Willis Towers Watson

Understanding and measuring
the illiquidity risk premium

Asset Research Team, March 2016

Introduction

Not every investor has the benefit of a long time horizon, but many of Willis Towers Watson's institutional investment clients do. Due to their liability structures, a significant number of pension funds, most sovereign wealth funds/endowments and some insurance companies, have the ability to 'lock up' their capital to some degree. This represents an important source of competitive advantage, with these investors able to harvest the higher returns on offer from illiquid assets.

However, even for our long-term clients, it is important to understand how much is on offer from a given asset class from time to time, and whether that is sufficient or not. After all, 'what gets measured gets managed', so if we do not have a good way of measuring illiquidity risk premia, how are we to manage and prioritise our exposure?

Over the years, the importance of illiquidity to our clients has led us to undertake significant work in this area. This paper aims to summarise this work, and is structured as follows:

- In **Section 1,** we take a step back and ask the question: how do we judge the minimum level of additional return required to encourage an investor to accept the increased uncertainty and the reduction in flexibility that illiquidity demands?
- **Section 2** then discusses our approach to determining what a particular asset class is actually offering in compensation for its illiquidity from time to time. We summarise this information in our proprietary Illiquidity Risk Premium Index.
- **Section 3** summarises the main messages and discusses the drivers and implications of what appear to be lower levels of illiquidity risk premia despite lower liquidity.

What is liquidity?

First of all, what exactly do we mean by liquidity and its opposite, illiquidity? For us, liquidity involves the presence of three factors:

i) The ability to trade in sufficient volume
ii) Without negatively impacting price
iii) All with some level of confidence

The absence of any of these factors renders an asset – to some degree – illiquid. The additional return an illiquid asset offers the investor over a liquid alternative we refer to as the illiquidity risk premium (IRP).

1. The required illiquidity risk premium

The IRP compensates the investor for tying up its capital. When an investor accepts illiquidity, it accepts an increase in the uncertainty of end outcome because it is less able to liquidate the asset should something not turn out as expected. Even if the asset can be liquidated, its illiquidity manifests in lower certainty over the price available.

Related to this is that the investor also accepts a reduction in flexibility since it cannot replace the asset with a 'better' alternative (or it is potentially more costly to do so) should that become attractive. As a consequence, illiquid assets should provide a greater level of return to compensate investors for these negative aspects.

The level of IRP an investor should demand for a given asset – or required/fair-value IRP – has three dimensions in our experience:

1. The investor's utility function: some investors will have a higher tolerance for illiquidity risk, perhaps because their time horizons are longer, because they are less likely to take advantage of the flexibility higher liquidity offers, or because the rest of their assets happen to be highly liquid and it is total portfolio liquidity with which they are concerned. These investors should be willing to accept a lower IRP than investors with a lower tolerance for illiquidity.

2. The level of illiquidity: different assets will demand that capital is locked up for different lengths of time. This time period is often not straightforward to define: some assets are completely illiquid for a given period, whereas others may at first glance be illiquid, but in practice a secondary market (of varying depth and robustness) is available. Other assets are liquid, until they are not (which often coincides with times when one would wish to sell or buy). Furthermore, pretty much anything can be sold if there is complete flexibility over price (including negative prices), therefore there is a 'pricing confidence' aspect here. The return needed should increase with the degree of this illiquidity, which is multidimensional.

3. The volatility and uncertainty of the underlying cash flows: if an illiquid asset's cash flows are more volatile, or in a broader sense less 'knowable', than an alternative, equally illiquid asset, the investor should demand a higher required return. This is because the chances of negative outcomes and hence the value of being able to liquidate that asset are higher for more volatile assets (analogous to an option prices increase with the volatility of the underlying asset). Note, because we are often working with assets that are not marked-to-market, it is not sufficient to gauge this by historical price volatility (because the volatility of the prices of illiquid assets will be depressed by the absence of trading). Therefore, we need to make a broader assessment of the volatility of underlying cash flows.

It would be useful to establish a required illiquidity 'surface' showing the necessary level of illiquidity premium along these three dimensions. However, the challenge of precisely defining any of these variables should be apparent from the definitions above. In particular, an investor's utility function is very difficult to define, being dependent on a great many idiosyncratic and qualitative factors specific to that institution and other assets it is holding in its portfolio.

