MILLIMAN REPORT

Setting discount rates under IFRS 17: Getting the job done

Paper 1: An overview of the process

October 2020

Pierre-Edouard Arrouy, FFIA
Charles Boddele, FFIA
Thomas Bulpitt, FIA, CERA
Grzegorz Darkiewicz, PhD, FRM
Russell Ward, FIA
Freek Zandbergen
Sihong Zhu, FIA

Milliman

Table of Contents

Introduction

International Financial Reporting Standard (IFRS) 17 is the new insurance accounting standard issued by the International Accounting Standards Board (IASB). After its initial publication in May 2017 there were a number of areas that led to further discussion within the industry and ultimately resulted in the publication of a revised standard in June 2020. The effective date of the standard is now 1 January 2023.

IFRS 17 requires preparers of accounts to derive discount rates for the valuation of the cash flows associated with their insurance contracts. Paragraph 36 of IFRS 17 outlines the requirement that the discount rates should:

- Reflect the time value of money, the characteristics of the insurance contract cash flows and the liquidity characteristics of the insurance contracts to which they are applied
- Be consistent with observable current market prices (if such prices exist) for assets with cash flows whose characteristics (such as timing, currency and liquidity) are consistent with those of the insurance contracts
- Exclude the impact of any factors that are inherent in the observable market prices but do not affect the cash flows of the insurance contracts

Conceptually setting a discount rate to reflect the time value of money and thereby to allow an expression of amounts to be paid or received at different future times in terms of a single consistent “currency” is relatively straightforward.

In mathematics, complexity theory considers that the behaviour of complex systems arises from the interplay of deeper, relatively simple, rules. Setting discount rates feels like it turns this on its head with a relatively simple concept concealing a tangle of potential underlying complexity.

A good deal has already been written on this topic from summaries of the IFRS 17 regulations themselves to very deep dives into specific aspects of the calculations. Our aim is not to repeat these here in full but to take the perspective of a practitioner who:

1. Already has a good working knowledge of the requirements laid out in the IFRS 17 regulations.
2. Accepts that there will be areas where investigations are needed to determine methodology and processes at a low level. In reality many of these areas will be specific to individual insurers in any case, as they will depend in part on a firm’s exposures, capabilities, objectives and priorities.
3. Would welcome a clear picture of the steps involved in the overall process, how they fit together with simplified examples and a discussion of the key challenges likely to be encountered and practical suggestions as to how they might be overcome.

If our points 1 to 3 above resonate with you, then we hope you will enjoy the introduction in this paper and the more detailed discussions in the later papers. Ultimately, the aim is to set out what we feel is a pragmatic road map to guide the discount rate setting process as follows:

- The second paper will cover three topics in more detail:
    i. Discuss the process around setting the risk-free discount rate.
    ii. Consider how the degree of illiquidity inherent in life insurance liabilities can be assessed and how these liabilities can be linked to suitable reference portfolios of assets to provide a basis for the quantification of any illiquidity premium.
    iii. Describe approaches which can be taken to quantify the illiquidity premium.
- The last paper completes the series with a consideration of the operational aspects of implementation in terms of embedding the required processes into broader business-as-usual (BAU) activities and will also consolidate all the information into a single paper.

Surveying the landscape

A number of steps are required to set the discount rate for any particular block of business. Under IFRS 17 one of two approaches can be taken, the so-called “top-down” and “bottom-up” methods.[1] Our view is that, whilst the selected method may influence the order and content of some of the steps, the steps themselves are relevant to both approaches. As an illustration, Figure 1 provides an overview of the steps required. The order represents where a bottom-up approach is adopted.

