in the euro area[3], but also in

[2] Liquidity can be divided into two types. Funding liquidity is the ease with which a market participant can meet its obligations as they fall due. Market liquidity is the ease with which an asset can be bought or sold, at something close to its current market price. For the purposes of this paper, our main focus is on market liquidity.

[3] For example, at the end of 2018, the German sovereign yield curve was negative out to seven years.

other major jurisdictions. This environment raises significant income challenges for fixed income focused investors such as central banks or other public investment institutions. In response to these challenges, public investors have explored alternative means of generating investment returns, albeit in a manner consistent with the typically conservative risk appetite associated with public investors.

This trend can be observed in the public investor survey data, such as UBS's *Annual Reserve Management Survey*[4], which shows that the share of reserve managers that are investing in a broader range of asset classes (other than sovereigns and supranationals) has been increasing in recent years. In doing so, public investors may enter into less populated parts of the financial markets universe. Hence, it becomes even more important for these institutions to have a means of monitoring the liquidity risk profile of the investment assets in a manner commensurate with their specific risk tolerance levels and investment criteria.

Secondly, central bank policies in recent years, in the form of quantitative easing and large-scale asset purchase programmes, have stabilised various fixed income markets, which has helped contribute to a lower financial market volatility backdrop[5]. Coupled with a search for yield in the low interest rate environment, this can induce risk-taking behaviour by financial market participants – perhaps even in a complacent manner – and this can sometimes be a precursor to a financial crisis, in the form of a so-called "Minsky moment" (Danielsson et al (2018)). The realisation of the effects of a financial crisis as a consequence of the moral hazard of a Minsky Moment may come too late for investors to exit particular areas of the market[6], due to the heightened liquidity risk at that time. This highlights the importance of a forward-looking approach to assessing liquidity, at an asset allocation level, as part of a holistic risk assessment of new investment proposals and asset allocation strategies in the context of overall risk appetite.

2.2 Liquidity risk management of public investors

Public investors can use a number of liquidity proxies to indirectly manage liquidity risk. As noted above, their traditionally conservative risk appetite has meant that they have tended to invest in large, highly rated issuers such as sovereigns, whose debt will always likely be relatively liquid and in high demand. In addition, concentration risk in particular holdings can be managed, as a proxy for liquidity risk, by applying limits on how much of a particular bond or issuer is held.

Other market risk management measurements, such as duration targets with limited deviation parameters, also seek to reduce the nominal amount of illiquid investments. In addition, the articulation of a certain level of liquidity risk appetite can

[4] See www.ubs.com/global/en/asset-management/insights/asset-class-research/asset-allocation/2018/annual-reserve-management-survey.html.

[5] For example, the VIX index reached an all-time low in November 2017.

[6] It may be that other aspects of a central bank's mandate, such as maintaining financial stability, may lead to a decision not to contribute to pro-cyclicality in markets through the disposal of assets, although investment decisions of one's own reserve assets are normally viewed through the lens of normal investment principles where practicable.

facilitate a hold-to-maturity[7] accounting approach on a portion of a public investor's reserves. Despite such risk management measures that seek to indirectly limit liquidity risk, it is not the same as specifically measuring liquidity risk in a manner that reflects the investment criteria associated with management of public sector investments and consistent with its risk appetite. Given the developments set out in the foregoing section, it can be useful for public investors to initiate a detailed evaluation of various liquidity risk management tools while considering the bespoke requirements for liquidity risk measurement appropriate for its own investment criteria and requirements.

2.3 Liquidity risk measurement requirements – some considerations

It is important that any means of measuring and managing risk (including the liquidity risk scoring approach considered in this paper) is appropriate for an institution's own needs. More specifically, it should be consistent with both its mandate and its risk appetite in the context of its investment policy principles and criteria. For example, a central bank with an explicit requirement to intervene in the foreign exchange market to manage the national currency would likely have heightened sensitivity to liquidity needs. Similarly, a public pension fund with regular cash outflow requirements would have equivalent liquidity sensitivities.

In the case of many public investors, its investment policy principles remain relatively conservative with an emphasis on avoiding investment losses taking primacy, such that the generation of return is subject to adherence to capital preservation priorities. In accordance with this, they might invest predominantly in sovereign and sovereign-like fixed income bonds. Therefore, some market available liquidity scoring methodologies that are calibrated to the broader bond market universe may not provide the best fit for public investor's investment management requirements. More specifically, when using such broadly calibrated metrics, some instruments or asset classes could be considered highly liquid, but, from a traditional central bank's (or public investors') perspective, could be considered less liquid when evaluated against its self-determined investable universe. It is therefore clear that, given its investment objectives and risk appetite, a tailored liquidity scoring methodology that could be specifically calibrated to a public investor's own investment universe might be more appropriate.

In this regard, when considering whether to develop a tailored liquidity scoring approach, the following characteristics might be considered important:

- **Transparent and easily explainable to internal, senior stakeholders.** There are many different ways to measure and manage liquidity risk, which can range from a relatively simple scoring approach to a more opaque "black box" approach, which may include machine learning techniques. Each approach has pros and cons, and can incorporate both quantitative and qualitative aspects.

7 While such an accounting treatment does not preclude selling an HTM bond, it is possible that such an action could, in certain circumstances, require marking all HTM assets to market, which would not be desirable in almost all circumstances.

A possible approach to take could be to develop a relatively transparent mechanism that could be easily explained to senior stakeholders for it to be utilised credibly in the context of risk monitoring and risk assessment considerations. Therefore, one could create a scoring model that includes a number of different observable and well understood indicators of a bond's liquidity, which combine an overall measurement of liquidity.

- **Flexible to changing requirements and risk appetite.** It is important that any risk measurement approach is fit for purpose and is robust over time. While commercial providers of liquidity risk solutions offer many benefits, they are often calibrated in a standardised way and not easily customisable.

  It might be considered important to have an approach that allows for some flexibility in how a public investor ultimately measures liquidity risk. For example, a changing liquidity environment in financial markets may necessitate a change in how the different components are calibrated – particularly during periods of systemic change. In addition, it could be considered important to have an approach that could respond to potential changes in the liquidity risk appetite of a public investor. This could be implemented, for example, by changing the composition of the institution's "investable universe" in order to generate a more representative liquidity score distribution over time.

- **Sufficiently granular for the public investor's holdings.** It is important that a liquidity measurement approach generates useful information about the liquidity of investment assets, both in and of itself, but also in comparison to other investment assets. This allows judgements to be made about each asset and its relative liquidity characteristics. In this way, the methodology should allow for the comparison of liquidity characteristics among various factors of the investment portfolio. For example, comparison among asset classes, sectors and issuers.