However, whilst we cannot make specific estimates of required or fair-value IRPs, we can study the academic literature to determine, according to theory and empirical evidence (for example, of the trading behaviour of equities of differing levels of liquidity), what sorts of IRP investors have demanded on average. We are also able to consult our experience of investing in illiquid assets over the years, and judge what sort of IRP has typically been demanded by ourselves and our clients over time. We provide the estimate of required/fair-value IRPs, based on the 'average' investor utility function, in the table below:

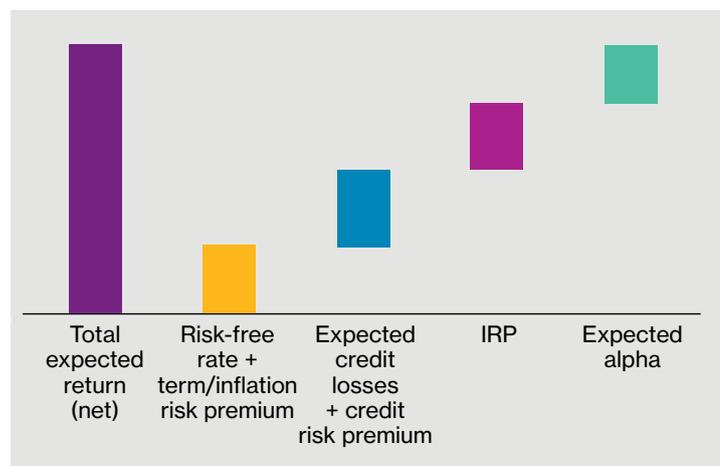| Estimate of average required IRP		Asset cash flow volatility	
		Low	High
Level of illiquidity	Low	Lowest: 50-100 bps	Higher: 100-150 bps
	High	Higher: 100-150 bps	Highest: 150-250 bps +

2. The expected illiquidity risk premium

Having an idea, even approximate, of the required or fair-value IRP is one thing, but how are we to identify the IRP of a given asset at a given time? In addition, how can we determine which assets provide the most attractive areas in which to "spend" whatever tolerance for illiquidity we have? Broadly speaking, we follow the straightforward process set out below:

- Establish some form of prospective yield or return: Critically this must capture, as much as possible, known uplifts to cash flows, such as those related to inflation linkage. This yield should be net of costs such as transaction costs, management costs and depreciation costs.

- Subtract an appropriate risk-free rate: This will strip out the market's estimated cash rate over the term of the asset, plus whatever bond risk premium (to compensate for inflation and interest rate uncertainty) it is currently demanding.

- Adjust for other risk premia: The focus here is on adjusting for credit risk (expected losses from defaults/migration, plus additional credit risk premium) and, where possible, equity-like risk (such as anything that causes cash flow uncertainty, for example voids/vacancies in commercial real estate ownership).

The residual expected return we assume represents the IRP. *Figure 1* below illustrates this further.

This approach is simple in principle, but complex in practice:

- It rests on a reliable estimate of expected return, which requires a significant amount of judgement. This also means it lends itself better to assets that generate more return from income than capital appreciation (that is, it works better for debt-like assets than equity-like assets).

- Expected cash rates and bond risk premia we take from bond yields, but neither expected credit losses (from defaults and downgrades) nor credit risk premia are directly available from market pricing. We use structural models of credit spread composition in public credit markets to estimate these (see Box 1 on page 6). This market-consistent measure must be adjusted for differing levels of seniority and security of the asset in question relative to a vanilla, publicly quoted corporate bond and any other important idiosyncratic factors.

Figure 1. **The building blocks of expected returns**

- For an actively managed strategy (often required in illiquid asset classes), additional returns may be expected from net alpha. In practice, 'alpha' may be hard to isolate in an illiquid asset as it often manifests in a lower expected default rate/higher expected credit or IRP. However, given we are trying to isolate the prospective IRP, we attempt to make this distinction.

- In treating the IRP as a residual and using market prices to determine other risk premia, we may capture any mispricing of these risk premia as mispricing of the IRP, depending on how our expected return is generated. At times when the market is demanding very low (high) levels for credit/inflation/term risk premia in the prices of bonds and credit assets, this may emerge as an attractive (unattractive) IRP. Furthermore, this residual is subject to 'noise' in the sense that small relative movements in the prices of other assets may cause larger moves in the measured IRP.

