Research Note

A Measure of the Liquidity of Insurance Liabilities

Delme Pritchard, Craig Turnbull
Version 1.0
December 2009

barrie+hibbert

Contents

1 Overview

Rewind to December 2008 when economies around the globe were recoiling from the collapse of Lehman Brothers and market conditions resembled an insurer's particularly comprehensive 'worst case' stress test[1]. It is difficult to envisage how the timing could have been worse for the group of 20 European insurance companies forming the CFO forum who were due to report their embedded value on a market consistent basis for the first time as of the year-end. On the 19 December, perhaps not surprisingly, the CFO forum issued a statement to the effect that the European Insurance CFO Forum Market Consistent Embedded Value Principles[2] (MCEV Principles) would be reviewed given the turbulent market conditions and that, in particular:

> " ...[the] areas under review include implied volatilities, the cost of non-hedgeable risks, the use of swap rates as a proxy for risk-free rates and the effect of liquidity premia."

In the event, only a handful of companies reported their year-end 2008 embedded value in a manner consistent with the MCEV Principles and even then, in all cases, practices deviated from the existing guidance. This note focuses on the latter of these areas identified by the CFO Forum – the application of a liquidity premium, in addition to the risk-free rate, for the valuation of insurance liabilities.

Without guidelines, or even established market practice it was perhaps inevitable that discordant approaches would emerge resulting in widely varying estimates, justified in different ways and with different adjustments for similar business lines. Looking at these in turn:

- Absolute value of the liquidity premium. The magnitude of the liquidity premium, where used, varied from 50bps to 300bps. A difference, when applied to the calculation of the present value of a 20 year cash flow, of over 30%.
- Justification for the incorporation of a liquidity premium. Two fundamentally different camps emerged for the justification of using a liquidity premium. The first argued that the liquidity premium should be based on an estimate of that embedded in the actual assets held to back liabilities on the basis that they would be held to redemption and hence the LP would be earned 'for free'. The second that it should be based on the theoretical portfolio of replicating assets and not the actual backing assets. We discuss this in more detail later.
- Application by business line. Some firms only applied the liquidity premium to the most illiquid liabilities (e.g. fixed annuity business), while others applied it more widely, in some cases adjusting it downwards to allow for the relative illiquidity of the liabilities.

The CFO Forum has recently (October 2009) released an update to the Principles to reflect the inclusion of a liquidity premium:

"Where the liabilities are not liquid the reference rate should be the swap yield curve with the inclusion of a liquidity premium, where appropriate… In evaluating the appropriateness of the inclusion of a liquidity premium (where liabilities are not liquid) consideration may be given to regulatory restrictions, internal constraints or investment policies which may limit the ability of a company to access the liquidity premium."

We note that the above may be interpreted as suggesting that the LP may depend on the actual assets backing the liabilities, given the indication that firms may want to factor in the impact of actual investment restrictions. Alternatively, given that the goal of the Principles is to set a framework for reporting the difference between the market value of assets and the third-party transfer costs for the liabilities, it seems more likely that the reference is in regard to the constraints that may apply to a typical third party acquirer of the liabilities.

Also within the updated MCEV Principles, with respect to the assessment of the liquidity of liabilities, it is stated:

> "A liability is liquid if the liability cash flows are not reasonably predictable."

Despite this clear statement of support, without a standard methodology for its estimation or application, there may be a danger that the inclusion of an LP in the valuation of liabilities becomes considered by external users of insurance companies' financial reports as a subjective 'fudge' factor, effectively reducing the comparability of results across companies and ultimately the credibility of the approach. Indeed, in a

[1] For a summary see "B+H Insights: After the Storm: Market-consistent measures of cost and solvency in 2009.", Craig Turnbull, Jan 2009.

[2] Copyright© Stichting CFO Forum Foundation 2008

Solvency II consultation paper by CEIOPS[3] where they ask the question on which liabilities should be considered sufficiently illiquid and how the LP should be quantified, they state there is more support for the exclusion of the liquidity premium for Solvency II and that there is currently "***no generally acknowledged method which will derive the illiquidity premium in a prudent, reliable and objective way.***" Reading between the lines, this suggests the main objections are based more on the practical implementation side rather than from a theoretical standpoint. Indeed, at the time of going to press, CEIOPS in their cover letter to the European Commission[4] seem to have softened their stance with regards to the inclusion of a liquidity premium in the risk-free rate, acknowledging that further work is necessary in this area:

> "Although the choice of the reference rate is a technical issue, CEIOPS recognises that, due to the level of quantitative impact on the level of technical provisions, this has strong political implications. This applies, in particular, to the inclusion of an illiquidity premium in the risk-free rate. Hence, further work would have to be carried out with a clear concept and mandate in light of the framework contained in the advice. CEIOPS is prepared to take the lead in this area and continue to involve all the relevant stakeholders in a transparent process."