  Moreover, to allow for a more dynamic monitoring, the methodology should be sufficiently granular to capture the daily movements in market liquidity, as having a score or metric that is too rigid might result in larger, more volatile shifts in the time series. Finally, the methodology should allow for the differentiation among the various aspects of market liquidity, as this would help to identify the key drivers behind any changes in the market conditions.

- **Built and maintained in-house.** As with any risk management tool, consideration must be given as to whether a product could be procured from a commercial provider or developed in-house. Important considerations relate to the cost of buying the solution, versus the internal resources required to build and maintain a model and tool, as well as some of the considerations related to flexibility discussed above.

Ultimately, there can be a number of strong benefits to developing the model in-house. Firstly, many institutions have been developing their data analytics capabilities in recent years, such that in-house development of a liquidity scoring measure could leverage from these capabilities. Secondly, in-house development provides a good opportunity to enhance internal knowledge of liquidity risk, both narrowly in the area of the public investor's financial assets, but perhaps in a broader sense, over time, in other areas of its mandate, such as prudential supervision or financial stability.

It is important to note that, while a methodology can be developed in-house, such an analysis can be informed by studying the methodologies and approaches already available elsewhere. While other solutions and approaches provide many benefits, they do not necessarily align with the public investor's requirements discussed above, mainly in terms of the transparency, flexibility and granularity in alignment with the public investor's risk appetite and investment mandate.

3. Liquidity scoring data

In this section, the paper outlines an approach to sourcing, specifying and augmenting the data used for scoring liquidity.

3.1 Data source and specifications

As opposed to financial instruments traded on public exchanges, fixed income securities are predominantly traded through decentralised OTC markets. Since there is no physical trading location, the participants conduct transactions via various modes of communication such as telephone, Bloomberg chat or proprietary trading systems. As a result, trades are often completed without other participants being aware of the transaction details, which normally makes OTC markets a lot less transparent and the degree of liquidity somewhat difficult to gauge. In order to calibrate a model of market liquidity scoring or liquidity risk, a detailed and granular data set is required, covering a range of variables relating to the underlying bond transactions. In the case of OTC transactions, it can be difficult to obtain reliable and consistent bond trading data from public sources and, therefore, specialist data providers must be used.

The data set used in the model presented is procured from IHS Markit, a professional market data provider[8]. The data set used is an extract from the corporate and sovereign bond universe, which, in turn, is a subset of Markit's fixed income data universe. The extract includes data since 2012 and contains daily data collected from over 300 market makers – covering over 100,000 corporate and sovereign bonds globally[9].

The data includes hundreds of bond-level fields, ranging from the instruments' basic parameters, such as coupon frequency and issuer region – to various liquidity measures, such as shadow liquidity[10]. For the liquidity analysis, the proposed model includes measures related to market depth, trading volumes, bid-ask (yield) spread and maturity. Table 1 outlines selected data fields and provides a short field description.

The proxies for market depth are given as (i) number of data sources and (ii) number of dealers, and both are relevant. For example, if the bond's data is available from numerous different sources, but only a limited number of dealers are quoting it,

8 See https://ihsmarkit.com.

9 See https://cdn.ihs.com/www/pdf/Pricing-and-Reference-Data.pdf.

10 Shadow liquidity is applied when an instrument does not receive any market data but shares characteristics with other instruments that are liquid.

it could be that the same dealers are quoting the bond in different places, in which case the market could not be considered deep. Alternatively, if there were many dealers quoting, but these were only available in a certain area of the market, this would not be considered a deep market either.

Bond parameters selected for liquidity scoring tool — Table 1

	Field alias	Field description
	Maturity	The date on which the principal amount of the security becomes due and payable, as stated in the terms of the security.
	Bid-ask yield spread	The difference between the ask yield and bid yield.
Market depth	Number of data sources	The number of distinct data sources that were received for the instrument set on a given business day.
Market depth	Number of dealers	The number of dealers quoting the bond averaged over a number of business days.
	Trading volumes	Trade volumes observed per specific bond, issuer or market sector.

Source: Markit.

3.2 Data Cleaning

Large data sets rarely contain complete information, creating a number of challenges for researchers and practitioners. According to Kofman and Sharp (2003), 28% of publications in finance for the period 1995–1999 used data sets where, on average, 20% of values were missing. The missing data in the data set available for the model presented made up about 15% of the data utilised, which is in line with the statistics provided above. Most data gaps or consistency issues encountered were addressed by either interpolation techniques or manual replacement using expert judgement.

Given the large volume of data used in the exercise, algorithmic tools and techniques are preferable to manual interventions, which were kept to a minimum. However, outlier data points can distort the efficient calculation of a liquidity score and such data required some form of intervention. For instance, the bid-ask yield spread field had a number of data points that were either a number of standard deviations from the mean or, in some cases, the entry displayed negative spreads. Nevertheless, the vast majority of the outlier data points were normalised using an internally developed interpolation algorithm – a key element of the data processing phase.

As opposed to deleting records, data correction using the interpolation algorithm was chosen as the preferred method, since deleting records will affect other fields with valid observations and, as a result, leave gaps in the time series. The interpolation method involves three major processing steps: (i) identification of missing/outlier data points, (ii) interpolation using the bond's own time series data and (iii) interpolation at an issuer level.

In the first part, the algorithm identifies missing data points in each data set field. Next, the nearest available data points are identified based on the bond's time series. The empty/error element is then populated by the nearest available data point if it is within a pre-defined specified date range, otherwise the algorithm proceeds to the issuer level matching. The matching on the issuer level is somewhat more complex. Firstly, a bond with the nearest maturity is identified from the same issuer. In the event that the maturities differ by an amount greater than a pre-defined specified range, the record is marked for review; otherwise, the missing point is matched to a nearest point in the time-series of the bond selected in the previous step. The nearest available point must be within a certain date range or the record is marked for review. Figure 1 provides a visual representation of the interpolation algorithm steps. This interpolation method proves to be a reliable data cleansing mechanism that allows detecting, fixing or removing errors and inconsistencies from the data set, and thereby substantially improves the usability of the data set.

The interpolation algorithm is by no means static and is, indeed, an area for potential enhancement. For instance, there are number of alternative methods to be explored to deal with missing data points such as multiple imputation or predictive imputation techniques. One example is a statistical framework-based method (Yuan (2010)) that allows an analyst to sample missing points from a deduced joint distribution function and preserve data statistical properties such as mean and variance. Another approach could be a range of methods, including machine learning techniques, that essentially imputes missing values from the observed ones (Bertsimas et al (2017)).

