AUTORITÉ DES MARCHÉS FINANCIERS

AMF

Scientific Advisory Board Review

Contributors to this issue:

- Eric Chaney, AXA
- Serge Darolles, Paris-Dauphine University

No. 3 – November 2015

AMF Scientific Advisory Board Review

No. 3 - November 2015

Publication director: Guillaume Eliet

Editor in chief: Stéphane Gallon

Editorial team:
Eric Chaney
Serge Darolles

Secretariat : Muriel Visage

ISSN : 2427-4461

Legal deposit : November 2015

The AMF Scientific Advisory Board

The AMF set up a Scientific Advisory Board composed of prominent figures from the academic and financial worlds to expand its research and strategic intelligence capabilities.

The members of the Scientific Advisory Board hold positions in leading universities, business schools and public and private research centres. The board includes foreign researchers of international reputation, which gives it a global perspective. It has a triple remit:

1. provide the AMF with information on ongoing academic research in the financial field,
2. identify developments that may have an impact on the AMF's areas of activity,
3. undertake research projects related to issues of concern to the regulator.

The Chairman of the Scientific Advisory Board is Gérard Rameix, AMF Chairman, and the Vice-Chairman is Christian de Boissieu, a member of the AMF Board. The other members, appointed by a Board decision, are:

Luc Arrondel, National Centre for Scientific Research; Paris School of Economics
Patrick Artus, Natixis
Bruno Biais, Toulouse School of Economics
Marie-Hélène Broihanne, Strasbourg University
Gunther Capelle-Blancard, Paris Panthéon Sorbonne University
Eric Chaney, AXA
Serge Darolles, Paris Dauphine University
Thierry Foucault, HEC
Olivier Garnier, Société Générale
Carole Gresse, Paris Dauphine University
Philippe Ithurbide, Amundi
Charles-Albert Lehalle, Capital Fund Management
Albert Menkveld, Vrije Universiteit Amsterdam
Patrice Poncet, ESSEC
Nicolas Véron, Peterson Institute for International Economics; Bruegel

To help disseminate the Scientific Advisory Board's research, the AMF organises an annual conference that brings together the financial and academic communities to examine a specific topic.

CONTENTS

Insurance companies in a low-yield world
Eric Chaney, AXA ..... p. 5

Insurance companies in a low-yield world
Bond yields are very low by historical standards and may stay so for several years, if the global economy is in a 'secular stagnation' period and if duration premiums remain depressed. Within the business portfolio of insurance companies, life is the most exposed to the low-yield risk especially when policyholders benefit from guaranteed returns and when the asset-liability duration gap of the insurer's balance sheet is large. This is the case for a significant number of life insurance companies in Europe. To mitigate the low-yield risk, insurers may diversify their business portfolio in favour of less yield-sensitive activities such as health insurance, reduce guarantees, make their ALM more resilient by reducing duration gaps and actively manage their assets.

Equity market liquidity and long-term investing
Serge Darolles, Paris-Dauphine University ..... p. 16

Equity Market Liquidity and Long-Term Investing
This article addresses a range of issues related to the analysis of liquidity in equity markets from the perspective of long-term investors. The results of a survey of French institutional trading desks are the basis for the presentation of a model capturing the effects of information and liquidity on volumes and changes in daily prices. Using this approach allows us to create four major asset classes, depending on their exposure to different types of liquidity frictions. An empirical application shows that even in the supposedly liquid markets, many securities are subject to a deficient short-term liquidity provision, and/or strategic behaviour of long-term investors.

Insurance companies in a low-yield world

Eric Chaney[1], AXA

Abstract

Bond yields are very low by historical standards and may stay so for several years, if the global economy is in a 'secular stagnation' period and if duration premiums remain depressed. Within the business portfolio of insurance companies, life is the most exposed to the low-yield risk especially when policyholders benefit from guaranteed returns and when the asset-liability duration gap of the insurer's balance sheet is large. This is the case for a significant number of life insurance companies in Europe. To mitigate the low-yield risk, insurers may diversify their business portfolio in favour of less yield-sensitive activities such as health insurance, reduce guarantees, make their ALM more resilient by reducing duration gaps and actively manage their assets.

Résumé

Les rendements obligataires sont historiquement bas et pourraient le rester, si l'économie mondiale est entrée dans une période de stagnation et si les primes de duration restent déprimées. Au sein des activités d'assurance, l'assurance vie est la plus touchée par le risque de taux bas, particulièrement lorsque les assurés bénéficient de taux garantis et que l'écart de duration entre actif et passif dans le bilan de l'assureur est élevé. C'est le cas d'un nombre important d'assureurs en Europe. Pour réduire le risque de taux bas, les assureurs peuvent diversifier leur portefeuille d'activité en faveur de secteurs moins sensibles aux taux, comme l'assurance santé, réduire les garanties, rendre leur gestion actif-passif plus résiliente en réduisant l'écart de duration et en gérant leurs actifs plus activement.

[1] Eric Chaney is Chief Economist for the AXA group and Head of Research for AXA IM.
Contributors to this article: Zhili Cao, Gilles Dauphiné and Mathilde Sauvé at AXA IM.

Long-term nominal interest rates, as measured by 10-year yields on government bonds, have trended down almost since 1984. They were virtually zero in Germany in April 2015 and even sank below zero in Switzerland and Denmark. Rates declined more quickly and at an earlier date in Japan, where the 10-year yield fell to 0.5% in the spring of 2003, whereas inflation had stabilised at zero after several years of deflation. This would have little significance in a fully flexible world, one in which prices and policies adjust in continuous time and central banks are not reluctant to reduce key rates below zero. But while the Swiss National Bank has shown that this can be done with no apparent harm, central bankers remain deeply wary of entering uncharted territory, and policies, in particular for life insurance, are not or only partially flexible, in comparison with market volatility. Assuming eurozone rates remain extremely low for several years, let's say below 1%, life insurance companies, like pension funds, would face serious transition problems owing to the asset/liability duration mismatch and the returns promised to their beneficiaries. This article seeks first to shed light on the factors that pushed rates down. It then considers the European life insurance industry's sensitivity to these factors before outlining several possible coping and resilience strategies.