To address these issues, objective and transparent methodologies need to emerge for both the estimation and application of the LP. Considerable research and work has been invested in for former[5] but, as far as we can tell, very little if any work has been done to address the latter.

The aim of this note is to add to the discussion of how an objective, quantitative measure for the (il)liquidity of specific classes of insurance liabilities might be developed. The theoretical framework for analysing the liquidity of liabilities already exists and we provide a brief summary of the literature in this area in the next section. While these cover the theoretical issues, there is little or no comment on the difficult issue of how an insurer can practically apply the theory to assess the illiquidity of their liabilities. In the following section, Section3, we tentatively propose a practical method to help with this assessment. This is followed with some simple worked examples in Section 4 which we conclude with a discussion of possible extensions of the framework. We conclude with a summary in Section 5

.

2 Background

There seems to be a dearth of literature on the measurement of the liquidity of insurance liabilities. Having said this, the IAA's paper on the Measurement of Liabilities for Insurance Contracts[6] provides a good discussion on the inclusion of a liquidity premium in the valuation of insurance liabilities. It also references the minutes from the IASB's Insurance Working Group[7], which examines the treatment of a liquidity premium in the valuation of insurance liabilities. While both of these documents provide a high level discussion of the issues, neither provide detailed guidance on how one would actually estimate the LP or estimate the proportion of the LP to be used with reference to the underlying liabilities.

Liquidity is usually defined in relation to assets rather than liabilities where a typical definition may be "The ease and certainty with which an asset can be converted into cash". With reference to the liabilities one can think of the liquidity of the liabilities with respect to the position of the policyholder and where liquidity can be thought of as the ability to sell the liability (or financial instrument) in a short period of time. In the extreme case, a liability which has known (i.e. entirely predictable) future cash flows could be considered to be illiquid (to the policyholder). For example, a plain vanilla 'annuity certain' would meet these criteria since the policyholder is denied access to liquidity – there is no second-hand market and there is no surrender value. The greater the uncertainty in the cash flows, the more liquid the liability could be considered to be. This uncertainty in the liability cash flows can come from two distinct sources:

- Policyholder actions e.g. the option of policyholders to lapse their policies.

[3] CEIOPS consultant paper no. 40.

[4] http://www.ceiops.eu/media/files/publications/submissionstotheec/CEIOPS-Cover-Letter-on-final-advice-to-EC-SII-1st-and-2nd-wave.pdf

[5] For a good summary see: 'Summary of Liquidity Premium Estimation Methods', John Hibbert et al., October 2009.

[6] International Actuarial Association (IAA). (2009). Measurement of liabilities for insurance contracts: Current estimates and risk margins. An international actuarial research paper prepared by the ad hoc risk margin working group.

[7] International Accounting Standards Board (IASB). Discount rates and day one losses. Information for observers, Insurance Working Group , Nov 2008.

- Variances in demographic assumptions from expected e.g. experienced mortality turns out to be different from expected.

From the policyholders' point of view, only the first of the above will impact their assessment of the liquidity of the insurance product. However, both of the above will impact the liquidity requirements for the firm writing the business. The IASB assume that allowance for the latter variance is captured in a separate risk margin, and that only the former is explicitly allowed for in assessing the premium for liquidity.

The following simple observation provides a strong case for why a liquidity premium should be allowed for in a market consistent valuation of liabilities.

> If the liability cash flows can be matched exactly by a (hypothetical) portfolio of assets, then any illiquidity premium embedded in the asset valuation would automatically be included in the valuation o f the liabilities.

By matching here we mean exact replication of cash flows i.e. in every future scenario, in every time step the cash flows from the matching asset portfolio are identical to the liability cash flows. So, at one extreme (absolute matching), there is a strong case for the use of a liquidity premium – so the argument against using one in other cases may, as noted above, be one based on the practical difficulty of estimating and applying it, rather than from a theoretical standpoint.

To be consistent with the above reasoning, the liquidity premium estimate should be based on a theoretical matching portfolio rather than any actual portfolio held to support the liabilities. In further support of this approach, if the actual backing assets were to be used to estimate the liquidity premium this would have the unfortunate consequence that the market consistent value of liabilities would change depending on the insurer's investment strategy. This contrasts with the observation that if the liability valuation is supposed to represent the transfer cost to a third party, that third party would necessarily pay the same price to either firm, irrespective of the composition of the assets held to support the business[8]. Given the above reasoning, the approach we adopt is that of using a theoretical replicating portfolio rather than the actual asset portfolio.