Data interpolation algorithm | Figure 1

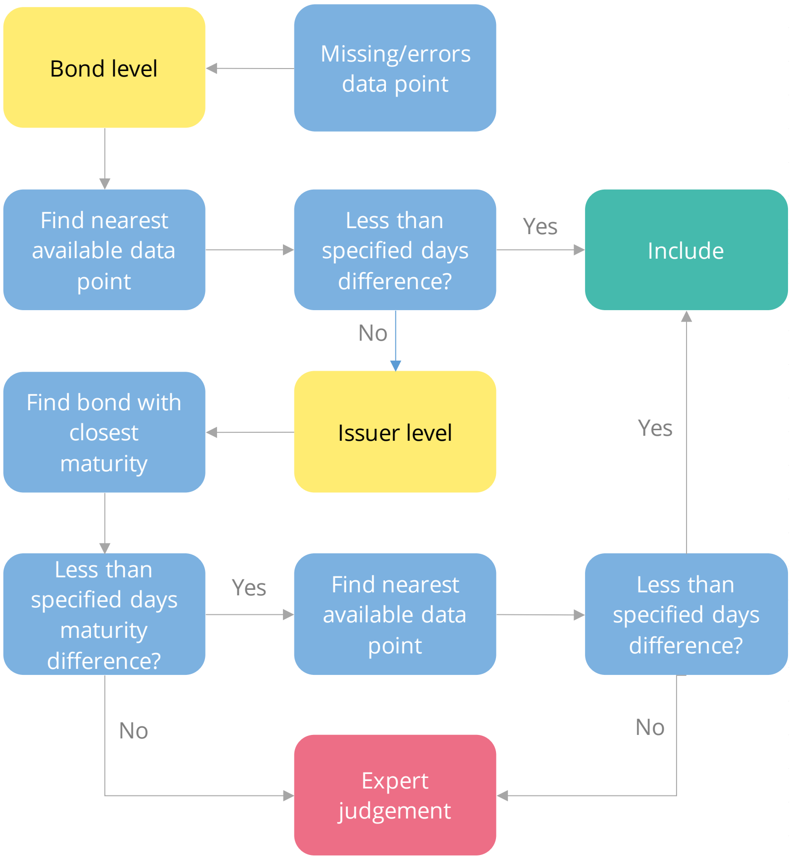

Source: Author's elaboration.

4. Calibration of a bespoke liquidity scoring tool

This section outlines the calibration of a potential liquidity scoring approach for monitoring and assessing the liquidity risk of fixed income investment exposures. The methodology ranks individual fixed income securities based on a number of liquidity indicators and aggregates these into a final liquidity score. The calculated scores allow comparison of exposures, not only at the individual security level, but also at issuer and portfolio levels. The tool provides a point-in-time estimate of liquidity, together with the ability to estimate liquidity for a historical time period, as well as the potential for more forward-looking analysis. This contributes to assessing trend and event analysis, with potential risk management applications that are considered in subsequent sections.

4.1 Methodology

To ensure that the approach was transparent and easily explainable to stakeholders, the methodology to estimate the liquidity score includes liquidity metrics that are widely accepted as affecting bonds' liquidity. Metrics such as the bid-ask spread of the bond, maturity, trading volumes and market depth (as defined in Section 3) were aggregated using a weighted scoring methodology, which is outlined in the subsequent sections. Through a combination of all these metrics, the calibrated liquidity score provides a robust liquidity measure for individual securities, from a public investors' perspective, as it incorporates the different aspects of liquidity that can be experienced in the financial markets. These four metrics can be grouped into the two key aspects of liquidity risk, which we have labelled "cost to sell" and "ability to sell".

Cost to sell identifies how much it will cost to sell any given security based on the observable data available. The approach captures this liquidity aspect by incorporating the bid-ask yield spread of the bond, as well as its residual maturity. With regard to residual maturity, the liquidity of a bond increases as the bond nears its redemption date, as the exposure becomes closer to cash the nearer it gets to maturity. This aspect of the score best maps to the common description of the "tightness" of market liquidity (CGFS (1999)).

Ability to sell identifies whether there is a ready market available for this bond, should a bondholder wish to sell it, and to what extent a bondholder might "move the market" if trying to enter the market as a seller (something that public investors generally try to avoid). This indicator comprises two elements, namely; trading volumes and market depth. Firstly, a measure of the market depth, which combines available dealer quotes and the number of pricing sources, aims to identify whether there is a buyer for the bond. Secondly, a measure of trading volumes has also been incorporated, to try to capture how much of a given bond could be sold over a particular period of time, without significantly impacting the bond's price. Figure 2 displays the summary of the composition of the presented liquidity score. This aspect of the score maps to the common description of the "depth" of market liquidity (CGFS (1999)).

Liquidity score composition Figure 2

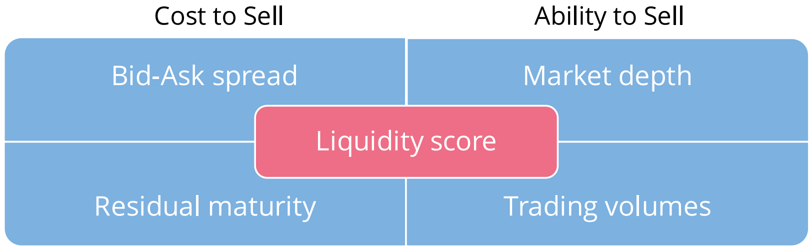

Source: Author's elaboration.

The final liquidity score for an individual bond is calculated by using the following simple approach. A security receives a score for each of the four key liquidity metrics. The scores for the four key metrics are determined by a number of predefined thresholds. For example, the bid-ask spread universe is divided into a number of thresholds or bands, where the band with the lowest bid-ask spread receives the highest liquidity score. The scores decrease in line with an increase in spreads. A similar methodology is also employed for the remaining liquidity components[11].

Next, the scores are aggregated to provide cost-to-sell and ability-to-sell scores, which are then combined to produce a final liquidity score for the instrument. The final score ranges from one to 100, with 100 being the most liquid. Once the unique securities are scored with an individual liquidity score, the final portfolio or issuer score is aggregated by using a simple weighted average approach, which is based on the relevant exposure size.

Furthermore, as well as calculating a score for all marketable fixed income exposures, such as bonds and treasury bills, other types of exposures may also be considered, such as repurchase agreements[12] and uncollateralised money market deposits. However, these exposures cannot be liquidated in the same way as one could do with a bond; hence a modified approach would be necessary. Given that the potential methodology for marketable and non-marketable exposures would not be homogenous, and would be based on different assumptions and factors, there would likely be difficulties in justifying the aggregation of these types of exposures in terms of an overall score. It is important that such differences are well calibrated, as well as being well communicated to senior stakeholders. A number of examples of how to incorporate these instruments into a liquidity score are outlined below, although it is acknowledged that they do not necessarily provide a perfect solution to the challenges raised by attempting to calibrate a portfolio level score.