FIGURE 1: HIGH-LEVEL PROCESS FOR SETTING IFRS DISCOUNT RATES

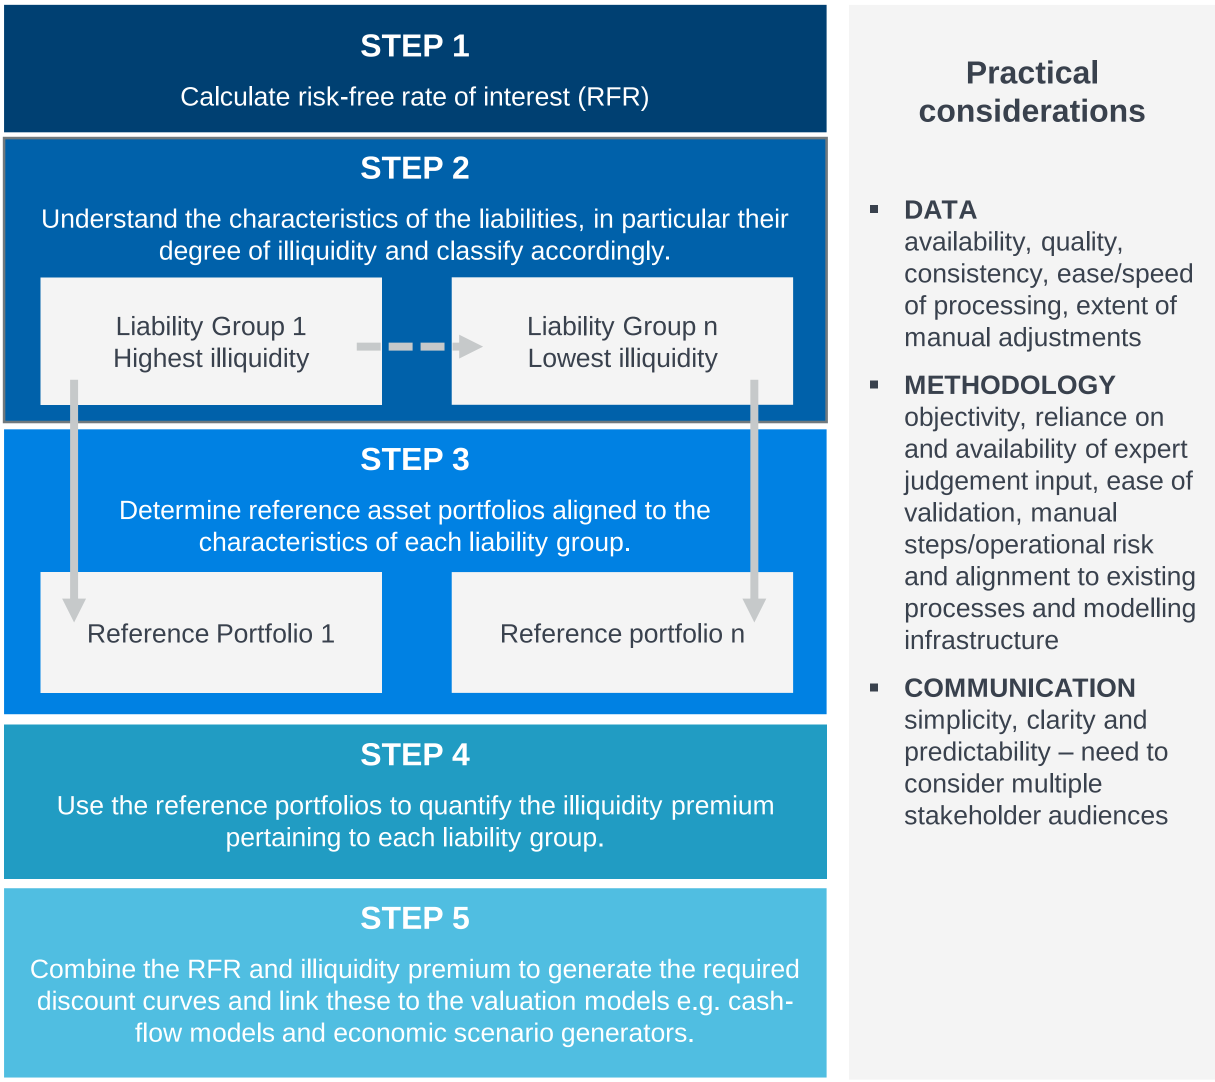

Planning the journey

Taking each step in turn, we now consider the tasks involved and call out particular tensions and challenges that are likely to be encountered.