Figure 1: A 30-year decline in long-term yields

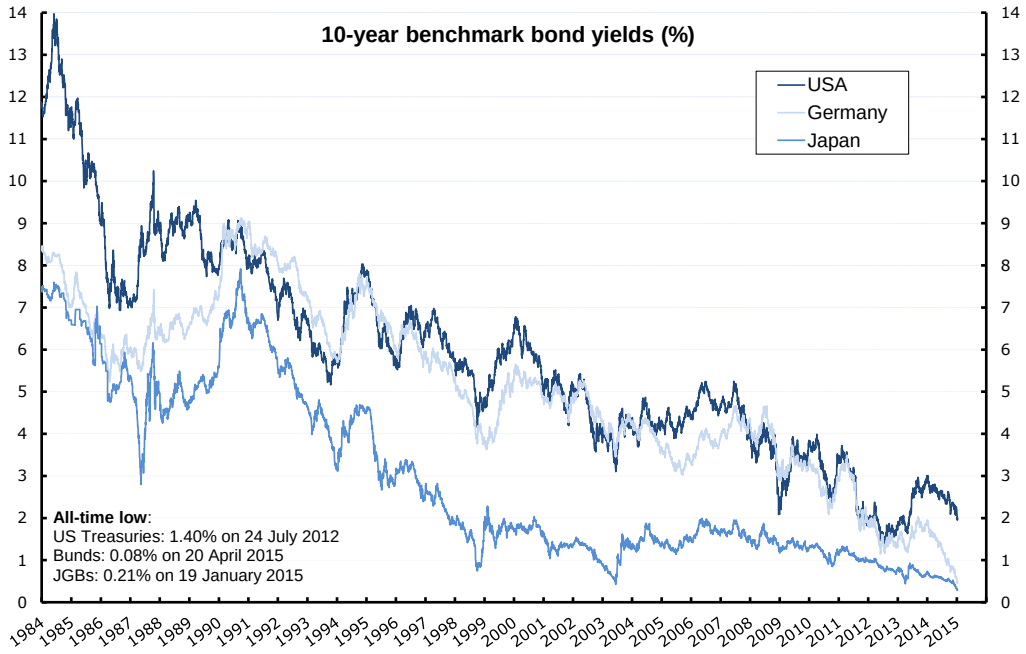

Sources: Reuters Datastream, AXA IM Research

The decline in long-term yields appears structural at first glance

An analysis of zero-coupon yield curves in the USA and Germany shows that the sharpest decline was in the longest-term forward interest rates, whereas they should, in theory, have been the least sensitive to cyclical fluctuations, whether in volumes, prices or central bank interest rate policies. The one-year forward interest rate nine years ahead fell by 230 basis points (bp) between end-December 2009 (start of the post-crisis recovery) and early September 2015 in the USA and by 265bp in Germany. These forward interest rates nine years ahead are meant to indicate what the market considers to be the equilibrium long-term rate. Yet, compared with their historical average in the previous cycle (1999-2007) of 6.35% in the USA and 5.40% in Germany, they are now 330bp lower in both regions. This coincidence suggests that common factors are at work worldwide, or at least in developed countries.

Figure 2: Long-term forwards: more than 3 points lower than in the previous cycle

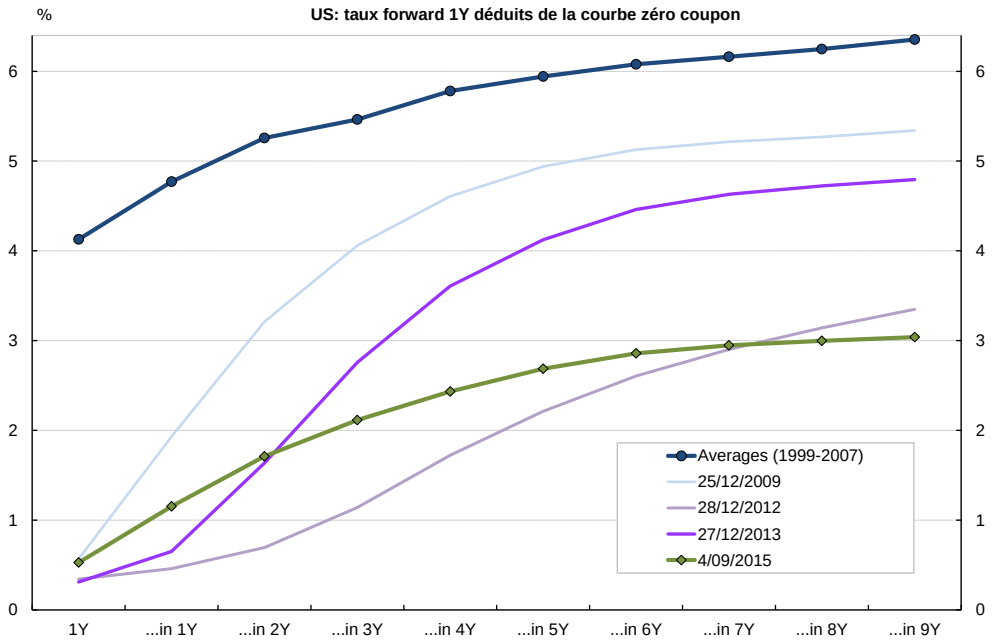

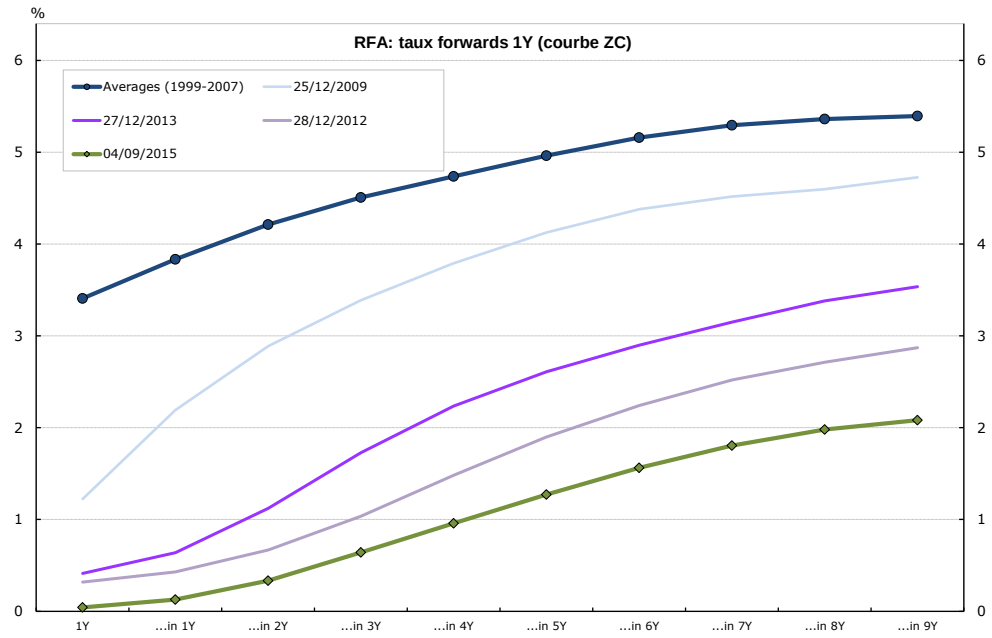

Sources: Reuters Datastream, AXA IM Research

Among the most convincing of the theories often cited to explain the structural decline in long-term rates are those of a global saving glut — as posited by Ben Bernanke as early as 2005 (1) — that has pushed the global equilibrium real interest rate down, and of a sharp slowdown in potential growth in developed countries, for both demographic and productivity slowdown reasons. These two possible explanations are not mutually exclusive and there has recently been some interesting theoretical and empirical support for the first, taking into account the disparity in access to credit in developed and emerging countries — see Coeurdacier 2015 (2). Concerning the second, some, such as Robert Gordon (3), believe the slowdown is not only

structural but also permanent. For others, such as Larry Summers (4), who opened the debate on secular stagnation, the roots of the slowdown lie in inadequate aggregate demand caused by real interest rates that, while undeniably low, are still too high because nominal interest rates cannot be negative. His remedy is a coordinated and expansionary fiscal policy at the global level. Still others, such as Joel Mokyr (5), believe the slowdown is overestimated (growth is underestimated) and temporary, given the extent of current and future technological innovation.

While it is difficult to draw a robust conclusion from the current academic debate, it would be prudent to consider the possibility that the decline in long-term nominal interest rates is permanent and, if it is temporary, that the normalisation period will be lengthy. However, a more detailed statistical analysis of yield curves shows that, while fundamentals should justify lower money-market rates in the future than in the previous cycle, long-term interest rates could nevertheless rise, if the distortions created by quantitative monetary policies were to subside. This can be illustrated by decomposing long-term interest rates into risk-adjusted rates and term premiums.