As a starting point for looking at transactions such as these, an assumption could be made that all instruments with a residual maturity of less than one month are treated as equivalent to cash. Given that we assume such securities to be perfectly liquid in this scoring approach, an automatic score of 100 would be applied once any security reaches a residual maturity of less than 30 days. It is important to note that

[11] See Section 4.2 for more detailed information on the calibrations of the thresholds.

[12] It is true to say that an offsetting repo transaction could be conducted to offset the liquidity impact. However, it could also be argued that a repo transaction could be used to meet any liquidity need, which may not be possible in all circumstances, and may only have a limited term in any event.

this would be based on likely liquidity requirements, which are assumed here to mainly revolve around needing to sell securities to rebalance a portfolio or to transfer from one portfolio to another. Of course, this assumption is heavily dependent on the specific mandate of a central bank or public investor. For example, a central bank that has an explicit mandate to be ready to regularly intervene in the foreign currency markets may need to have a more restrictive view of what "cash-like" means to them and their requirements. Such an active currency intervention strategy may require a residual maturity to fall within a week in order to be considered cash equivalent, for example.

Having established this assumption, a simple scoring approach could be applied to all non-fixed income exposures, on a sliding scale from 100 at one-month maturity, down to 50, that would be aligned with the institution's maximum allowed maturity for such collateralised and uncollateralised exposures. The scoring methodology for these instruments can be informed by an analysis of the likely overall size and compositions of the exposures, primarily focusing on maturities and position concentrations. This can also be supplemented by an overlay of expert judgement. Another alternative methodology could include the liquidity score estimated based on the collateral that is being utilised in the transaction, usually a sovereign bond. Once a liquidity score on the collateral pledged is estimated, a discount factor could be applied to the score, to decrease the overall score due to non-marketability.

4.2 Calibration approach to score components

To meet the previously outlined granularity requirements, the model inputs and settings have to be appropriately calibrated. For example, it is important to ensure that the bid-ask spread and trading volume thresholds are granular enough to capture the daily liquidity changes of a particular bond. Separately, indicators such as market depth and bond maturity are more stable – hence the calibration for these variables can be less granular.

The calibration of the bid-ask thresholds should be based on the analysis of the bid-ask spreads of the public investor's investable universe. This analysis indicated that more than 90% of all the observable spreads of bonds within a traditional central bank's investable universe would fall within a very tight range of between zero and 10 basis points (see Graph 1). This meant that, if wider spread bands or thresholds were calibrated, the model would be less able to capture the changes in the bond's daily liquidity. Therefore, the calibration of bid-ask scoring bands were calibrated in a way that allowed analysts to better rank the bid-ask spreads within a public investor's investment universe. Interpolation was used within the bands to avoid discontinuous jumps in the scores, which might be based on small changes in variables.

Bid-ask spread calibration Graph 1

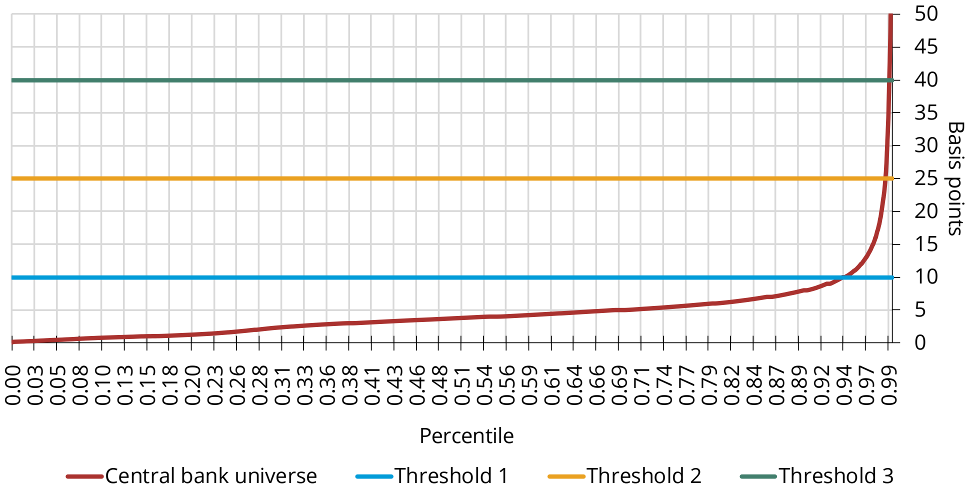

Source: Markit; Authors' calculations.

Graph 2 demonstrates how wider bid-ask bands and higher maximum allowable bid-ask spreads (once the maximum spread is reached, the security receives a score of zero for this aspect of its liquidity) would not be as responsive to daily liquidity changes of the bond. These types of wider calibrations would be more aligned to some market liquidity risk providers, which are typically calibrated to a wider bond market universe than a traditional central bank reserve manager, as discussed earlier. For this example, we use data for a specific corporate bond that experienced significant deterioration in risk profile and caused the bond's bid-ask spread to increase substantially over the period displayed. As can be seen by the graph, the significant changes in the bond's bid-ask spread, and the subsequent deterioration in liquidity, would not have been captured promptly by a liquidity metric that employs relatively wide bid-ask bands. For instance, a score based on the 50 bp wide band, depicted in the graph, does not indicate a significant drop in the bond's liquidity until November 2017. On the other hand, a more granular methodology, such as that described in this paper, displays greater sensitivity to developments in the bid-ask spread. The methodology presented captures substantial deterioration from July to November 2017, until the bond's bid-ask spread level increases beyond the final threshold range and is assigned the lowest score of zero, making such measurement more adequate for a public investor's investable universe.

Liquidity deterioration of a bond – bid-ask spread threshold comparison

Graph 2

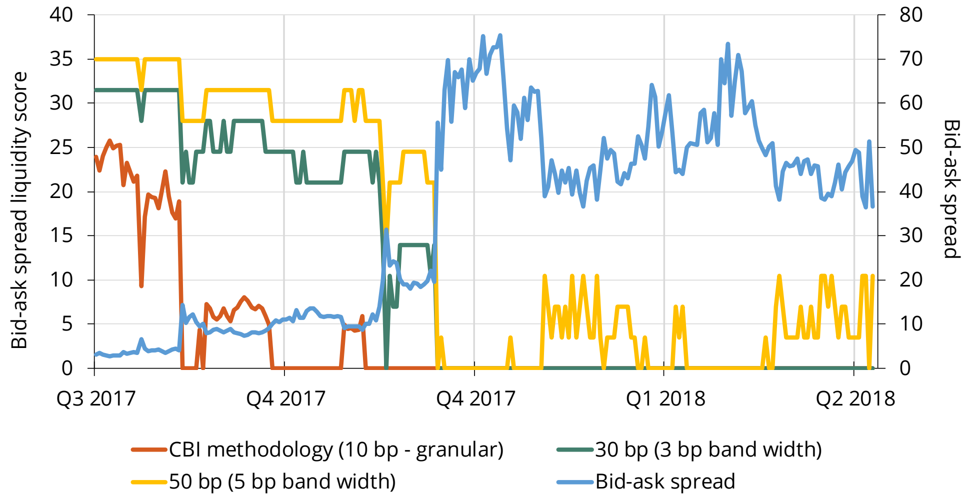

Source: Markit; Authors' calculations.