Whilst it has its limitations, we believe this approach is extremely powerful as it is a) intuitive, b) fairly simple in principle and c) the only one we have! Box 1 overleaf discusses the modelling approach we use to decompose publicly listed corporate bond spreads into their component parts, in order to identify the IRP available from public credit markets. This is worth special attention as the market-consistent expected losses, credit risk premia and IRPs of publicly traded corporate debt are important under our approach to determining the appropriate risk premia to deduct from the expected returns of other assets exposed to credit risk.
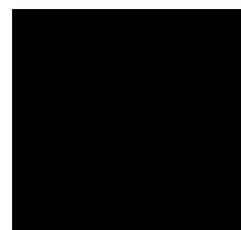
Box 1 – Decomposing credit spreads

The additional yield a corporate bond provides over a risk-free (or mostly risk-free) alternative of the same maturity is known as the credit spread. This is paid to/demanded by investors in order to compensate for:

- Expected credit losses, which can arise due to defaults on the payments due on the bond, or in some cases due to a credit rating downgrade causing a permanent loss of capital on the bond (for example, a BBB-rated bond could get downgraded, forcing investment-grade investors to liquidate at a time of falling prices).

- Further risk aversion around credit risk. In addition to these expected credit losses, risk-averse investors will demand an extra premium to compensate for the uncertainty of those losses, and the left-tailed nature of credit as an asset class (lots of modest positive returns, but the potential for large negative returns).

- The lower liquidity of corporate bonds versus risk-free equivalents such as US treasuries or UK gilts – an IRP.

Structural credit models enable us to split credit spreads into these component parts by making a limited number of assumptions. Our preferred implementation of these models uses the Leland-Toft model as its foundation, which takes Merton's observation that a firm's total equity value can be modelled as a call option on the firm's assets. It follows that the firm's debt value should be the present value of the firm's assets minus the value of this option.

By valuing this option using the market's expectation of future volatility (taken from equity option-implied volatility), we are able to estimate – in a market-consistent manner – how much spread the market is demanding for credit losses and the additional risk aversion-related credit risk premium. Further, by valuing this option in a risk-neutral world we are able to isolate how much of this spread relates to expected credit losses alone. Finally, the residual 'unexplained' spread is the market-consistent amount of IRP from corporate bonds.

The charts below show the decomposition over time of US investment-grade corporate bond spreads over US (including non-investment-grade) bonds at the most recent date. Having a decomposition by ratings band is an important innovation that enables us to size the credit risk premia of other assets if we are able to estimate the credit quality of that asset.

US investment-grade decomposition over time
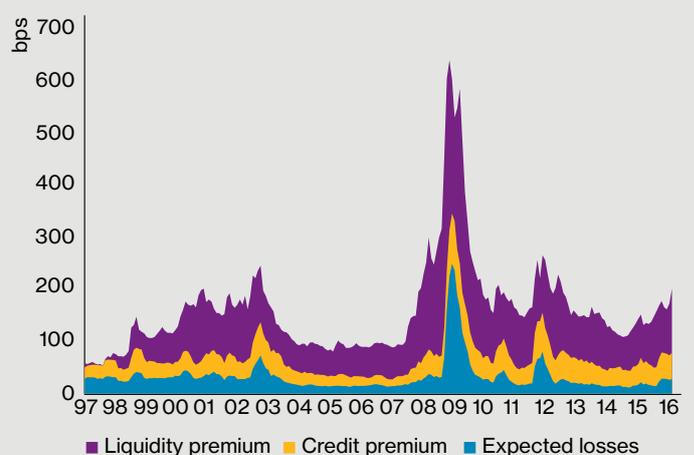
Decomposition by ratings band as at 31 January 2016
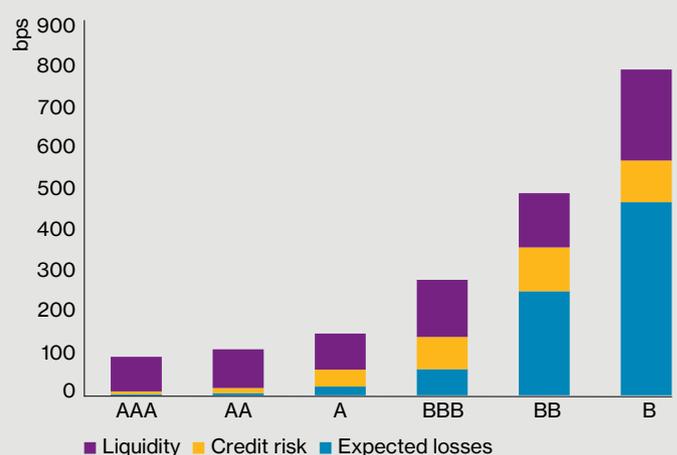
Expected IRPs in practice – European direct lending