STEP 1: CALCULATE RISK-FREE RATE OF INTEREST (RFR)

Reference rate

The first step will be to decide on the underlying instruments to be used in constructing the risk-free curve. The prime candidates are government bonds or interest rate swaps with the choice influenced by the range of instruments available by term, the level of market liquidity and the need (or not) to adjust yields to allow for credit

[1] IFRS 17 refers to two main approaches for computing the discounting yield curve. The bottom-up (resp. top-down) approach starts from the risk-free curve (resp. the reference portfolio yield curve) and adds (resp. subtracts) layers in order to reach an appropriate discounting yield curve. For further discussions see https://www.milliman.com/en/insight/ifrs-17-discount-rates.

risk. Even within an asset class, though, there can be variations between instruments with very similar parameters. An example is the slight variation in yields observed between “on-the-run” and “off-the-run” US Treasury bonds.

Develop a set of filtering criteria that can be applied readily to the universe of candidate instruments in order to ensure a consistent selection of benchmark assets over time. Note that this does not imply the same instruments will always be selected but simply that those selected will possess a set of pre-agreed desirable characteristics. For example, divide instruments into set duration buckets and, within each bucket, select the instrument with the top score, which might be driven by factors such as trading volume, level of bid-offer spread, issue date (we might prefer a more recent issue to an older one) etc. Ideally, this process will be set up so as to facilitate automation and reduce time, cost and operational risk.

Interpolation

Many actuarial valuation models evaluate cash flows in monthly time-steps. Such a model will require a risk-free interest rate curve with rates specified at these monthly points. However, even in markets with very rich selections of instruments there will not be a market observable rate aligning neatly to every monthly duration point. Thus, an approach is needed to take a set of rates derived from market data and interpolate them to develop the monthly rates needed for the model.

Extrapolation

Whatever instruments are selected, it is likely that the length of the insurance liabilities will require discount rates that extend beyond the boundaries of liquid trading in the financial markets. This aspect brings into play the toughest challenges associated with developing the risk-free yield curve. In particular:

- How to determine the longest duration at which market data is sufficiently reliable, often referred to as the “last liquid point” (LLP). This point can vary significantly between markets and for different instruments within those markets.
- How to extrapolate the risk-free yield curve from the LLP to an ultimate horizon by which the insurance liabilities are expected to have been extinguished.

There is no single accepted approach to extending the risk-free yield curve and the choices made can have a significant impact on the valuation of some long-term insurance liabilities. We can look across to how other regimes handle this but the ongoing debate over the position of last liquid points and the level of the ultimate forward rate (horizon value) under Solvency II (SII) illustrate the challenge.

FIGURE 2: EXAMPLES OF SII RISK-FREE INTEREST RATE CURVES[2]

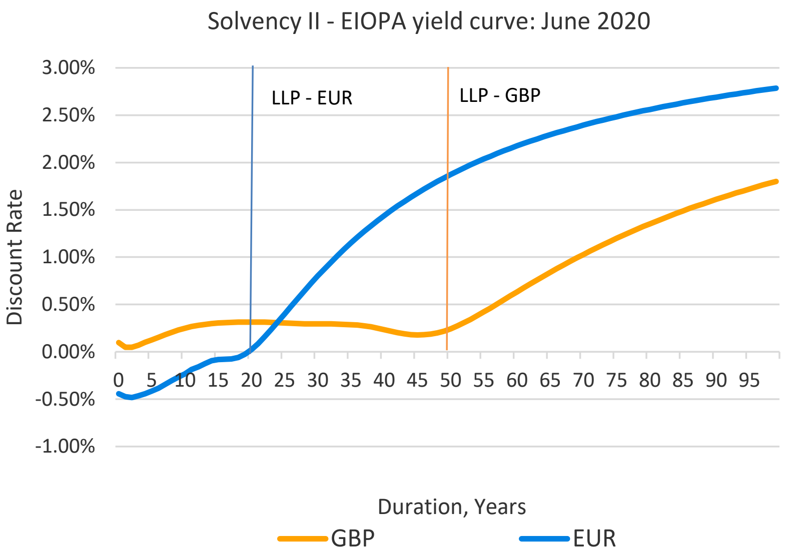

[2] Source: Milliman analysis using data from EIOPA.