The calibration of the other variable thresholds can be based on a more qualitative analysis, where expert judgement was applied. For example, the calibration of maturity scoring bands can be linked with the public investor's target investment horizons and duration targets. The calibration of the trading volume thresholds can be estimated by reference to the standard composition of a public investor's exposures and nominal amounts of holdings.

5. Potential applications of the liquidity scoring approach

Incorporating liquidity risk monitoring into a holistic risk monitoring framework with other risks such as market risk, credit risk and interest rate mismatch risk, may allow for the identification of previously unidentified risks and/or a more complete view of the overall interaction of the investment risks faced by the investor. Armed with these insights, risk managers would be in a better position to provide more granular information to portfolio managers, and the institutional governance bodies, in support of risk- informed decision-making.

However, there are a number of additional, potentially beneficial, approaches that a central bank, or other public institution, might take to utilise liquidity tools – such as the one presented in this paper – to better incorporate liquidity risk into their overall risk management framework. A number of such approaches are proposed in the subsequent sections, beginning with the simplest applications, such as market analysis, followed by more complex uses, such as portfolio liquidity benchmarking and optimisation approaches that can be built upon the scoring approach.

5.1 Market liquidity analysis

To complement risk monitoring and reporting of its own investment exposures, risk managers must also perform and communicate risk intelligence to the institution's senior management. Liquidity tools, such as the one presented in this paper, can be valuable assets that allow for monitoring of financial market liquidity more broadly. In particular, the tool could aid in analysing market liquidity trends or conducting event analyses.

Euro area benchmark liquidity score (2012–2018) Graph 3

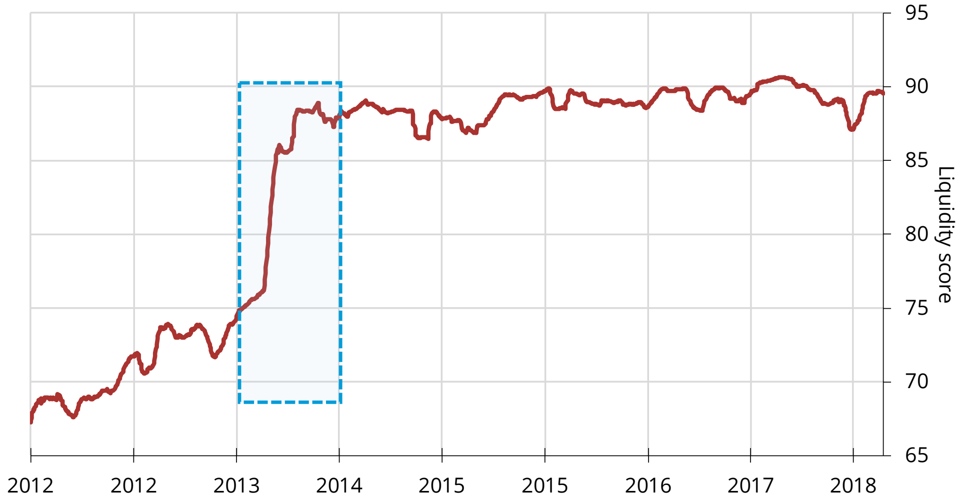

Source: Markit; Authors' calculations.

For example, using the liquidity scoring model, a simple analysis was performed to estimate whether liquidity in the euro area has changed in a meaningful manner during the period 2012–2018, which incorporates the introduction of the various ECB asset purchase programmes. An equally weighted euro area benchmark was developed, consisting of six countries (France, Germany, Ireland, Italy, Portugal and Spain) and four tenors (two-year, five-year, 10-year and 30-year). Graph 3 provides the summary snapshot of the analysis.

As can be seen, the score indicates that market liquidity increased significantly during the period 2012–2014, while the period 2014–2018 exhibited no major changes in market liquidity. The key driver in the increased liquidity in the presented model was the decrease in euro area bid-ask spreads. While this analysis is consistent with a number of recent publications (Jurkšas et al (2018), Larkin et al (2019)), there is some mixed evidence on the topic (ESMA (2018)). There are a number of reasons offered for the increase in liquidity in the European bond market, including the introduction of quantitative easing by the ECB and the overall increase in the credit quality of euro area sovereign issuers.

In addition to the longer-term trend analysis, the tool can also be utilised when evaluating liquidity changes in specific issuers surrounding a risk event. Graph 4 presents the evolution of the French 10-year bond around the time of the 2017 election in France, with the newly developed tool used to produce a liquidity score of between one and 100 in the manner described previously. As can be seen from the increased uncertainty surrounding the French election outcome in late 2016, the liquidity of the bond began to decline as expectations for a Marine Le Pen victory grew[13]. It is important to note that the bond was still considered very liquid according to the model; however, there were some notable changes. Once the outcome of the election became clearer, the bond returned to its previous liquidity levels. Analysis of historical events such as these provides useful information on how market liquidity may react to similar events in the future. For example, if the uncertainty persisted for a longer period of time, or if the outcome of the elections was not perceived to be market friendly, the liquidity of the bond may have decreased further.

Liquidity score of French 10-year bond around the 2017 election in France — Graph 4

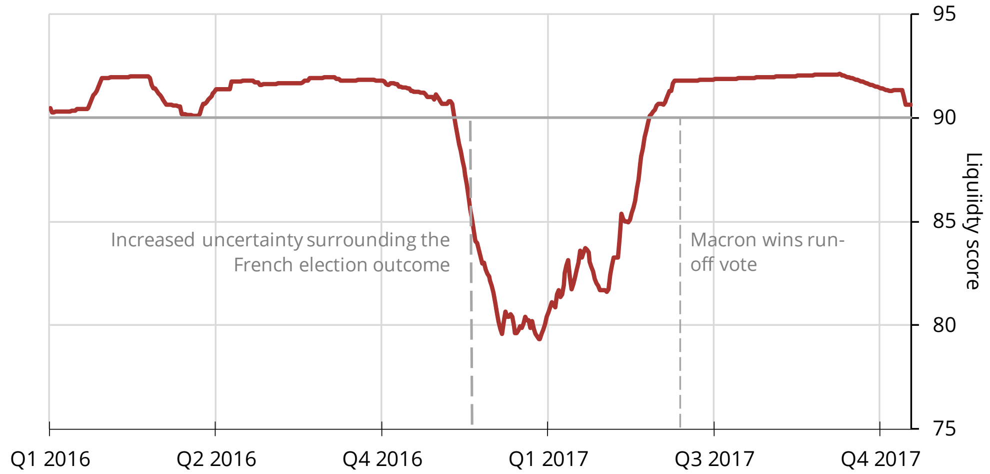

Source: Markit; Authors' calculations.