This generic approach can be brought to life with an example. In recent years, we have been advocates of allocations to a specific European direct lending strategy. The economic conditions in the eurozone and the UK and the continued regulatory pressure being applied, particularly to the balance sheet size of highly levered European banks, have increased banks' cost of lending and created sufficient room for the well-connected investor to lend to mid-sized companies at attractive rates of return.

Figure 2 provides the breakdown of our forward-looking expected returns at two different points in time. The decomposition uses the same categories (and colour coding) as Figure 1.

Details behind these numbers:

- The risk-free rate is the market-discounted cash rate at the time, including risk premium, for the maturity of a typical private loan (five years) in the UK and the eurozone.

- Expected credit losses and credit risk premia are taken from our Leland-Toft models applied to European and UK credit markets, weighted by the estimated credit quality of the counterparties to any loans and adjusted lower to reflect the increased seniority of direct lending relative to unsecured corporate bonds.

- Alpha in this context is inexorably wrapped up with a lower default exposure and loss given default. Therefore, this number is our estimate of how much lower default experience will be if implemented through our preferred (and we believe skilled) manager.

Figure 2 illustrates how our changing expectations of returns from this asset class interact with changing credit risk and alpha potential, and deliver a shifting estimate of IRP on offer. The overall conclusion for European direct lending is that whilst the IRP on offer has declined significantly since initially recommending the strategy, it remained reasonable at the start of 2015.