Half the decline in long-term interest rates is due to the fall in term premiums

For the purposes of the analysis, we decompose the long-term interest rate into a “risk-free” rate, the average of short-term interest rate expectations once the investor risk aversion effect is removed from observed forward interest rates, and the term — or duration — premium. For a life insurance company, the concept of duration premium is very real: assuming a policyholder agrees not to surrender his policy until a set date, the insurance company can cover the commitment with an asset having the same duration. If this asset has no credit risk (sovereign bond), the insurance company would then not have to concern itself with market fluctuations. But if, as is the case in reality, the policyholder can surrender his policy at any time, the uncertainty about the asset's ability to cover the liability gives the insurance company good reason to seek additional return to offset the risk incurred on the duration of its liability (the duration premium) and, secondarily, results in an asset/liability duration gap that is tilted towards the liabilities side.

In practical terms, we replicate the methodology developed by Adrian, Crump and Moench of the Federal Reserve Bank of New York (6), for US, UK, Japanese and German zero-coupon curves. The German curve stands in for the eurozone curve. The example of Germany is particularly relevant, given the very low level of long-term rates observed in that country and its yield curve's leading role in swap rates.

Figure 3: Record fall in the term premium of German bonds (%)

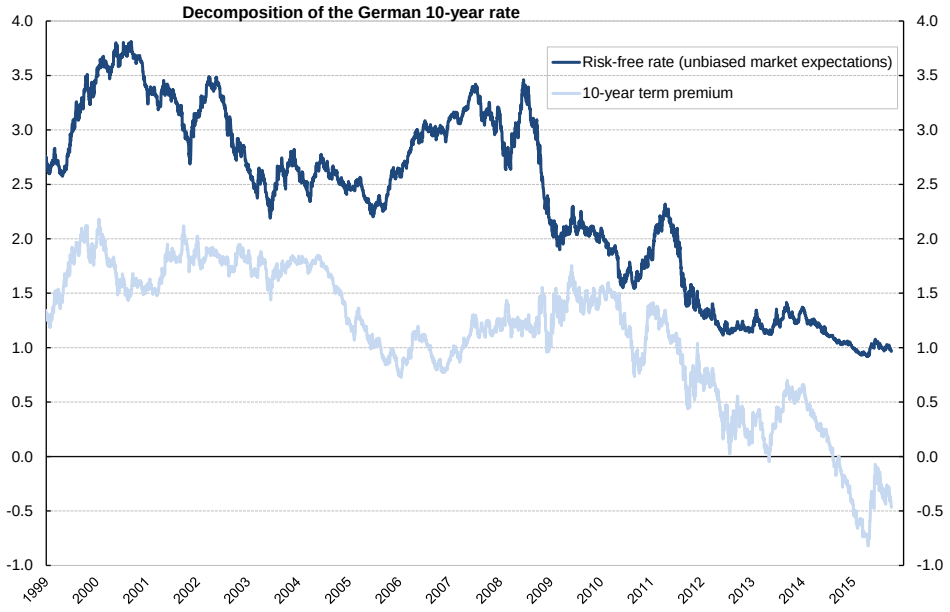

Source: Bloomberg, AXA IM Research

From the launch of the euro until end-2010, or roughly until the election of George Papandreou in Greece, which signalled the start of the eurozone's sovereign crisis and restored German government bonds to their safe-haven status, the 10-year term premium averaged 150bp, with relatively little variance. Our calculations indicate it is now negative. At the same time, the risk-free rate, consisting of ECB interest rate expectations for the next 10 years, has fallen fairly steadily since peaking in 2008 and has ultimately stabilised at around 1%. As inflation expectations have also fallen, albeit less so, expected real rates have decreased to a lesser extent. This is of interest because, in the eurozone's case, it shows that the markets are sympathetic to the theory of a protracted slowdown in the area's potential growth. But what is important for investors is that the fall in the term premium (from 1.4% on average before the euro crisis to -0.5% on average year to date) accounts for half of the fall in the 10-year yield which, over the same period, decreased from 4.2% to 0.5%. A move in term premiums back to their historical averages would lift eurozone 10-year rates by nearly 200bp, all else equal.

Does it then follow that once things are back to normal, 10-year rates will rise to at least around 3%? That would be a stretch, first because eurozone stability is far from assured, but mostly because the US example shows that central banks' policies of purchasing sovereign securities cause the term premium to fall to the same extent as the increase in the central bank's balance sheet that results from these purchases. The collapse in the German term premium is easily explained by the ECB's quantitative policy, which the markets expect to last at least until September 2016 and which will therefore increase the ECB's balance sheet by an amount comparable to the increase in the Fed's balance sheet. Empirical studies have shown that the Fed's asset programmes have depressed the term premium by about 110bp, according to Federal Reserve Vice Chairman Stanley Fischer (7). As it is likely that the purchased securities will remain on the ECB's books for a long time after the end of the purchasing policy, it is entirely possible that the term premium will remain below its long-term average for several years.

For insurance companies, the challenge of low yields is therefore most likely a lasting one.

Life insurance more affected than other products

Of the products offered by the insurance industry, life insurance is clearly the most affected. Whether it is closer to a retirement benefit, like in Germany, or a savings product, like in France, these policies come with guarantees (formal guarantees or those resulting from intense competition), the extent and rigidity of which determine rate sensitivity.

Savings managed in unit-linked form are less affected because, with these types of products, the policyholder bears most of the risk. Also, this form of savings is generally more exposed to equity investments, whose sensitivity to falling rates is not deterministic and depends mostly on the reasons for the decline in rates. We might expect, for example, that the fall in the term premium due to quantitative policies is mostly positive for equities whereas a protracted slowdown in potential growth is mostly negative.

As for non-life insurance (health, property and casualty, personal), rate sensitivity depends chiefly on the time horizon for the insurance product.

Table 1: Life insurance is five times the size of non-life insurance in the European Union

(€bn)	Non-life insurance	Life insurance
Liabilities (2011)	890	4,490
Premiums (2013)	447	670
Benefits and claims paid (2013)	628	324
Underwriting income	67%	20%
Interest income	33%	80%

Source: Insurance Europe – European Insurance in Figures, 2015

What are the main impacts of low yields?

We take an all-else-equal approach, meaning irrespective of the reasons for the decline in bond yields. While useful for the analysis, this partial equilibrium reasoning could be misleading if we were to attempt to use it prescriptively. Although the central banks' firmly expansionary policies are partially responsible for low yields, their aim is nevertheless to prevent a deflationary depression. In that scenario, the damage to the real economy would likely be far worse than the damage that might result from persistently low interest rates in a context of moderate growth, but growth nonetheless.

To better understand the impact of low yields on insurance companies' balance sheets, it is useful to start with a thought experiment, that of a fully flexible world. In such a world, prices and interest rates would adjust to macroeconomic fluctuations without delay or viscosity and the insurance company's entire balance sheet, both assets and liabilities, would also be flexible, as it would be indexed to rates and therefore to inflation. So, rate hikes or cuts would have the same impact on assets and liabilities, and would therefore be neutral.

The impact of low yields stems from the differences between actual policies and this ideally flexible world:

- On the liabilities side, **guarantees** are embedded in contracted policies to varying degrees, while assets are subject to market fluctuations (asymmetry).

- Because of (i) problems finding long-term assets and (ii) management of early surrender risk, in life insurance there is generally a **gap between the duration of the**

asset (A) and the duration of the liability (L), with A < L. The greater the L-A duration gap, the more sensitive the balance sheet is to the risk of falling rates, because the return on the asset will fall more quickly than the return expected by policyholders, on the liabilities side.

Life insurance and pension activities are therefore more affected by low-rate risk than property and casualty (P&C) insurance, where prices are generally flexible and which involves shorter-term commitments, resulting in a negative duration gap (L<A). At the other end of the scale, unit-linked (UL) savings and health insurance have little or no sensitivity to interest rate risk.