5.2 Liquidity benchmarking

Once the monitoring and analytical skills and expertise are developed from consistent use of a scoring tool, a more formal introduction of liquidity risk management in the investment management process, in the form of liquidity benchmarking, can be contemplated.

As with performance evaluation or market risk measurement, absolute liquidity risk figures provide useful information, but may not tell the full story. Therefore, liquidity must also be evaluated on a relative basis; for example, the liquidity risk of an issuer might be compared with its peers, a fixed income class such as corporate bonds might be compared with other classes such as sovereign fixed income, or active portfolio management decisions might be compared with a more neutral portfolio liquidity risk score. Therefore, there is a strong argument to use liquidity benchmarking when evaluating the relative liquidity of a portfolio manager's holdings or, indeed, the overall asset allocation of a public investor.

A role for portfolio liquidity benchmarking is supported by two observations. Firstly, a central bank or any other public institution must invest in something, as public funds left idle would forego potential returns (notwithstanding negative yields

[13] For an example of the sentiment at that time, see Financial Times (2017).

on some investment options in the current low interest rate environment). Secondly, the overall liquidity of the market inevitably changes over time caused, for instance, by various structural and regulatory factors, which are outside an investor's control. In this way, at a portfolio level, an institution's risk managers should not be "punishing" portfolio managers for systemic trends that are not within their control. Hence, similar to a performance benchmark used for portfolio manager's investment performance evaluation, a liquidity benchmark might be used to evaluate the investment portfolio's liquidity relative to a benchmark. This would allow the risk manager to determine whether changes in portfolio liquidity is being driven by the portfolio manager's actions or by the overall change in liquidity of the market or market sector to which the portfolio has to invest, according to their mandate.

There are a number of options that can be considered when choosing the appropriate liquidity benchmark. Ideally, the chosen benchmark should be aligned to the investment manager's return and risk objectives. Hence, the benchmark that is currently employed for the purpose of relative performance and risk measurement might be the best candidate. However, if risk managers want to evaluate liquidity risk of the overall market rather than the specified sector, a broader market benchmark might be more preferable.

Following the construction of the appropriate liquidity benchmark, the starting point might be the estimation of liquidity of both the individual portfolios and the specific benchmarks. An example of a liquidity benchmark and portfolio liquidity score are illustrated in Graph 5, where the decrease/increase in the liquidity score indicates an increase/decrease in associated liquidity risk.

Tracking portfolio liquidity versus a benchmark liquidity — Graph 5

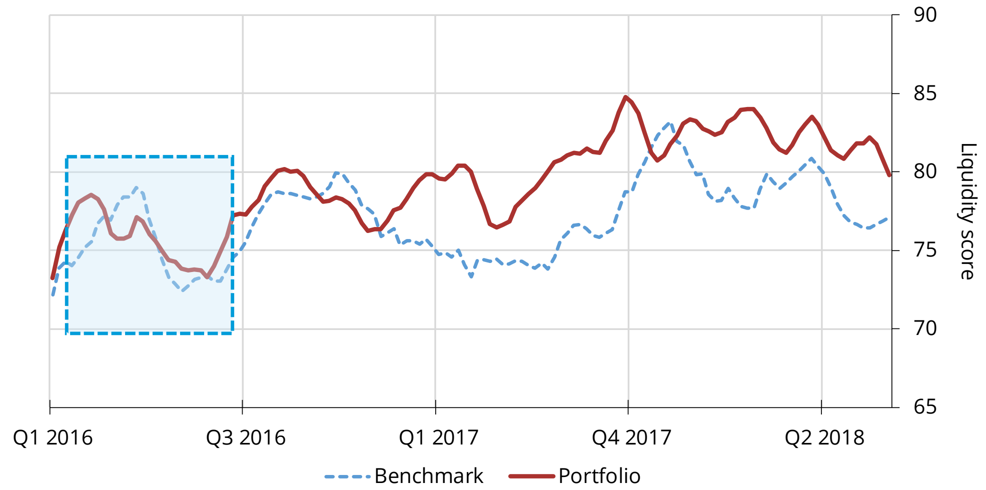

Source: Markit; Authors' calculations.

As can be seen from this example, the portfolio manager is tracking the benchmark relatively well; in fact, during significant periods of time, the portfolio manager's liquidity score is higher than the assigned benchmark. This information could be interpreted in two ways; positively, from a risk management perspective, as the portfolio manager is taking on less risk than the benchmark, or it could also have potentially negative aspects from an investment perspective, as the portfolio manager may not be fully utilising and capturing potential liquidity premia allocated to them by the benchmark.

Graph 5 also highlights a number of other observations relating to market developments. For example, from February to July 2016 (in the run up to the UK Brexit vote), portfolio liquidity risk increased noticeably (as the overall portfolio liquidity score decreased), which might raise some questions for a risk manager if portfolio liquidity risk was estimated in isolation. However, analysing this information in tandem with the benchmark liquidity metrics indicates that benchmark liquidity risk also increased (as the overall benchmark liquidity score decreased). Given that the benchmark tracks the portfolio's investable universe, a conclusion can be made that the overall liquidity of the invested portfolio decreased due to more systemic factors such as political uncertainties, rather than portfolio managers' actions.

The natural follow up question is how to calibrate a means of limiting the extent that a portfolio manager (or the overall investment assets) is permitted to deviate from the relevant benchmark? One way of answering this might be to allocate a deviation budget to the portfolio manager, whereby deviation beyond this budget would be limited. This would be similar to the type of active risk budget that is allocated to investment managers in respect of market and/or credit risk.