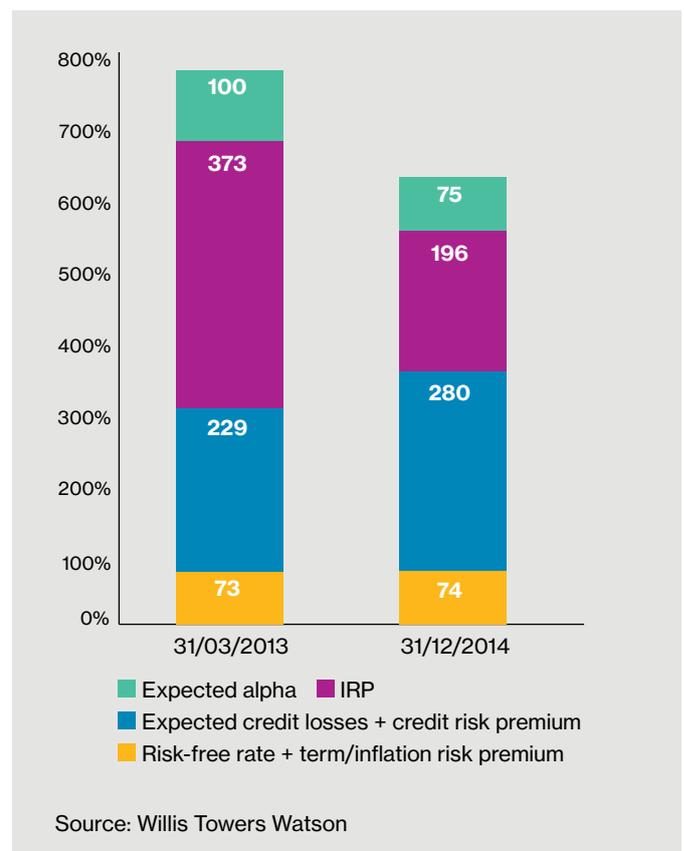Figure 2. **European direct lending expected return decomposition (bps)**
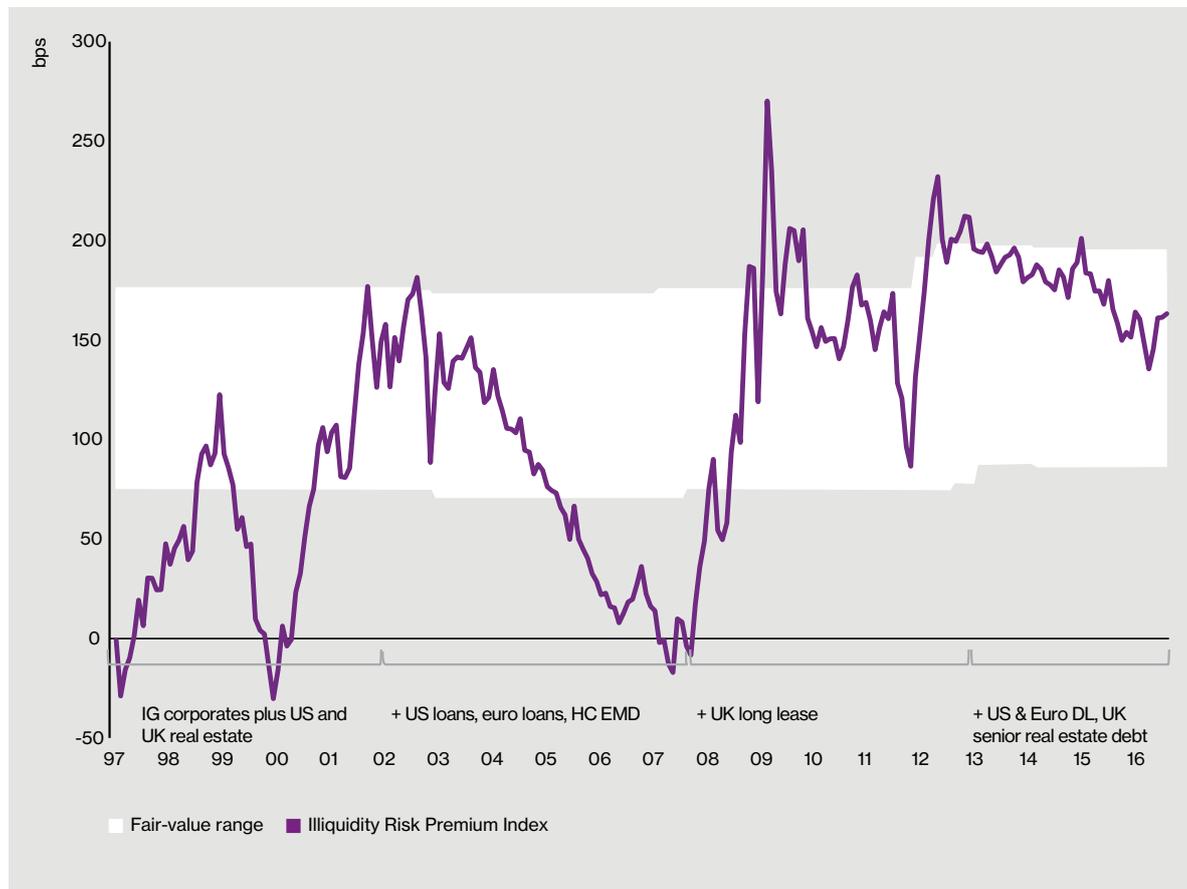

Source: Willis Towers Watson

Figure 3. **The Willis Towers Watson Illiquidity Risk Premium Index**

Combining measurements of IRPs

By following this process, we are able to estimate, over time, the IRP on offer from a variety of different assets of varying liquidity. First of all, this enables us to compare IRPs across assets on a consistent basis, and make relative-value statements about the attractiveness of taking illiquidity risk across those assets. This is a useful input in portfolio construction when we are determining how to 'spend' a given illiquidity budget across the opportunity set.

It is beyond the scope of this article to cover all the assets we have implemented this for, but we have combined all these measures in our own Illiquidity Risk Premium Index, which is shown in *Figure 3* above. This takes a simple average of all the assets we have data for at a particular point in time, and plots this against the approximate fair value (based on the ranges developed in Section 1).