Use of a relatively high ultimate forward rate, based on long-term historical data, places a lower value on long-term insurance liabilities and can provide immediate balance sheet relief. However, this approach places significant reliance on a mean reversion mechanism for interest rates whereby they will rise back to historical average levels in the future (see Figure 2). Given that, for a peacetime period, public debt burdens were already high coming into 2020, and that the fiscal measures taken to mitigate the health and economic impacts of the COVID-19 pandemic have increased them further, there is now a very big question mark over the appetite, or indeed feasibility, of raising rates in the foreseeable future in light of the clear consequences for debt sustainability (see Figure 3).

FIGURE 3: UK PUBLIC SECTOR NET DEBT AS A PERCENTAGE OF GDP[3]

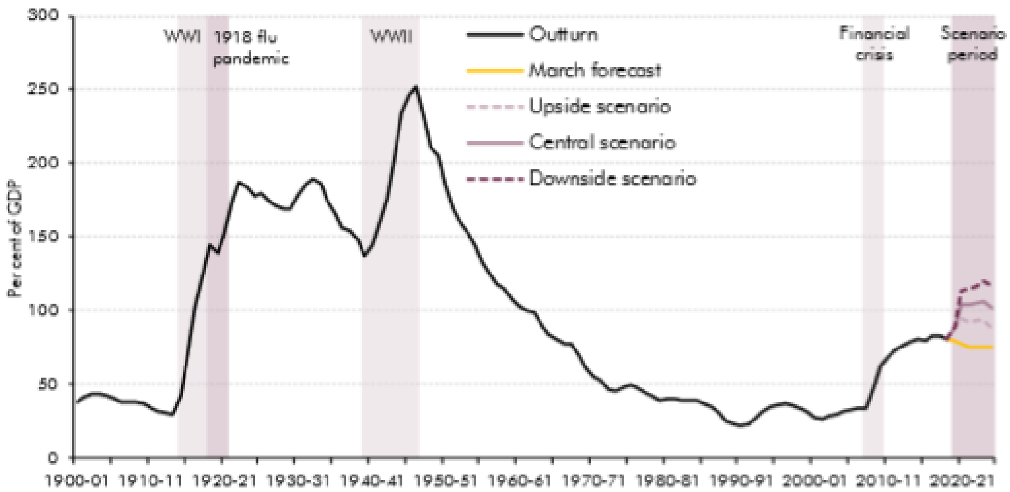

Source: Back of England, ONS, OBR

Should interest rates remain very low for a very long time then use of an elevated ultimate forward rate will place a strain on future earnings, as the value of liabilities will increase at a higher rate than can be earned on risk-free assets. A high discount rate would then initially lead to a high contractual service margin, but over time this will then have to be netted with systematic financial losses on the interest accretion.

For insurers operating in markets where the extrapolated part of the risk-free interest rate curve has a material impact on results, the derivation, ongoing monitoring and adjustment of a suitable horizon rate is perhaps the most challenging and significant area of expert judgement required in setting the risk-free interest rates. We will return to this issue in the second paper of this series.

STEP 2: CLASSIFY THE LIABILITIES

One of the key IFRS 17 principles related to discount rates is that they should reflect the characteristics of the liability cash flows to which they will be applied. "Characteristics" here refers to timing but also, critically, the liquidity of the underlying insurance contracts.

This principle applies to both top-down and bottom-up approaches to determining discount rates. In the top-down approach, it is implicitly reflected in selecting a reference portfolio of assets whose cash flow characteristics, including liquidity, are similar to the liability profile. In the bottom-up approach, the alignment is directly considered in the explicit calculation of the illiquidity premium.