Benchmark deviation thresholds Graph 6

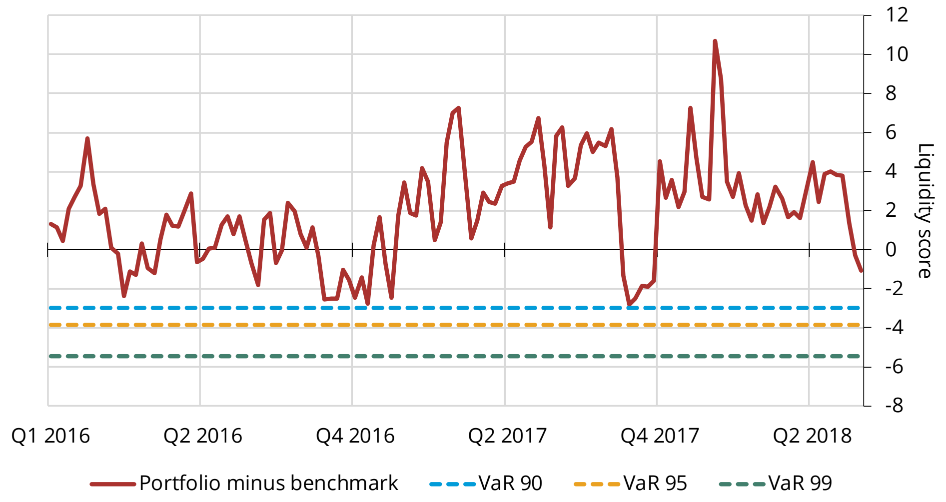

Source: Markit; Authors' calculations.

This can be achieved by examining the volatility of the liquidity score of the benchmark itself, and the extent to which the benchmark liquidity score deviates from its mean value. This could be used as a proxy for the acceptable deviations of the portfolio liquidity score from the benchmark liquidity score. Portfolio liquidity deviations are estimated by subtracting the overall portfolio liquidity score from the overall benchmark liquidity score. Subsequently, these estimates might be employed to set limits against which the portfolio manager might be evaluated. For example, the limits could be estimated by choosing an appropriate confidence level (eg VaR 90), on either a historical or parametric basis. Graph 6 illustrates potential threshold levels. The choice of percentile can be guided by, amongst other things, the liquidity risk appetite of the investor.

As can be seen from the example in Graph 6, the portfolio manager would not have breached any of the assigned limits during the past number of years. In order to operationalise this methodology, there are additional considerations that may need to be incorporated. It may be necessary to include some sense of the persistency of the limit breaches. For example, a portfolio manager might be informed about the breaches and corrective action might be requested only if the limits breach is persistent rather than a once-off occurrence. Again, these types of additional methodological approaches would have to be calibrated based on the institution's risk appetite[14].

5.3 Liquidity optimisation

There are additional approaches that can be borrowed from other risk domains when designing risk management applications utilising the liquidity score. One such approach is to incorporate liquidity risk scoring into an optimisation exercise. This could be used in the context of performing risk assessments of new investment proposals. It should be noted, however, that such an exercise could be incorporated as part of a holistic risk assessment that also considers other types of risks (eg credit and market risks). It is, therefore, not proposed that asset allocation decisions are calibrated with respect to an optimisation exercise in terms of liquidity risk only.

In order to illustrate this approach, it is useful to look at a practical example. Consider a portfolio that consists of four issuers, which are equally weighted in terms of holdings (this example could easily be adapted to that of asset classes or currencies etc.). This portfolio has an associated liquidity score that can be tracked and monitored over time. Consider that an investment proposal is produced that wants to add a fifth issuer, with a subsequent reduction in the holdings of the first four issuers so that all issuers are once again equally weighted. A hypothetical new portfolio can then be created, which includes the benchmark holdings of the proposed new issuer together with the holdings already in the portfolio, and which measures its hypothetical historical liquidity score over the same period as the old portfolio. These two portfolios can be seen in Graph 7.

[14] The benchmarking approach can also be applied at a global level, when making asset allocation decisions such as entering new asset classes, geographical areas or currencies. In this case, it is necessary to construct a global liquidity benchmark, which would include a combination of financial instruments from the most liquid currencies. For example, such a benchmark might include euro area and US dollar sovereign and sovereign-like securities, as well as other currencies, if appropriate.

Current portfolio liquidity versus proposed portfolio liquidity	Graph 7

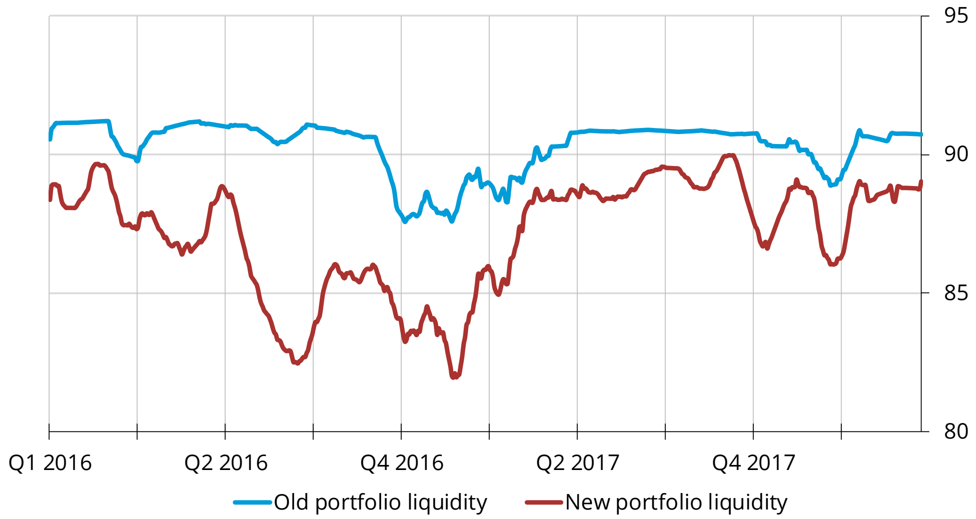

Source: Markit; Authors' calculations.

The graph shows that, on a historical basis, the effect of the addition of the new issuer would be to both reduce the overall level of the liquidity score, as well as increase its volatility over time. Similar exercises can be performed on a more forward-looking basis, such as in stress scenario analysis, if required. The question arises as to whether a risk manager would be happy with the outcome of changing the liquidity profile in this way. If not, it is possible to perform an optimisation of the liquidity score when considering an appropriate liquidity risk exposure.

There are a number of different bases on which to perform such an optimisation. For example, it is possible to construct an asset allocation that exhibits a generally higher liquidity score (either maximising or aiming for a particular value) or a lower score volatility (either minimising or aiming for a particular value). Another option is to limit the extent of deviation from a constructed global liquidity benchmark, similar to that described in Section 5.2. Examples of the output of such an optimised approach are illustrated in Graph 8.

As mentioned above, such an exercise could be performed as part of a comprehensive analysis of any potential new investments. Additional constraints could also be added to specific issuers that warrant an overweighting based on a consideration of the risk/return trade-off and, as such, many different liquidity profiles can be constructed and assessed in accordance with the public investor's specific risk appetite.