In practical terms, the decline in rates and in financial returns generally:

1. reduces margins due to the difference between the (declining) return on new assets and the guarantees or expectations of the customer, which can be exacerbated by competition between insurance companies;
2. increases the value of the liability more than the asset due to the duration gap (if A<L). An excessive gap can threaten the insurance company's solvency;
3. increases the risk of capital loss in the event of a large wave of policy surrenders against the backdrop of a swift, unexpected increase in rates;
4. in the latter scenario (rate hike) life insurance products will be less competitive than traditional bank savings products, which will offer more favourable returns while insurance products will take longer to align.

Guarantees are important…

According to a survey conducted in 2013 by the European Insurance and Occupational Pensions Authority (EIOPA), published in its analysis of the low rate environment (8), the bulk of the distribution of guarantees contracted by EU insurance companies (excluding unit-linked) is close to 3% with a not insignificant distribution tail above 4%. EIOPA cautions that its survey was voluntary and not exhaustive, and that it may not be representative. Furthermore, responses do not seem to have been weighted by the size of the respondents' balance sheets. The troubling conclusion that can be drawn from these data may therefore be skewed by the weight of small companies in the sample.

Figure 4: In 2013, most return guarantees were between 3% and 4%

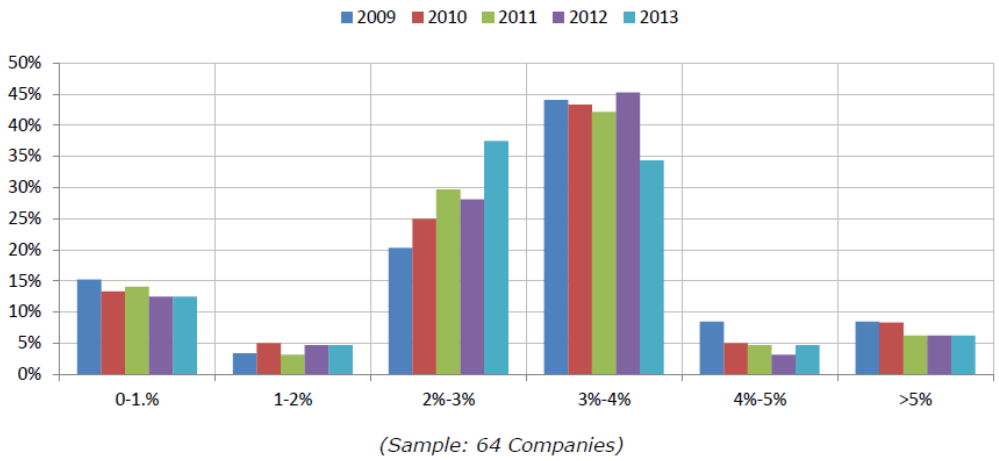

Source: EIOPA – 2014 (7)

We also see life insurance companies starting to adapt, with a shift in the guaranteed rate from 3.5% to 2.5% in 2013. This shift very likely continued in 2014 and 2015.

… but the main risk comes from the duration gap

According to EIOPA's 2013 survey, the duration gap for European insurance companies leans heavily towards liabilities (**A<L**), with a highly asymmetric **L-A** distribution spread between -1 year and +10 years. Even accounting for the doubts raised about the sampling and the non-weighting of the responses, the size of the positive duration gap nevertheless appears to show that a large number of European life insurance companies are massively exposed to low interest rate risk. Moreover, the level of guaranteed returns does not appear to be negatively correlated with the duration gap, meaning the two risks do not offset each other.

Figure 5: In general, the duration of liabilities greatly exceeds that of assets

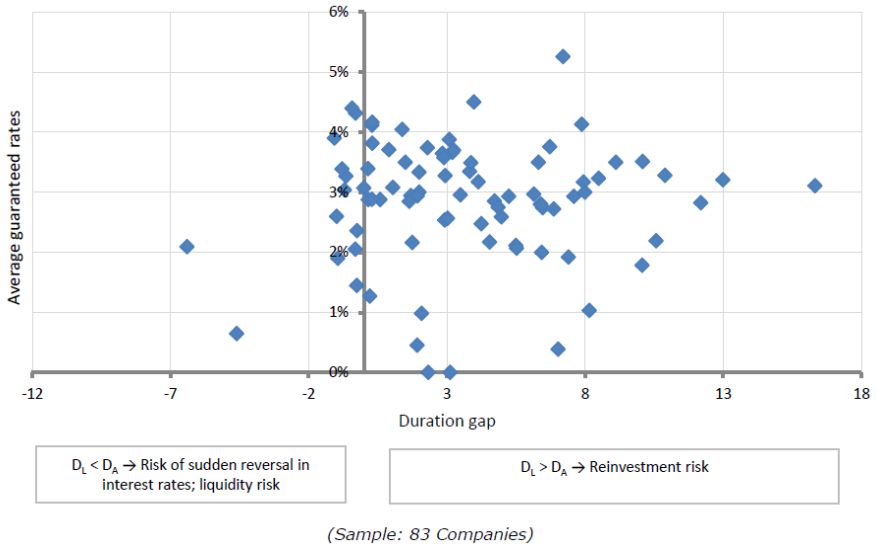

Source: EIOPA – 2014 (7)

Note that a minority of life insurance companies and most P&C insurance companies have a negative duration gap (**L<A**), which conversely exposes them to a high risk of capital loss if rates rise, and to a high solvency risk in the event of massive surrenders.

Coping strategies

1. Strategic diversification

The best strategic protection against low interest rate risk, since it mainly affects life insurance, is a balanced diversification of the business portfolio between protection/health, savings/retirement, unit-linked versus non-unit-linked products, etc. In Japan, the companies that survived the 2003 rate cuts were precisely those that had achieved the greatest diversification, in particular into health insurance. The major European companies are moving in the same direction, as can be seen in the AXA example.

Figure 6: Shift in AXA's business mix between 2010 and 2014 towards health and unit-linked

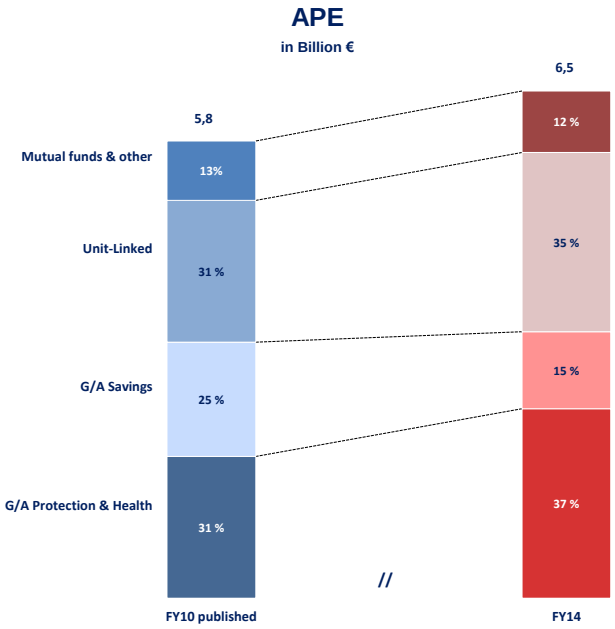

Source: AXA, 2014 results. APE = Annual Premium Equivalent

2. Reduction in guarantees

The Japanese precedent is once again instructive. One of the Japanese companies that has been the most successful in adapting to the deflationary and low-rate environment is Nippon Life (Nissay), now the country's second-largest insurance company after the failure of eight mid-sized insurers. The director of Nippon Life's New York office explained during an investor presentation that the secret to Nissay's resilience was how quickly it lowered guaranteed rates for its life insurance policies, from 6% in the early 1990s to 1.65% in the early 2000s. It is worth noting that in the early 2000s, the Japanese Diet passed a law, applicable retroactively, authorising insurance companies to renegotiate policies that had been contracted previously, given the extremely high risk of a string of bankruptcies, as well as the high degree of rigidity of these policies.