The liquidity of insurance contracts is in principle defined similarly to the liquidity of other financial instruments. In practice it boils down to the ability to convert the asset into cash or terminate the insurance contract upon demand, quickly and without material loss. In the case of assets, a loss can be related, e.g., to a bid-offer spread and in the case of insurance contracts to exit costs.

[3] UK Office for Budget Responsibility (14 July 2020). Fiscal Sustainability Report – July 2020. Retrieved 2 October 2020 from https://obr.uk/fsr/fiscal-sustainability-report-july-2020/.

Looking from a different perspective, more illiquid liability contracts would have more predictable cash flows. This predictability can be viewed through the lens of the level of exposure to underwriting risks and in particular voluntary discontinuance. Here, 100% illiquidity is achieved when cash flows are fixed and insensitive to any stress scenarios on life underwriting risks. For example, pay-out annuities in force have strong illiquidity characteristics. Although they are not free of biometric risks (longevity risk), they are not exposed to surrender risk. For this reason, companies with pay-out annuity liabilities can comfortably invest in less liquid assets to harvest a higher illiquidity premium. The higher illiquidity of the assets is acceptable as the risk of liability cash flows departing significantly from the expected profile is sufficiently low for this type of business. This is the main rationale to increase discount rates with illiquidity premia corresponding to the degree of illiquidity of the liabilities.

In practice, different features may influence the liquidity of an insurance contract. These features can be of a contractual nature (for example: are surrenders allowed or not? what are surrender penalties?) or of a more general nature, which defines whether discontinuing an insurance contract may be attractive to clients or not. Examples of such features could include the remaining term of a contract, fear of losing initial underwriting, enjoying a high guaranteed interest rate in the low interest rate environment, the possibility of losing certain tax advantages and many others.

Although there could be many features defining exposure of the contract to unexpected lapses, in reality—following publication of the "Exposure Draft of the Proposed International Actuarial Note (IAN) 100 on Application of IFRS 17 Insurance Contracts"—liquidity of insurance contracts is driven by two values:

- **Exit value:** The higher the exit value, the higher the liquidity of the contract (contracts with higher exit value are more attractive to lapse than contracts with lower exit values). The exit value is impacted by contractual definition of surrender value and all explicit (contractual) and implicit costs which need to be incurred by the policyholder in case of discontinuation of an insurance contract. For example, losing tax advantage can be considered to be an implicit cost upon surrender.
- **Inherent value**: If a contract’s pricing or construction is such that there is negligible or no inherent value (or value build-up), or it is lower than the exit value, then it is likely to be considered liquid. Despite being able to price the inherent value, in reality it is much more subjective than the exit value, as it reflects how policyholders perceive the value of holding the contracts. Different features can impact inherent values such as embedded guarantees, the health of the policyholder, who might be afraid to undergo another underwriting process, the state of the economy etc.

Although this distinction is quite intuitive, it does not provide any quantification of the level of illiquidity. We will discuss two possible approaches.

Homogeneous liquidity buckets

One possible approach to achieve that is to gather insurance contracts into broadly homogeneous liquidity buckets and select reference asset portfolios with liquidity characteristics similar to each liability bucket. This may not always be straightforward, as different illiquidity characteristics might not be directly comparable—how should different features be weighted or scored? An additional practical challenge could arise if the level of illiquidity strongly depends not only on contract features, but also on time to maturity, adding a further dimension to the illiquidity classification and making it likely that contracts will move between illiquidity buckets over time. Last but not least, it is very challenging to decide which assets should be used in a reference portfolio to reflect the level of illiquidity of liabilities.

Methods based on stress testing

In the "Consultation Paper on the Opinion on the 2020 Review of Solvency II," the European Insurance and Occupational Pensions Authority (EIOPA) discusses the possible changes in Solvency II related to the Volatility Adjustment. In particular, EIOPA investigates approaches to align a reference asset mix to the company-specific duration of liabilities and to capture liquidity characteristics of insurance contracts. As the forthcoming changes are practically in line with IFRS 17 requirements, similar ideas could be considered also in the context of setting the illiquidity premium for IFRS 17.

