Craig Turnbull
Craig.Turbull@barrhibb.com

Liquidity premiums and contingent liabilities

The liquidity premium – the concept that illiquid assets have lower prices than equivalent liquid ones – has recently emerged from relative obscurity to become a major issue in financial reporting (MCEV) and regulatory capital assessment (Solvency II) in European insurance. Most research on this topic has focused on how to observe the level of liquidity premiums embedded in the market prices of assets such as corporate bonds. Less work has been done on how to apply these liquidity premium estimates to the market-consistent valuation of (complex) insurance liabilities.

This note explores how market-consistent liability valuation methodology can be adjusted to allow for liquidity premiums, and the implications this has for the behaviour of the valuations of different types of liabilities. The process of incorporating an allowance for liquidity premium into the market-consistent valuation of insurance liabilities can be considered in three stages:

1. Estimation of the liquidity premiums embedded in asset prices.
2. Estimation of the illiquidity of the insurance liabilities, so as to derive the liquidity premium applicable to the liability under consideration.
3. Application of the liquidity premium estimate (stage 2) in the market-consistent valuation of the liability.

We first provide a re-cap of the work done to date on the first two stages before exploring the final stage in greater detail.

Estimating the liquidity premium embedded in asset prices

Considerable work has been done in this area and a number of methodologies have emerged for the estimation of the liquidity premiums that can be found in asset classes such as corporate bonds. These approaches derive liquidity premium estimates by comparing the prices of two assets that are equivalent except for liquidity. The methods differ by the choice of assets that are used for this comparison, and whether these assets have market prices that can be directly observed or need to be estimated by a model. Barrie & Hibbert's Research Report "Summary of Liquidity Premium Estimation Methods" provides a fuller discussion of these methodologies.

Measuring the relative liquidity of Insurance liabilities

The next stage of the process involves determining the liquidity of the liability. This isn't necessarily a binary decision, although interestingly, the draft QIS5 specification has taken a simplifying approach by assuming that liability liquidity can be grouped into three buckets:

1. Liabilities maturing in less than one year (assumed to be perfectly liquid, so get 0% of the liquidity premium);
2. Immediate annuities (assumed to be perfectly illiquid, so get 100% of the liquidity premium);
3. And everything else (gets 50% of the liquidity premium).

More recently there has been a joint proposal from the CRO Forum, CFO Forum and the CEA that an additional category of business be introduced that gets 75% of the liquidity premium.

Barrie & Hibbert published research in December 2009 that explored a more granular and sophisticated modelling approach to measuring liability liquidity (see Research Report "*A Measure of Liquidity of Liabilities*" for a fuller discussion). The central idea of the paper is that the liquidity needs of a liability can be assessed by measuring how often a matching asset portfolio needs to be rebalanced due to unexpected changes in the liability cash flow outcomes (e.g. due to unexpectedly high mortality rates in an immediate annuity book).

Relative to the QIS5 method, it is a modelling approach that inevitably involves more model and parameter choices, but offers a less arbitrary and more granular approach to liability liquidity assessment.

The liquidity needs of a liability can be assessed by measuring how often a matching asset portfolio needs to be rebalanced due to unexpected changes in the liability cash flow outcomes.

Market-consistent valuation methodology for illiquid liabilities

Once armed with an estimate of the applicable liability liquidity premium, how should it be incorporated into the valuation of insurance liabilities?

For liabilities whose form is sufficiently simple that market-consistent valuation only requires the discounting of expected liability cash flows using a risk-free curve, the adjustment is trivial: add the liquidity premium to the risk-free yield curve and discount those expected cash flows using the adjusted yield curve.

What about more complex liabilities such as those that require market-consistent simulations to value them?

Possible areas of confusion

In QIS5, the approach is again quite simple: the risk-free yield curve used for the valuation is directly adjusted for the addition of the liquidity premium. The natural next step is for risk-neutral valuation to then proceed as usual – i.e. this adjusted yield curve is used both in the risk-neutral stochastic projection of the assets and in the discounting of the resultant contingent liability cash flows.