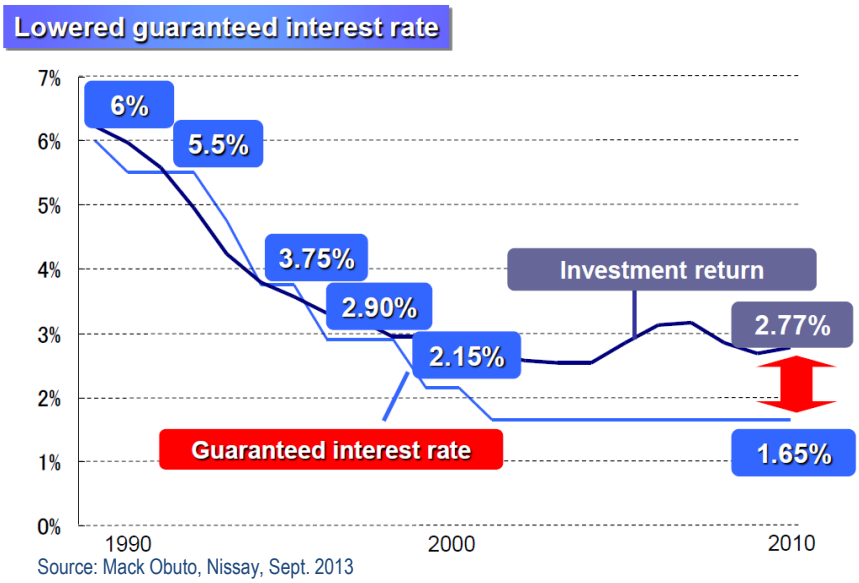

Source: Mack Obuto, Nissay, Sept. 2013

3. Asset management

Neither diversification nor active liabilities management, like guaranteed rate flexibility and incentives for unit-linked savings, can by themselves protect the life insurance segment from low-rate risk, due to the lingering rigidity of the contracted policies. These strategies are supplemented by management of the assets themselves, within the constraints of Solvency 2. Protection can be achieved in two key ways: extending asset duration and seeking higher returns for a given duration. Some of the most frequently used tactics are:

- Increasing the share of “peripheral” sovereigns in bond portfolios, which requires a detailed analysis of the sovereign risk within the Monetary Union, whose States are no longer truly sovereign within the monetary and fiscal meanings of the term;

- Moving higher up the risk/return curve within the investment-grade segment and, to a moderate extent, into high-yield. The entry into force of the European directive on bank recovery and resolution provided new opportunities from this perspective, for example, by raising the attractiveness of subordinated bonds relative to uncovered senior bonds;

- Shifting into non-traditional bond assets (loans, ABS, etc.);

- Taking advantage of illiquidity premiums (real estate, infrastructure, etc.).

It is worth noting that, in the case of infrastructure, the interests of insurance companies align with those of the States. The former, whose liabilities are by nature long term, are seeking to manage their assets and the latter would like to promote the long-term growth potential of the EU while reducing the short-term demand gap that many economists believe is responsible for low interest rates. A more growth-friendly regulatory approach would likely be able to ignite a virtuous circle, which would be positive for growth and interest rates while reducing the systemic risk cited by the IMF with respect to the European life insurance model (9).

References

(1) Bernanke, Ben (2005), "The Global Saving Glut and the U.S. Current Account Deficit", Sandridge Lecture, Virginia Association of Economists, Richmond (VA)

(2) Coeurdacier, Nicolas, Stéphane Guibaud and Keyu Jin (2015), "Credit Constraints and Growth in a Global Economy", American Economic Review 2015, 105 (9)

(3) Gordon, Robert J. (2014), "The Demise of U.S. Economic Growth: Restatement, Rebuttal, and Reflections", NBER WP 19895

(4) Summers, Lawrence H. (2014), "Reflections on the 'New Secular Stagnation Hypothesis'", in "Secular Stagnation: Facts, Causes, and Cures", edited by Coen Teulings and Richard Baldwin, CEPR Press, 2014

(5) Mokyr, Joel (2014), "Secular stagnation? not in your life" in – see (4)

(6) Adrian, Tobias, Richard K. Crump and Emanuel Moench (2013), "Pricing the Term Structure with Linear Regressions", Journal of Financial Economics 110

(7) Fischer, Stanley (2015), "Conducting Monetary Policy with a Large Balance Sheet", Speech at the 2015 U.S. Monetary Policy Forum, New York, February 27, 2015

(8) EIOPA (2014), "Low interest rate environment stock taking exercise", EIOPA-BoS-14/103, 28 November 2014

(9) de Bock, Reinout, Andrea Maechler and Nobuyasu Sugimoto, "European Life Insurers: Unsustainable Business Model", IMF Blog, 5 May 2015.

Equity market liquidity and long-term investing

Serge Darolles, Paris-Dauphine University

Abstract

This article addresses a range of issues related to the analysis of liquidity in equity markets from the perspective of long-term investors. The results of a survey of French institutional trading desks are the basis for the presentation of a model capturing the effects of information and liquidity on volumes and changes in daily prices. Using this approach allows us to create four major asset classes, depending on their exposure to different types of liquidity frictions. An empirical application shows that even in the supposedly liquid markets, many securities are subject to a deficient short-term liquidity provision, and/or strategic behaviour of long-term investors.

Résumé

Cet article aborde une série de questions liées à l'analyse de la liquidité des marchés actions du point de vue de l'investisseur long terme. Les résultats d'une enquête menée auprès d'institutionnels français servent de base à la présentation d'un modèle capturant les effets de l'information et de la liquidité sur les volumes et les variations de prix journaliers. L'utilisation de cette approche permet de créer quatre grandes catégories de titres, en fonction de leur exposition à différents types de frictions de liquidité. Une application empirique montre que, même sur des marchés supposés liquides, de nombreux titres sont exposés à un apport de liquidité court terme déficient, et/ou à un comportement stratégique de la part des investisseurs long terme.

Introduction

Equity market liquidity and the different ways it can be measured have been the focus of much academic research[1]. However, some aspects of liquidity are covered more often than others. These include immediacy, i.e. the ability to carry out a transaction instantly, and market depth, i.e. the possibility to place large orders. These are vital issues for short-term investors. Knowing whether they can very quickly enter and exit large-scale positions is critical ex ante information. But for long-term investors building a portfolio position over time, other aspects of liquidity are far more important. These include, for example, resilience, i.e. the markets' ability to quickly correct price deviations arising from transactions. Without this resilience, asset prices at a given point in time might no longer reflect their fundamental value, which makes investment fund valuations more difficult and undermines the equal treatment of unitholders.

It is not easy to measure the aspects of liquidity relating to long-term investing. Briefly, liquidity supply/demand and the arrival of new information both play a role in asset price formation. But neither information nor liquidity is directly observable. Only completed transactions can be used to measure their impact on prices, although it is not possible to separate the information effect from the liquidity effect. Over a very short period of time, one can always make the assumption that no new information will arrive. Changes in prices therefore depend only on the balance between the supply of and demand for liquidity. But this assumption does not hold up over the long term: prices therefore fluctuate based on information and liquidity. Because these two effects are fundamentally linked, it is particularly difficult to differentiate between them: an asset tends to be illiquid when it is the subject of substantial information flows.