What is notable is that EIOPA prefers methods based on stress testing over creating homogeneous liquidity buckets. EIOPA argues that the method based on stress testing implicitly takes into account all contract features and at the same time is much simpler to implement.

The EIOPA approach is interesting and the prospect of using a single quantitative measure to assess illiquidity is certainly attractive. Nevertheless, some challenges remain and we intend to explore the relative merits of this and other approaches in Paper 2 of this series.

STEP 3: DETERMINE REFERENCE ASSET PORTFOLIOS

As there is no market information on the illiquidity premia of liabilities, a link needs to be made towards an asset portfolio to have observable data.

To determine the illiquidity premium to apply for liability bucket *k*, the general form is as follows:

*Liability Illiquidity premium (k) = application ratio (k) * asset portfolio illiquidity premium (k)*

Where:

- Application ratio (k) is the proportion of any illiquidity premium to be applied to liability bucket *k*.
- Reference asset portfolio illiquidity premium (k) is the illiquidity premium assessed based on the backing assets (own portfolio or reference portfolio).

In general, there are two main approaches to select an asset portfolio to start from:

1. Use of an *actual portfolio* of the insurer's own assets which back the liabilities on the balance sheet (similar to the matching adjustment approach under Solvency II).
2. Use of a *theoretical reference portfolio*, selected to provide a close match to the liabilities but unrelated to the actual assets held. The use of a theoretical portfolio is adopted, though in a broader sense, to underpin the Volatility Adjustment approach under Solvency II.

When making a choice between using an insurer's own portfolio versus a theoretical reference portfolio, there are a number of factors to consider carefully, below we elaborate on three notable ones.

First, it is critical to understand the level of granularity at which the asset portfolio has been segmented. In an ideal world we would have a separate asset portfolio for each liability bucket. The application ratio would then simply be 1 in every case. In reality, this is likely impractical, in particular where the insurer's own portfolio is used, as the business is unlikely to hypothecate assets to the level required. This practical constraint then raises a risk of double-counting the haircut on the illiquidity premium when using the own portfolio. For example, the asset portfolio linked to a product which can lapse without a penalty will be invested in relatively liquid instruments from a liquidity risk management perspective. Consequently, the liquidity premia will be on the low end. If application ratios were derived following the stress test-based classification noted in the previous section, the application ratio will also be on the low end. The combination of the two would lead to a double-whammy reduction on the level of the liquidity premia to be applied in the valuation.

Second, using a theoretical reference portfolio can potentially lead to mismatches and volatility in the profit and loss (P&L) statement due to differences between assets and liabilities, in particular in the case of significant differences in the asset allocation between the theoretical reference portfolio and the actual asset portfolio. The current Solvency II regime suffers from a similar issue where the Volatility Adjustment is based on a theoretical reference portfolio and leads to high volatility in the levels of the Own Funds, known as the overshooting and undershooting effect. In particular, in times of spread widening, several insurers have experienced a larger offset in their liability valuations compared to the actual losses in their asset portfolios.

Third, one should think about management incentives in the case where an insurer's own portfolio is used. Without boundaries it can potentially trigger the wrong incentives to invest in particular assets, which will increase the illiquidity premia to be applied in the discounting of liabilities. For example, there may be a temptation to invest in assets with which the insurer has insufficient origination and management experience. In any case, the illiquidity of the assets used should remain aligned to that of the underlying liabilities and a regular test might be applied to verify this.[4] Note that this is the most important reason for Solvency II to stick with the use of a theoretical reference portfolio in determining the Volatility Adjustment.

[4] Note that maintaining a reasonable correspondence between the liquidity characteristics of the liabilities and those of the assets used to derive the discount rate for those liabilities is a requirement under IFRS 17.