This implementation sounds simple, but one area of potential confusion and ambiguity that can arise here is over whether the stochastic model's volatility parameters should be re-parameterised so that the calibration 'undoes' the impact of the liquidity premium adjustment on the model's fit to option prices. Where liquidity premiums have been used to date (e.g. in MCEV and in UK regulatory capital assessment), the standard practice has been to *not* re-calibrate the model's volatility parameters and we expect this to be the case in QIS5 and SII implementation. This is broadly intuitive in the sense that if the intention was to keep the valuation basis consistent with the pre-liability-premium basis, there would be no reason to change it in the first place. Put another way, a swaps + 20bps calibration isn't consistent with swap prices, so why should it have to be consistent with swaption prices? We are using option prices to derive market-implied volatility levels, and there is no obvious reason for this derivation to assume the market is using the same liquidity premium adjustment in its derivative valuations as QIS5. But it should be noted that retaining a fit to option-implied volatilities instead of option prices arguably contravenes TP.1.263 of draft QIS5, although it is consistent with TP.1.320.

The approach of merely adjusting the starting yield curve has the benefit of apparent simplicity, but it will produce a counter-intuitive effect: the inclusion of the liquidity premium will reduce the value of some forms of contingent liabilities (as expected), but will result in an *increase* in the valuation of some contingent liabilities (options and guarantees that increase in value when interest rates increase). The following example illustrates this effect.

Example: standard fixed deferred annuity contract

This case is straightforward: the market-consistent valuation of a fixed deferred annuity is calculated by discounting the expected liability cash flows using the risk-free yield curve (adjusted for the inclusion of the applicable liquidity premium).

Now let's suppose that this contract also provides the policyholder with the option to exchange the annuity income for a lump sum at the annuity vesting date. This form of policyholder guarantee will typically be valued using market-consistent simulations. Note that the guarantee will be worth more to the policyholder when interest rates are very high: in such scenarios, the annuitant will be able to secure a higher annuity income by exercising his lump sum option and re-investing it at the market annuity rate.

So, the market-consistent valuation of this option will increase if we include a liquidity premium in the starting yield curve in the way prescribed by QIS5. Meanwhile, a policyholder option that is more valuable when interest rates fall (e.g. an option to exchange a lump sum for annuity income at a pre-determined rate) will be reduced in value by the inclusion of the liquidity premium. Are these opposite effects arbitrary or logical? *If the rationale for the inclusion of liquidity premiums in liability valuations is that illiquid things should cost less, does it make any sense for the value of a liability to be increased by an allowance for its illiquidity?*

This counter-intuitive result is arguably indicative of the tenuous link between the liquidity premium earned on assets such as corporate bonds, and those earned by more complex contingent assets such as non-linear derivatives like options and swaptions (which most resemble the contingent liabilities under consideration). Whilst the link between the liquidity premium that can be earned by investing in illiquid assets such as corporate bonds and the valuation of a liability that can be replicated through a buy-and-hold strategy that uses those illiquid bonds is intuitive, how this impacts on the valuation of options (that are replicated through dynamic strategies that necessitate the use of liquid assets) is far less obvious.

In the QIS5 methodology, the impact on the contingent liability valuation is defined by the adding the liquidity premium to the underlying interest rate used in the derivative valuation. As we have discussed, sometimes this will result in the valuation going up, and sometimes its will result in the valuation going down. These valuation impacts are hard to interpret as meaningful estimates of the impact of illiquidity on the option value.

However, the behaviour of the valuation of the deferred annuity's lump sum option may be more intuitive when considered in aggregate with the behaviour of the underlying deferred annuity. The total liability value will be equal to the sum of the value of the standard fixed deferred annuity and the lump sum option. When a liquidity premium is added to the valuation in the QIS5 way, the value of the standard component of the deferred annuity will fall, whilst the value of the option will increase. Viewed from this perspective, the increase in the option value could be considered as a 'claw-back' of the discount produced for the standard component of the deferred annuity by the inclusion of the liquidity premium. This 'claw-back' might be viewed as reasonable since the inclusion of the option makes the deferred annuity more liquid for the policyholder. But there are more direct and transparent ways of reducing the liquidity premium impact on these types of contracts if that is the intention.