The Index does have a few limitations that are worth noting. The sample of assets included is restricted by data availability, therefore the fair-value level (itself uncertain) does change over time albeit not by much. In addition, our sample of assets is prone to 'positive-selection bias' in the sense that the sample of assets included is biased towards those we have conducted sufficient research on, which will tend to be those assets that appear to offer attractive risk premia a priori. Furthermore, the Index is focused on measuring the IRP on offer from the more illiquid end of the opportunity set – there are other measures of illiquidity generated from more liquid assets (TED spread, small versus mid equity bid-offers and so forth) that are not included.

With these limitations in mind, the messages emerging from this aggregation of IRPs in the assets we cover are fairly intuitive:

- Historically, there have been better and worse times to take illiquidity risk.

- The fact that there are longer periods when our measure is below fair value (illiquidity is poorly rewarded) than above may imply that:
  - Investors' average required illiquidity risk premia are too high relative to what is available.
  - We have misestimated that average.
  - Or, it may reveal that the maxim 'what gets measured gets managed' holds true: we believe our approach to assessing the IRP is fairly unique, and therefore it seems likely that the IRP was too low simply because the average investor did not appreciate it was too low.

- In the years immediately prior to the global financial crisis, illiquidity was poorly rewarded.

- In recent years, the additional returns on offer from illiquid assets have been relatively attractive and should have encouraged the long-term investor to exploit them.

- The level of this attractiveness has reduced in recent years, although this may understate the true level of IRP decline due to the sample of assets we measure (which, as mentioned above, is biased towards assets likely to offer more attractive IRPs).

As to whether illiquid assets are in aggregate providing returns that are too low, that really comes down to the investor's specific utility function, as discussed in Section 1. However, on average we would regard the general level of available illiquidity risk premia of the assets we measure to be moving into the lower end of the range of fair value.

What about private equity?

So far, we have been silent on measuring the IRP in private equity, which may seem odd given its highly illiquid nature. To some degree it is debatable whether private equity should provide an IRP. Whilst investors are required to lock up capital for many years to access private equity – which should be compensated via some level of additional expected return – there is nothing inherent to many assets held within private equity vehicles that demands they be held privately/in illiquid form. Indeed, companies are owned publicly with great ease by many investors, so the additional return from private equity is not exclusively in compensation for illiquidity.

Instead, additional expected returns are likely to compensate for a variety of factors as well as illiquidity, which are difficult to disentangle. A significant element of further expected returns is generated by the investor (or their agent) taking control, and altering the operational or financial structure of the business in some fashion. Clearly, taking control of a business and successfully realising value requires a considerable amount of skill, suggesting much of the extra return is a skill premium. Finally, in some instances, higher expected returns may reflect an inherently higher equity risk premium as a consequence of the increased operational, financial or idiosyncratic risks the investor adopts when investing in some forms of private equity. Therefore, whilst we suspect private equity does provide an IRP to some degree, it is onerous to disentangle it from the control, skill and other risk premia that drive the additional expected returns of private equity over its listed sibling.

3. The drivers of changing IRPs, and what they mean for portfolio strategy

The Willis Towers Watson Illiquidity Risk Premium Index suggests illiquidity risk premia are currently at the low end of fair value. What are the top-down forces acting on IRPs that might explain this move? We make a few observations below:
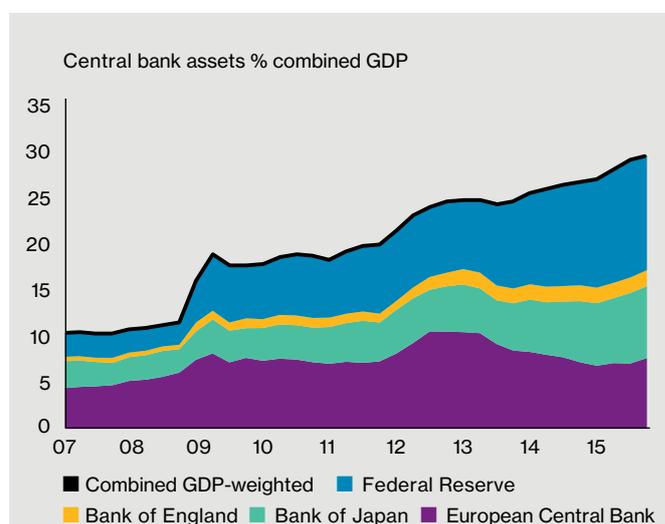
- Section 1 discusses the importance of utility functions in determining the level of risk premium demanded by investors. Over time, the changing composition of the marginal investor base and/or the preferences of these investors could change the form of that utility function and alter the 'fair' level of the IRP.