STEP 4: QUANTIFY THE ILLIQUIDITY PREMIA TO BE APPLIED

Liquidity refers to the degree of ease with which an asset can be converted into cash. It can also correspond to the ability to sell large quantities of the asset without any adverse repercussion on the price. Therefore, an illiquidity premium is a premium, demanded by investors, to compensate them for the risk that they may be forced to hold an asset for a certain period and/or accept a significant loss of value in the event of an early sale.

The expected return of an asset over the risk-free rate can then be split into two parts:

1. Remuneration of the credit risk, often split into expected cost of default and credit risk premium.
2. Remuneration of the illiquidity of the asset.

The illiquidity premium is then obtained by a direct calculation or by difference as soon as the credit risk is identified.

The literature has addressed several methods to compute the illiquidity premium. The most popular ones are the following:

- **Historical probabilities of default:** Transition matrix and recovery rates are the main inputs to derive the cost of default (and possibly the cost of downgrade). It is a well-known method as it’s used for the Volatility Adjustment and Matching Adjustment under Solvency II.
- **Credit default swap (CDS) negative basis:** The difference between the CDS premium and the corporate bond spread with the same characteristics forms a negative basis. The latter can be interpreted as a direct quantification of the illiquidity premium.
- **Covered bonds:** It is supposed that such bonds do not embed the risk of default, therefore the difference between their yield and the risk-free rate can be considered as a direct quantification of the illiquidity premium.
- **Proxy models:** The most common proxy is for the illiquidity premium to be a linear function of the spread. Such an approach was introduced under Quantitative Impact Study (QIS) 5[5] and has been refined by EIOPA for the computation of the Volatility Adjustment. More recently, the "Consultation Paper on the Opinion on the 2020 Review of Solvency II" suggested alternative proxies to compute the illiquidity premium.
- **Regression models:** The illiquidity premium is computed with a function of different explicative variables.
- **Structural models:** Such models use option-pricing techniques on a firm’s own asset value to calculate both expected and unexpected credit risks. Expected credit risk refers to the historically observed probability of default while unexpected credit risk arises from the real-world credit risk premium. Subtracting the sum of expected and unexpected credit risks from the actual market spread then leaves the illiquidity premium. The Merton model and the Black-Cox model constitute two famous structural models.

In the second paper of this series, all these models will be analysed from different angles. In fact, some models can be interesting because of their very precise computations of the illiquidity premium, but the market can lack in transparency and liquidity so that the data necessary is unreliable. Some approaches can be easily used but rely too much on expert judgement and are very subjective, while others are very theoretical and their application can require the use of specific data at granular levels or simplifications which might undermine the robustness of the results.

The choice of the approach can have an important impact on the volatility of the illiquidity premium driving procyclical effects, and it could also depend on the kind of liability that the assets of the reference portfolio are backing. Besides, if most of the above-mentioned approaches establish a clear link between corporate spreads and their liquidity, one should also address the measurement of liquidity for assets other than bonds. It is then necessary to study what is already used in different sectors (finance, insurance etc.) and what can be done for every specific asset.

[5] See Appendix 2 of QIS 5 Technical Specification: Risk-free interest rates, at https://www.thecroforum.org/wp-content/uploads/2013/05/cfo-forum-cro-forum-paper-risk-free-rates_en1-1.pdf.

STEP 5: OPERATIONALISE THE PROCESS

Although this is the final stage of the process, as Figure 1 above neatly demonstrates, the operational aspects of the discount curve derivation need consideration throughout. A key factor for many insurers will be the speed at which they can produce the relevant discount curves in order to meet working day timetables.

Factors that will influence this efficiency will be largely focussed on:

- Data
- Methodology
- Communication

We will consider each of these in turn.

Data

By its very definition, an illiquidity premium is derived from the spread above risk-free that is available on assets for which there is not a fully liquid market. Market data is needed in order to robustly determine a market value of assets from which a market-implied spread can be calculated. This presents a conundrum as, in order to define an efficient process, the ability to quickly access up-to-date market data is important. This is likely to influence the insurer’s decision over the choice of assets used in the calculation.