A second factor makes it difficult to address these aspects of liquidity. The structure of the equity markets has evolved significantly in recent years, which is changing the way liquidity is provided. New technologies and the prevalence of automated trading have given rise to new liquidity players. The traditional market maker, which provided liquidity continuously and non-strategically, now has company. High-frequency traders (HFTs), with their ability to trade very quickly and on several markets at the same time, have become major players in liquidity provision. But their arrival raises a number of questions, mainly as to the quality of the liquidity provided. On the one hand, competition between HFTs and traditional market makers helps reduce transaction costs and therefore benefits short- and long-term investors. On the other, HFTs behave strategically and do not provide all the necessary liquidity when the risks related to this activity are too high. Yet it is precisely during these periods that the markets need it the most. This strategic behaviour triggers illiquidity spirals and increases systemic risk [see Brunnermeier and Pedersen (2008)]. This new environment also changes the behaviour of long-term investors. Will they also behave strategically in their search for liquidity? Will they behave defensively towards the HFTs' liquidity supply? And, by extension, how might these developments alter the liquidity supply/demand balance, and therefore the way in which liquidity is measured?

Against this backdrop of strategic behaviour on both the supply and demand sides, we need to rethink how we analyse liquidity. Trading volumes are a telling example. On the one hand, activity by HFTs will increase volumes on listed markets, without improving liquidity. On the other, strategic behaviour by investors tends to reduce volumes, either because they use alternative trading platforms or because they spread their transactions over time. It is difficult to determine the overall impact on volumes, which skews any volume-based liquidity measure. In this article we attempt to answer these questions by exploring two complementary approaches. Section 1 presents a summary of interviews conducted with major French institutional trading desks in September and October 2014. Topics discussed during this series of interviews included how to ascertain available liquidity, the solutions used to capture this liquidity and the impacts of the recent changes in the markets on their behaviour. This summary paints a picture of ideal liquidity for the long-term investor and offers some thoughts on how to improve the

[1] See for example Goyenko et al. (2009) who compare a large number of liquidity measures.

current situation. Section 2 draws on a model proposed by Darolles et al. (2015a, 2015b) and discusses how strategic behaviour by different market participants affects liquidity. Joint modelling of the impact of information and liquidity on volumes is used to separate the two components. By merely observing daily data, it is possible to propose a simple methodology for estimating the impacts of the strategic behaviour of providers and consumers of liquidity. The impact of activity by HFTs on liquidity provision and the consequences of strategic behaviour by long-term investors are analysed using the model. New ways of analysing liquidity that account for the strategic behaviour of the different participants may then be proposed.

1. Interview summary

The world of asset management is changing. Active management is contending with the surge in passive management and exchange-traded funds (ETFs). Asset managers focus on their value added and often delegate execution of buy/sell orders to dedicated departments, i.e. trading desks, which simultaneously process orders from different managers in the management company[2]. Centralising execution generates economies of scale, but can also be a source of conflicts of interest. For example, what guarantees are there that two orders concerning the same security, given simultaneously by two managers, will be processed in the same way, or that one manager will not systematically be given preferential treatment over another? Information about the nature of the orders could be lost due to this centralisation. Is a particular sale tied to profit-taking on a temporary price change or an automatic adjustment of the weight of different assets in a portfolio? The priority of the order depends directly on its nature, and this information is not systematically provided to the desk.

These desks have nevertheless become a vantage point for observing the major factors affecting the liquidity supply/demand balance. Before we explore different models, we thought it would be useful to find out what these intermediaries think about liquidity from a long-term investing standpoint.

[2] Some trading desks may also process external orders, from the management company's clients, for example.

1.1 How is liquidity ascertained?

Even trading desks find it impossible to accurately determine an asset's liquidity at a given point in time if no trade has occurred. Some of the liquidity is revealed once the trade has been made and observers have seen how the price reacts based on the characteristics of the order (side, size, etc.). The desk therefore only measures ex post liquidity, i.e. once this liquidity has been consumed. We like to use an image to convey this idea. Imagine that the market starts off as an empty container, which is filled up with an order. The market is liquid if the container is immediately emptied. It is clear from this example that, without the initial order, it is impossible to observe the system's ability to respond to the imbalance triggered by the order. The system responds immediately if liquidity is provided non-strategically. Conversely, if the liquidity provider acts strategically, for example, by letting the container fill up, it can then take advantage of a price mismatch and subsequently make a gain on future price convergence. The major challenge for the desk is then to make an *ex ante* forecast, based on the characteristics of the order, of the price mismatch (referred to as "market impact") and therefore of the cost of the trade. Traditional liquidity measures attempt to forecast this cost based on observable quantities, such as volume. But if this volume is altered by the strategic behaviour of the liquidity provider, it gives a biased view of liquidity. However, volumes generated by HFTs on the US markets are estimated at 70% (30-35% in Europe). This figure is attributable to the many transactions carried out between HFTs, trades that increase volumes and create liquidity frictions in the short term (intraday) but do not really provide liquidity.

Desks have various ways to ascertain liquidity. Those that rely on quantitative approaches use historical liquidity indicators, calculated from observations of past transactions and used to forecast future liquidity. These desks have access to sophisticated trading cost analysis (TCA) tools, which are used to anticipate the volatility generated by the arrival of an order. But these tools are still based on a set of assumptions that are not easily verifiable. We can see very quickly that these are black boxes and that we have no knowledge of their internal workings. Other desks take more qualitative approaches and analyse how liquidity is shown to the market. The market participants' experience is therefore a critical factor in the liquidity analysis.

All the desks agree, however, on differentiating between the pre- and post-MiFID periods[3]. Pre-MiFID, or before November 2007, execution was a simple task. Above a certain size, it was difficult to find enough liquidity on the single listed market to place the order. The desk's role was then to find blocks large enough to make the trade. Post-MIFID, execution became more complex. Liquidity can be found in many places: listed markets, alternative trading platforms, dark pools, etc., each with its own characteristics and requiring different expertise. The post-MiFID market fragmentation has forced investors to adapt to new conditions and to treat liquidity differently depending on where the trade is made.

Two separate concepts of liquidity come into play. The first corresponds to *available*, or *active*, liquidity. This is the liquidity posted to the order book in listed markets. It can be quantified without difficulty, although it is not so easily captured. It is possible to cancel buy/sell orders, and a portion of the posted liquidity can be made to disappear very quickly. The second corresponds to *hidden*, or *passive* liquidity. This is liquidity that must be sought on alternative trading venues, but also from counterparties who do not necessarily express their interest to the market. It is of course more difficult to measure because of fragmentation, and achieving it comes at a cost to the investor. In the end, fragmentation affects the trading desks' behaviour and their execution strategies regardless of the type of liquidity in question. Dealing with posted liquidity requires efficient access to the markets to be able to immediately capture this liquidity, by using sophisticated execution algorithms, for example. Dealing with hidden liquidity requires

[3] MiFID stands for the Market in Financial Instruments Directive (Directive 2004/39/EC) governing the provision of investment services in financial instruments by banks and the operation of traditional stock exchanges and alternative trading venues

reliance on market traders' experience to limit the costs associated with seeking out this liquidity.

1.2 How is liquidity found?

Interviews with the desks show a lack of consensus on this point. The strategies used by these desks to find liquidity depend on their trading culture and operating constraints, which relate mainly to how the desks' efficiency is measured. A “quantitative” desk executes orders received based 70-95% on algorithms. This has many advantages. First, it reduces transaction costs and improves the desk's financial performance when measured solely in terms of cost. It is also a way to address market fragmentation. An algorithm can easily test the liquidity of several markets simultaneously and execute the order on the market that has the best option. An algorithm can also break up large orders over time. The unit size of the orders decreases, which makes them less visible and keeps the desks away from strategic liquidity providers. Controlling execution through an algorithm requires access to the Smart Order Router tool. In general this tool is provided by the broker that gives the desk access to the markets. The desk is responsible for setting the final parameters.