Sample optimised portfolio Graph 8

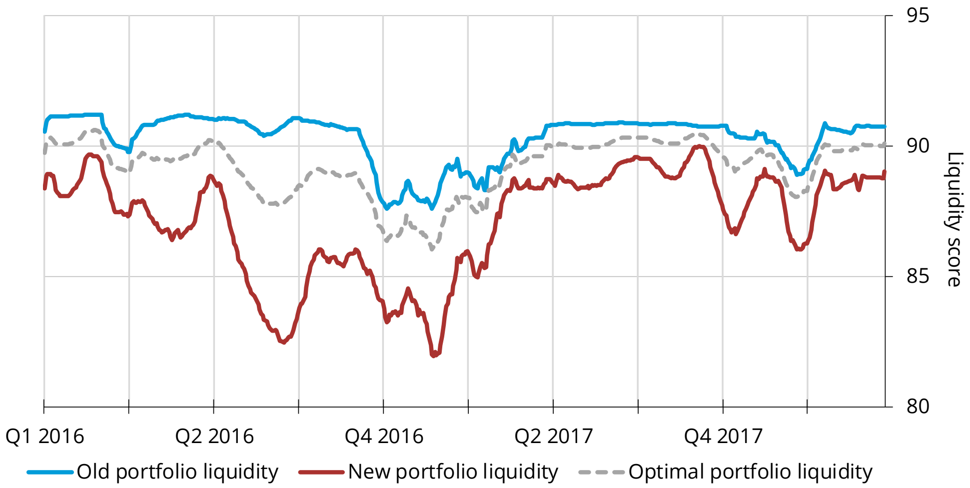

Source: Markit; Authors' calculations.

Similar to the use of mean-variance optimisations in asset allocation decisions, the assumption that the past history of an asset's liquidity is a good guide to the future is a significant assumption to make, and comes with some risks. Therefore, it may be useful to consider other means of calibrating such a liquidity optimisation approach, such as attempting to add a forward-looking element. This could be achieved by estimating the forward-looking parameters, with the addition of a simulation approach to the components of the liquidity score. This can also be complemented through the use of scenario analysis, where liquidity conditions of relevant asset classes in certain reference time periods could be utilised to derive potential outcomes for new asset classes. Finally, the utilisation of stressed liquidity factor correlations, either hypothetical or based on historical data, could also be incorporated in the optimisation process.

6. Conclusions

Since the GFC, there has been an increasing focus on liquidity risk as a risk factor in its own right. It is in the context of this increased recognition of liquidity risk that investors, including public investors such as central banks, have become more interested in measuring and monitoring the liquidity risk of discretionary investment asset portfolios.

The development process of the liquidity risk management framework, for a public investor, might identify that the nature of public investors' holdings necessitates a tailored liquidity scoring methodology, calibrated to an institution's own investment universe and risk appetite. Some characteristics that might be considered when developing such frameworks are transparency, flexibility and granularity.

This paper presents an approach to measuring and monitoring liquidity risk that can be tailored specifically to a public investors' investment assets. The methodology presented in the paper was used to estimate whether liquidity in the euro area has changed in a meaningful manner during the period 2012–2018. Analysis using this approach shows that liquidity in the European sovereign bond market increased significantly between 2012 and 2014, while the period 2014–2018 exhibited no major changes in market liquidity.

Moreover, the paper considers a number of approaches that a central bank, or other public institution, might take to utilise liquidity risk scoring tools – such as the one presented in this paper – to better incorporate liquidity risk into their overall risk management framework. These approaches to managing liquidity risk, including concepts such as liquidity benchmarking and liquidity optimisation, can offer a number of potential benefits, such as a better evaluation of the portfolio liquidity position relative to the market, and a comprehensive assessment of the impact on the liquidity profile from potential changes in portfolio composition.

Finally, some next steps to develop practical applications of the methodology could range from enhancements of data pre-processing techniques to refining the applicable threshold and calibration approaches. The latter could include further development of the liquidity optimisation approach or the development of the look-through approach to assets linked to repo transactions. For public investors, the use of liquidity risk measurement and monitoring approaches is likely to become a standard feature of their risk management toolbox. By developing a better understanding and management approach for liquidity risk, alongside market and credit risk, public investors can enhance investment decision-making and ultimately limit potential losses. This paper has sought to contribute to this ongoing development path by presenting a liquidity risk scoring approach along with potential applications.

References

Amihud, Y and H Mendelson (1988): "Liquidity and asset prices: financial management implications", *Financial Management*, pp 5–15.

Basel Committee on Banking Supervision (2010): *Basel III – International Framework for Liquidity Risk Measurement, Standards and Monitoring*.

Bertsimas, D, C Pawlowski and Y Zhuo (2017): "From predictive methods to missing data imputation: an optimization approach", *Journal of Machine Learning Research*, no 18, pp 1–39.

Committee on the Global Financial System (1999): "Market liquidity: research findings and selected policy implications", *CGFS Papers,* no 11, May.

——— (2016): "Fixed income market liquidity", *CGFS Papers*, no 55, January.

Danielsson, J, M Valenzuela and I Zer (2018): "Learning from history: volatility and financial crises", *The Review of Financial Studies*, vol 31, no 7, pp 2774–805.

European Securities and Markets Authority (2018): *Liquidity in EU Fixed Income Markets – Risk Indicators and EU Evidence*, September.

European Systemic Risk Board (2016): *Market Liquidity and Market-making*, European System of Financial Supervision, October.

Financial Times (2017): "Could France's Marine Le Pen deliver Frexit?" www.ft.com/content/d37b6d90-fdd1-11e6-8d8e-a5e3738f9ae4.

Favero, C, M Pagano and E Von Thadden (2010): "How does liquidity affect government bond yields?" *Journal of financial and quantitative analysis*, vol 45, no 1, pp 107–34.

Galliani, C, G Petrella and A Resti (2014): "The liquidity of European corporate and government bonds: drivers and sensitivity to different market conditions", *JRC Technical Reports*, European Commission.

Jurkšas, L, D Kapp, K Nyholm and J Von Landesberger (2018): "Euro area sovereign bond market liquidity since the start of the PSPP", *ECB Economic Bulletin Boxes*, no 2.

Kofman, P and I Sharpe (2003): "Using multiple imputation in the analysis of incomplete observations in finance", Journal of Financial Econometrics, vol 1, no 2, pp 216–49.

Laganá, M, M Perina, I von Köppen-Mertes and A Persaud (2006): "Implications for liquidity from innovation and transparency in the European corporate bond market", *ECB Occasional Paper*, no 50.

Larkin, J, P Anderson and S Furlong (2019): "The Irish government bond market and quantitative easing", *Central Bank of Ireland Quarterly Bulletin*, no 2, April.

Pasquariello, P and C Vega (2009): "The on-the-run liquidity phenomenon", *Journal of Financial Economics*, vol 92, no 1, pp 1–24.

Yuan, Y C (2010): *Multiple imputation for missing data: concepts and new development (Version 9.0)*, SAS Institute Inc, Rockville, MD.