There is no rationale for the liquidity premium to be wrapped up in the calibration of stochastic interest rate models.

What's the answer?

What is the right way to allow for illiquidity in the market-consistent valuation of interest rate-contingent liabilities such as the lump sum option for the deferred annuity contract discussed above?

The process could be simplified if the impact of the liquidity premium was separated from the risk-neutral simulation modelling: there is no rationale for the liquidity premium to be wrapped up in the calibration of stochastic interest rate models, and it will further complicate the process of calibrating to interest rate swaptions and validating their quality of fit (particularly when the addition of the liquidity premium will mean models that have been fitted to implied volatilities as per standard practice will now no longer reproduce market prices). It will also require multiple calibrations, scenario sets and ALM model runs to be produced for no particular reason.

A simpler solution would be for market-consistent simulations to be used to value liabilities prior to any adjustment for liquidity. The liquidity premium adjustment could then be made by simply multiplying the pre-liquidity market-consistent value by a factor that is based on the applicable liquidity premium and the duration of the liability. This would not address the question of whether adjustments to contingent liability valuation for the effects of illiquidity should always be positive, always be negative, or either depending on the direction of its rate sensitivity. But it would simplify

the valuation process by reducing the economic scenario calibrations, validations and scenario sets required for market-consistent calibration of contingent liabilities.

An additional complication to consider is whether certain liabilities, which are classified as 50% liquidity premium liabilities under QIS 5, should be reduced by the liquidity premium. An example of such a liability type might be a typical UK with-profits policy which is often considered to consist of an asset share and a value of guarantees. The asset share has many of the features of a unit linked contract and it is unclear why the asset share value itself should be adjusted for illiquidity.

In summary

The use of liquidity premiums to lower the market-consistent valuation of insurance liabilities has become a high-profile issue in Solvency II. This has led to a considerable amount of work being done by the industry (and Barrie & Hibbert) on the estimation of liquidity premiums embedded in market asset prices. Significantly less work has been done on measuring liability liquidity, and how market-consistent liability valuation methodologies should be adjusted to allow for the estimated liquidity premium.

This latter topic is particularly complicated when considering the valuation of contingent liabilities such as options and guarantees. In these cases, the directional interest rate sensitivity of the liability may be opposite to that of fixed liabilities (e.g. a lump sum option attached to a deferred annuity contract will increase in value when interest rates are assumed to be higher). The proposed QIS5 methodology of simply adjusting the starting yield curve by the size of the liquidity premium will therefore produce counter-intuitive results for liabilities such as these – i.e. the intended lowering of liability values to allow for their illiquidity will not occur, and instead the value of these liabilities will be increased by the inclusion of the liquidity premium. For some liability types even applying this adjustment needs careful consideration.

This effect highlights the dubiety of incorporating an allowance for liquidity into contingent liability valuation by changing the starting yield curve that the stochastic interest rate model is calibrated to. It would be simpler and more transparent to use risk-neutral simulations to value the liability prior to any adjustment for illiquidity. The illiquidity adjustment could then be made by applying a factor to the pre-liquidity-adjusted value, based on the applicable liability premium and the duration of the liability.

Disclaimer

The information in this document is believed to be correct but cannot be guaranteed. All opinions and estimates included in this document constitute our judgment as of the date indicated and are subject to change without notice. Any opinions expressed do not constitute any form of advice (including legal, tax and/or investment advice).

This document is intended for information purposes only and is not intended as an offer or recommendation to buy or sell securities. The Barrie & Hibbert group excludes all liability howsoever arising (other than liability which may not be limited or excluded at law) to any party for any loss resulting from any action taken as a result of the information provided in this document. The Barrie & Hibbert group, its clients and officers may have a position or engage in transactions in any of the securities mentioned.

Barrie & Hibbert Inc. and Barrie & Hibbert Asia Limited (company number 1240846) are both wholly owned subsidiaries of Barrie & Hibbert Limited.