Conversely, a “qualitative” desk does not restrict itself to posted liquidity and seeks to spark interest among potential counterparties. For large orders in particular, the desk's role is to determine the best way to find hidden liquidity. In that case, it is critical to have strong bilateral relationships with brokers so desks can ask what the brokers are offering for a given size, similar to trades on over-the-counter (OTC) markets. For smaller orders, the only choice variable is transaction costs, and in that case the algorithms are used. We also note desks' increasing interest in liquidity netting solutions (Liquidnet, POSIT, etc.). The idea is to take the economies of scale afforded by the desks one step further and create clubs made up of investors with an interest in trading among themselves. These solutions are clearly designed to hide some of the liquidity from certain market participants, primarily HFTs. Consequently, the liquidity no longer all appears on the market immediately; it is only reported ex post to all participants. This is problematic for algorithms, which need an ex ante big-picture view of the liquidity offered by the various trading platforms. This solution also raises an issue of price fixing for the transaction. If liquidity has completely disappeared and no market price is available, it is difficult to set a price without a natural price discovery process. Yet these illiquidity situations are those in which these alternative solutions are desired or needed. An incorrect transaction price can also lead to unequal treatment of fund unitholders entering into this type of OTC transaction.

1.3 Strategic investor behaviour

When the desks are asked about the impacts of market fragmentation with the introduction of alternative trading platforms and dark pools, the responses vary widely. Some consider that fragmentation increases the likelihood of being detected by strategic liquidity providers, while others believe it protects against them, as it is easier to hide orders on different platforms. The response often depended on the desk's trading culture.

A “quantitative” desk that mostly uses algorithms tends to generate small orders by spreading trading volumes over time (called “slicing”), a practice that aims to make these orders harder to detect. However, gaming is the main risk facing algorithms. This refers to the propensity of certain liquidity providers to filter systematic buy/sell behaviour and exploit this information. To combat these practices, trading algorithms have been updated to include anti-gaming strategies. Their aim is to conceal, to the extent possible, the source of orders sent by algorithms, by reducing order size on organised markets (slicing) or using dark pools as an alternative to these markets. These alternatives are used to prevent over-slicing and thus ensure that the time needed to place the entire order is not automatically extended. The specific way in which dark pools work also has an impact on order size. To avoid having HFTs as trading counterparties, it is much more efficient with a dark pool to increase order size than to reduce it. The amounts put in play by HFTs on each transaction are generally low, and they do

not participate in large transactions. As such, sending large orders keeps HFTs out of the transaction. In the end, desks that use algorithms behave defensively, and when they attempt to use their weapons to respond to HFTs they seem to be at a serious disadvantage, like a clean athlete trying to compete against an athlete on steroids.

For a “qualitative” desk taking a more diversified approach to liquidity, reducing order size does not seem to be the right solution. Such a desk may be pulled in two directions. The first is to use market fragmentation to go where HFTs do not. As many HFTs are found on alternative trading platforms, the solution is to use another well-identified long-term investor as trading counterparty through liquidity netting solutions. This amounts to establishing bilateral relationships among investors outside the markets, and thus cut off from the transparency they provide. The second is to work even more actively with the orders to avoid detection by the HFTs. This signals a “fight for liquidity” approach. In this combat, the “qualitative” desk is not necessarily on the defensive because it does not use the same weapons as HFTs. Its behaviour is more strategic than defensive, as all the information it has is not immediately passed along to the market. Its advantage is basically control over the flow, as opposed to HFTs, which rely on investors' orders. But this strategic behaviour is problematic and creates new types of liquidity frictions in the long term (beyond the day of the transaction). By slowing flows, investors reduce volumes, hurt the quality of the market and therefore alter the price formation process. The impacts HFTs have on liquidity (short-term liquidity frictions) must therefore be measured directly as well as indirectly, via the strategic behaviour of long-term investors resulting from the presence of HFTs (long-term liquidity frictions).

We conclude this summary of interviews conducted with trading desks by raising two key questions. While volumes alone provide biased information, it should be possible to measure the share of total volumes linked to liquidity providers. In other words, what is the share of liquidity volumes versus arbitrage volumes, i.e. that linked to the provision of liquidity by HFTs which artificially inflates reported volumes? The second question concerns the extent of the impact of investors' strategic behaviour on market quality. To what extent is the strategic behaviour of these investors at the origin of inefficiencies? The only way to answer these questions is to carefully model the relationship between information and liquidity, a topic we discuss in the following section.

2. A fresh approach to analysing liquidity

The interviews conducted with trading desks showed that volume alone was not enough to analyse liquidity and that participants' behaviour had a direct impact on liquidity. In this section we consider recent academic literature on the relationships between information and volume and between liquidity and volume. We then discuss new liquidity analyses that have been adapted to the long-term investment context and that enable us to account for the strategic behaviour of both providers and consumers of liquidity.

2.1 The literature

The empirical literature on the relationships between information, liquidity and volume draws on sound theoretical bases, and a number of academic articles address these topics. We begin our review of the literature with the model proposed by Tauchen and Pitts (1983), which explores the relationship between information and volume. We then discuss other models that incorporate the liquidity aspect into this study. Tauchen and Pitts (1983) consider a single risky asset, presumed to be extremely liquid, and a finite number of traders with heterogeneous reservation prices. These traders trade only in response to the arrival of new information on the asset's fundamental value. The price changes and volumes generated thus stem only from revisions to reservation prices, and therefore from information. A trading day consequently consists of a series of consecutive equilibriums initiated by a random number *I* of new pieces of information. For each day, price changes and daily volumes correspond to the sum of price changes and intraday volumes, and follow independent Gaussian distributions whose variance depends on unobservable latent variable *I* [see Tauchen and Pitts (1983)]. In particular, the greater the value of *I*, the higher the average trading volume and the larger the variance in volumes and price changes. Conversely, the average of the price changes remains unchanged regardless of the value of latent variable *I*.

The implications of this first model are intuitive and relatively easy to ascertain using a narrow set of assumptions. They also allow us to state that price changes follow a stochastic volatility model. It is therefore easy to estimate the parameters of the model using the usual econometric approaches, for example, the generalised method of moments. However, this approach does not account for any liquidity friction and cannot explain the correlations observed between price changes in financial assets over two consecutive days.

Richardson and Smith (1994) propose the first extension to the Tauchen-Pitts (1983) model. Here, information *I* for a given day affects both the average of and variance in price changes. Andersen (1996) and Liesenfeld (2001) assume persistent information over time to generate time dependence between changes in asset prices over two consecutive periods. Anderson (1996) also modifies the equation that gives daily volumes based on Glosten and Milgrom (1985) to account for liquidity via an increase in volumes. But, unlike Tauchen and Pitts (1983), most of these extensions are based on ad hoc assumptions that are hard to justify from a theoretical standpoint.

Darolles et al. (2015a) start from Tauchen and Pitts (1983) and use Grossman and Miller (1988) to enter price changes and daily volumes into a model with two unobservable latent variables, *I* (information) and *L* (liquidity). The proposed model thus replicates the behaviour of HFTs acting as strategic liquidity providers that do not systematically provide all the necessary liquidity. HFTs act when asset prices are sufficiently far from their fundamental value and make their money on the convergence of prices towards their fundamental value. This convergence is not immediate for reasons related to the liquidity providers' strategy and the funding constraints they face[4]. In this first extension, average volumes become an increasing function of information *I* and liquidity *L*, while the variance in price changes still depends only on information *I*. A simple measure of frictions arising from the strategic provision of liquidity in the short term can therefore be proposed. It is simply a question of the accessible share of the daily volume

[4] See Gromb-Vayanos (2010) for a survey on the limits of arbitrage.

relative to the volume corresponding to the variance in daily price changes. In particular, for a given level of variance, a security's exposure to short-term liquidity frictions increases as the daily volume rises.