As an example, where a portfolio of equity release mortgages is considered to be representative of the liquidity characteristics of the portfolios of insurance contracts, the valuation of the equity release mortgages relies to some extent on infrequent property survey data as well as mark-to-model techniques. Although this valuation may be required as part of the overall year-end valuation process (where the insurer actually holds the assets), it is unlikely to be available early in the process and, in any case, only provides the starting point from which to derive an illiquidity premium. Where the insurer does not hold the assets, it may be in a position where there is limited, or no, publicly available data, to use as a starting point.

Although determining the illiquidity premium using a reference portfolio of equity release mortgages may be desirable from a capital efficiency perspective (depending on the insurer's overall preferences and objectives), it may not be desirable from an operational efficiency perspective.

Where an insurer considers using bottom-up techniques to determine an illiquidity premium, insurers may want to factor in the availability of data such as covered bond spreads or credit default swap rates. Some of these data items may be more available than data on equity release mortgages but potentially not at the required durations. The manual workarounds that may be required to compensate for the incomplete data may make the approach prohibitive from an operational perspective.

Methodology

Secondly, the choices made when setting the methodology have the potential to affect the efficiency of the process. Insurers will want to consider the trade-off between technical accuracy and efficiency. It will be key for insurers to look for areas where estimates can be used and where simplifying assumptions can reduce the workload, accepting that applying such assumptions introduces a reliance on expert judgement.

A potential area where this might be a key consideration would be the granularity at which the calculation is performed. IFRS 17 requires that each group of insurance contracts is considered in isolation. However, for many insurers (and in particular multiline insurers), the complexity of the process to derive discount rate curves would become insurmountable if a separate calculation, and potentially reference portfolio, was performed for every group.

Insurers will want to look for portfolios of contracts that are sufficiently similar such that fewer calculations can be performed. Indeed, some insurers may consider whether a single calculation would be sufficient for the full liability valuation.

Leveraging existing tools that may have been developed for other reporting purposes would be a simple way to increase the efficiency of the process. For example, under Solvency II, a lot of work has already been performed in determining the Volatility Adjustment and whole frameworks have been built around the calculation of the Matching Adjustment. Where an insurer already adopts either of these measures, they may present a sensible starting point from an operational perspective.

Other factors to consider when setting the methodology would include the level at which expert judgement is used and therefore the balance of simplicity and objectivity. The techniques chosen will ultimately need to be validated and therefore it would be worth thinking about how easy that would be. Clearly, the more subjective the methodology, the harder this is.

Communication

Finally, and this ties in with complexity, the ease of communication of the methodology should be considered. Given that IFRS 17 is an accounting standard and used in financial disclosures, it will be essential for users of those financial statements to be able to understand, at least to some extent, the process followed to determine the discount rate curves.

Users of the financial statements for many insurers will include the board of directors, market analysts, shareholders, credit rating agencies and auditors. All of these stakeholders will be looking at the statements from a different perspective and therefore it will be important to consider the needs of each.

For example, the board may be more interested in predictability of results than comparability across insurers, but the latter may be more attractive from the perspective of the credit rating agencies. Auditors will also need to be able to clearly understand the methodology and make an assessment of the appropriateness of the approach.

When setting methodology it will be important to consider how easy it will be to communicate the ultimate methodology to all stakeholders.

Summary

In conclusion, there are many interesting and challenging areas to address along the road to establishing appropriate discount rates under IFRS 17. We hope you have found this first instalment interesting and will join us for the remainder of the journey as we issue subsequent papers in this series.

Milliman

CONTACT

Pierre-Edouard Arrouay
pierre-edouard.arrouay@milliman.com

Charles Boddele
charles.boddele@milliman.com

Grzegorz Darkiewicz
grzegorz.darkiewicz@milliman.com

Russell Ward
russell.ward@milliman.com

Freek Zandbergen
freek.zandbergen@milliman.com

Sihong Zhu
sihong.zhu@milliman.com