- Quantitative easing (QE) and other central bank asset purchases have been a crucial driver of shifting marginal utility functions. An explicit policy objective of QE has been to push down on risk premia more generally – IRPs are one example.

- Given the predicted tightening of US policy in coming years, one might expect the pressure on the IRP to be upwards.

- However, taking a lead from bond markets, it is global liquidity conditions under the broadest measure that are most likely to be important here – a simple measure based on the projected size of key developed world central bank balance sheets is presented in *Figure 4.* This indicates that global liquidity conditions have remained easy due to the size of Japanese and eurozone asset purchases, but US dollar liquidity has been declining at the same time. The extent to which this will impact our measure of IRP is uncertain.

- Another factor relevant to expected IRPs is the changing structure of liquidity from the sell side. As we have discussed previously [1], measures of the market-making capacity of key fixed income markets have declined significantly since 2008. Whilst these measures may over state the true decline in underlying liquidity provision (for example, they exclude off-balance sheet trading via 'dark pools'), we do believe the underlying liquidity of capital markets is structurally lower.

- Combined with the rapid rise in assets invested in fixed-income securities, this may leave fixed-income markets less able to meet the demands of investor redemptions should sentiment turn.

[1]. For example, see: https://www.towerswatson.com/en/Insights/IC-Types/Survey-Research-Results/2012/07/Corporate-Bond-Liquidity-Constrained-but-Active-Managers-Can-Still-Excel
Also "Credit market liquidity: changing return of market liquidity and the need for caution" (forthcoming)

In aggregate, the pressures on the global IRP are likely to continue to be lower for some time, reflecting the amalgamation of these forces and conditional on no downside event materialising. Investors are being encouraged to search for yield, and in a low-return world illiquidity risk is potentially an appealing means of generating additional yields. However, we believe there to be a higher chance than normal of a downside event occurring, pushing all risk premia – including IRPs – higher. This macro picture and the evidence presented in Section 2 of a declining IRP lead us to caution: pursuing a risk premium some may be ill-suited to bear, in combination with the picture of reduced structural market-making capacity in more liquid markets, risks 'ending in tears' should sentiment turn.

In part, this underlies our generally cautious outlook on risky assets because, for many risky assets, the amplitude of future price moves to the downside may be greater than normal, which is another way of saying risk-adjusted returns are lower. However, we are still of the view that for many of our clients, who are suited to taking illiquidity risk, certain assets still provide attractive risk premia on a highly selective basis. Moreover, if a negative event does occur and illiquidity risk premia do increase, these clients should stand ready to capture these attractive returns. Often this requires investment at a worrying time, but having a solid process to estimate the IRP – such as that detailed above – eases this decision.

Figure 4. **European Central Bank and Bank of Japan balance sheet expansion is driving easy global liquidity conditions**

Central bank assets % combined GDP

Legend: Combined GDP-weighted, Federal Reserve, Bank of England, Bank of Japan, European Central Bank

Further information

For further information please contact:

Martin Jecks
Martin.Jecks@willistowerswatson.com

About Willis Towers Watson

Willis Towers Watson (NASDAQ: WLTW) is a leading global advisory, broking and solutions company that helps clients around the world turn risk into a path for growth. With roots dating to 1828, Willis Towers Watson has 39,000 employees in more than 120 countries. We design and deliver solutions that manage risk, optimise benefits, cultivate talent, and expand the power of capital to protect and strengthen institutions and individuals. Our unique perspective allows us to see the critical intersections between talent, assets and ideas — the dynamic formula that drives business performance. Together, we unlock potential. Learn more at willistowerswatson.com.

The information in this publication is of general interest and guidance. Action should not be taken on the basis of any article without seeking specific advice.

To unsubscribe, email eu.unsubscribe@towerswatson.com with the publication name as the subject and include your name, title and company address.

willistowerswatson.com

Willis Towers Watson I**l'I'l**I