Darolles et al. (2015b) propose a second extension of Tauchen and Pitts (2013) to capture the strategic behaviour of investors as well. As with Darolles et al. (2015a), liquidity providers always act strategically with respect to liquidity, which leads to higher intraday volumes. But investors' strategic response in their search for liquidity affects volumes on the following day by spreading out orders, which creates liquidity frictions in the long term. Darolles et al. (2015b) assume that investors minimise their market impact by splitting up their orders and consequently pass the information they have on to the markets over several days. Based on that assumption, and comparable to Richardson and Smith (1994), information *I* becomes time-correlated and affects the price change average. This ultimately generates an autocorrelation between price changes over two consecutive days. The model therefore captures both the strategic behaviour of liquidity providers (through an increase in volumes due to short-term frictions) and the strategic behaviour of investors (through the autocorrelation of price changes). The challenge to the market efficiency assumption is directly tied to the strategic behaviour of the latter.

2.2 A fresh analysis of liquidity

Darolles et al. (2015b) thus propose a way to model volumes that accounts for both information and liquidity. The economic implications of the model depend on the values of the different parameters, and defining different categories permits an analysis of the liquidity characteristics of the assets. Estimating the parameters of the model for a given asset then makes it possible to categorise each security and consequently to understand what type of illiquidity it is exposed to. This is done by testing the statistical significance of a subset of the model's parameters.

The first category corresponds to all liquid assets. For these securities, volumes depend only on information. This puts us in the realm of Tauchen and Pitt (1983). Liquidity providers and investors do not behave strategically, there are no volumes relating to short-term liquidity frictions and there is also no autocorrelation of price changes. The second category is that of assets subject only to short-term liquidity frictions. Here, volumes depend on information and liquidity, and the price change average is constant, as in Darolles et al. (2015a). Liquidity providers behave strategically and there are some arbitrage volumes, but investors do not behave strategically, which is why there is no autocorrelation between price changes. The third category is the mirror image of the second: only investors behave strategically and consequently there are long-term liquidity frictions. Volumes depend only on information, such as the price change average. This puts us in the realm of Richardson and Smith (1994). Liquidity providers do not behave strategically and there are no arbitrage volumes, but investors behave strategically, which explains the autocorrelation between price changes. Lastly, the fourth category includes securities subject to short- and long-term frictions. Average volumes depend on information and liquidity, and price changes on information. Liquidity providers and investors behave strategically, which is why there are arbitrage volumes and an autocorrelation in price changes.

2.3 Estimation and filtering

The presence of two latent processes relating to information and liquidity makes the estimation of the model proposed by Darolles et al. (2015b) more difficult. It is done in two steps using the model's triangular structure. Price changes depend only on information, while volumes depend on both liquidity and information. The first step is therefore to estimate the parameters of the price change equation, then to filter the information process based on the prices observed on each date. Darolles et al. (2015b) therefore estimate the model's parameters for price changes, and filter information *I*. Traditional approaches to the estimation of stochastic volatility models

lead to very high estimated levels of volatility persistence. Considering a regime switch in the volatility parameter generally reduces these persistence effects [see Hwang, Satchell and Pereira (2007) for a description of the estimation method].

In the second step, the information reflected in the volumes is replaced by the information that was filtered in the first step. This gives us a volume equation that depends only on liquidity. A simple Kalman filter enables us to estimate the parameters of the volume equation and to filter liquidity for each date.

2.4 Initial empirical results

Darolles et al. (2015b) applied their approach to daily returns and volumes for 92 securities that were part of the FTSE100 index on 27 March 2014, for the period from January 2010 to December 2013, or 1,043 daily observations. The main results can be found in Table 1.

Table 1: Number of securities in each of the four categories of liquidity defined according to significant parameters in the Darolles et al. (2015b) model

	No long-term liquidity friction	Long-term liquidity friction
No short-term liquidity friction	36	15
No short-term liquidity friction	32	9

In this universe of presumably liquid stocks, they find that a number of assets (61% of all assets, or 56 out of 92 securities) experience different types of liquidity frictions. Of these illiquid securities, 57% (32 securities) should be classified only in the short-term friction category, 27% (15 securities) only in the long-term friction category, and 16% (9 securities) in the short- and long-term liquidity friction category. Another way to analyse these figures is to count up the securities associated with at least long-term liquidity frictions, or 24 securities out of the total of 92 (26%). The model suggests in particular that, due to investors' strategic behaviour, end-of-day prices (used for funds that are valued daily) for these securities do not reflect the fundamental value of the assets.

Beyond the direct liquidity analysis, these results can be used to develop arbitrage strategies that could improve the functioning of the markets and reduce liquidity frictions. Short-term frictions, for example, result from an insufficient number of liquidity providers. The latter can then slow their liquidity provision to create price mismatches and subsequently make money on the convergence. For securities that have these frictions, liquidity provision needs to be increased through the use of short-term mean-reverting strategies. This affects 41 of the 92 securities in the empirical study, or 44% of the total sample. Long-term frictions affect 24 securities, or 26%. Here, the strategic behaviour of investors who stagger their orders and slow the price discovery process is problematic. The remedy is then to implement long-term momentum-type strategies likely to foster the incorporation of available information into prices and thereby reduce the autocorrelation observed in the daily return series.

Conclusion

Our objective in this article was to respond to a range of issues related to equity market liquidity from the perspective of long-term investors. A survey of major institutional trading desks raised a series of questions about the strategic behaviour of different market participants in terms of

liquidity. Models in the academic literature were used to quantify short- and long-term liquidity frictions and thus provided answers to some of these questions. These models are also the springboard for analyses of different forms of liquidity based on price changes, volumes and autocorrelations. The implications of these approaches and the initial results obtained for securities traded on supposedly liquid markets are clear. Liquidity providers, just like investors, act strategically with respect to liquidity, and generate liquidity frictions that can undermine the orderly functioning of the markets and their ability to ensure a quick and efficient price discovery process. This article also suggests ways to reduce the impact of these frictions, either by promoting liquidity provision in the short term or by accelerating the incorporation of information in the long term.

References

Andersen, T. (1996). Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility. Journal of Finance, 51(1), p. 169-204.

Brunnermeier, M. and L. Pedersen (2008). Market Liquidity and Funding Liquidity. Review of Financial Studies, 22, p. 2201-2238.

Darolles, S., G. Le Fol and G. Mero (2015a). Measuring the Liquidity Part of Volume. Journal of Banking & Finance, 50, p. 92-105.`

Darolles, S., G. Le Fol and G. Mero (2015b). Financial Market Liquidity: Who Is Acting Strategically? Working Paper.

Glosten, L. and P. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics, 14(1), p. 71-100

Goyenko, R., C. Holden and C. Trzcinka (2009). Do Liquidity Measures Measure Liquidity? Journal of Financial Economics, 92(2), p. 153–181.

Grossman, S. and M. Miller (1988). Liquidity and Market Structure. The Journal of Finance, 43(3), p. 617–633.

Hwang, S., S. Satchell and P. Pereira (2007). How Persistent Is Stock Return Volatility? An Answer with Markov Regime Switching Stochastic Volatility Models. Journal of Business, Finance and Accounting, 34, p. 1002–1024.

Liesenfeld, R. (2001). A generalized bivariate mixture model for stock price volatility and trading volume. Journal of Econometrics, 104(1), p. 141-178.

Richardson, M. and T. Smith (1994). A Direct Test of the Mixture of Distributions Hypothesis: Measuring the Daily Flow of Information. Journal of Financial and Quantitative Analysis, 29(1), p. 101-116.

Tauchen, G. and M. Pitts (1983). The Price Variability-Volume Relationship on Speculative Markets. Econometrica, 51(2), p. 485-505.