It is worth recognising that in practice the replicating portfolio will only (at most) be able to match the market risk characteristics of the liabilities and not the demographic ones or ones relating to policyholder behaviour (or management action). Under this framework it will be the uncertainty in the demographic assumptions and policyholder behaviour which drives the 'liquidity' of the liabilities. For example, it will not be the **size** of the lapse rate that will be important it will be the **uncertainty** in the lapse rate and the consequent **predictability** of future cash flows. This is consistent with the view given by the CFO Forum in their update and is also supported in the IAA note:

> *"What matters most is the degree of uncertainty regarding the timing of the ultimate cash flow to be paid from whoever holds the assets to whoever is owed the obligations…."*

Further, assets may not exist which capture the characteristics of the liabilities. For example, consider a risk-free portfolio of bonds that is cash flow matched to the liability cash flows for an annuity certain (with no lapse options). The asset portfolio is a matching one by nature (fixed) and term but not by 'liquidity' because risk-free bonds are typically highly-liquid market instruments.

The remaining question is how does one practically assess the liquidity of a set of liabilities? This is touched on at a high level only in the IAA note:

> "*In theory, this depends on how illiquid the liability is. Any increase in the uncertainty in the timing or amount of liability cash flows should increase the liquidity premium.*"
>
> *[By* ***'increase the liquidity premium'*** we presume they mean ***'<u>decrease</u> the illiquidity premium']***

And also in the IASB note where they discuss the use of a replicating portfolio as discussed above:

> *"Apart perhaps from some components of participating and unit-linked insurance contracts, there are few if any cases in which a fully replicating portfolio exits. However, it may be possible to apply the notion of a replicating portfolio to some of the cash flows from a contract and then apply other techniques to the remaining cash flows."*

The next section looks at two possible approaches to applying this definition to a set of insurance liabilities.

---

[8] This argument only applies to business where the liabilities are not contingent on the backing assets (i.e. excluding participating business).

3 Methodology

We start in Section 3.1, by describing a hypothetical approach for assessing the liquidity of liabilities. On closer inspection, however, it quickly becomes apparent that the modelling and calibration requirements will create a significant obstacle to its practical implementation. This leads us, in Section 3.2, to put forward a simplified approach that removes a significant amount of the modelling and calibration burden, but at the cost of requiring some significant assumptions to be made.

3.1 Hypothetical Approach

For a given set of assets and liabilities, a stochastic projection could be employed which would allow us to project the asset and liability cash flows through time, and to calculate the net cash flows. In particular, by incorporating a bid / offer spread process in the modelling exercise it would provide us with a relatively straight forward way to estimate the distribution of the cost of forced liquidations. Now, if we were to base the analysis on a replicating portfolio[9] of assets, then this analysis would provide us with the framework for making objective quantitative statements about the 'liquidity' of the liabilities.

However, at the heart of this approach is the requirement to project the bid / offer spread which in turn would require deriving and calibrating a credible model of how the bid/offer process evolves through time allowing for dependencies on all the relevant risk factors. This is likely to be a very challenging modelling /calibration undertaking[10].

We are currently not aware of any work or research focused in the area of modelling the evolution of the liquidity premium prospectively through time. Perhaps, once more work is done in this area, and credible models are developed and validated this alternative approach will become a more realistic possibility. Meanwhile, how can we overcome this technical road block?

The liquidity of a set of liabilities will be a function of:

1) the predictability of the cash flows (i.e. how likely are forced liquidations); and
2) the cost of forced liquidations.

If we make the simplifying assumption that we can ignore the interaction between these two drivers of liquidity then we can focus on the evaluation of each separately. This is the approach we adopt in the next section, where we first focus on evaluating the predictability of the liability cash flows, and then discuss how these can be used to assess the overall liquidity of the liabilities.

3.2 A Candidate Approach

Given the discussion from the previous section, this section aims to outline one possible approach to calibrating the liquidity of a set of insurance liabilities. The focus in this section will be on assessing the predictability of liability cash flows and will we further restrict the focus to assessing fixed (nominal) cash flows only. In theory, one may argue that other asset classes and liability types may have an embedded liquidity premium in them (i.e. small cap equities and liabilities based on such assets), but this is beyond the scope of this note. The approach is:

1) Construct a cash flow matching asset portfolio for the fixed liability cash flows using risk-free bonds.
2) Project the matching assets[11] and liabilities stochastically for the term of the liabilities and for each projection time step calculate the net cash flow. Rules will need to be agreed for investing surplus and funding deficits. For example, surplus could be invested in cash while deficit could be funded by first using any existing cash funds, followed by disinvestments from the matching portfolio.
3) The previous step will give an amount disinvested of the original matching portfolio for each time step in each trial. This can be divided by the initial market value of the cash flows at each time step and subtracted from unity to get a percentage not disinvested for each trial and time step. Call this

[9] Even a perfectly replicating portfolio would only be able to match market risks and not demographic and policy holder behaviour risks.

[10] It has proven challenging to get a consensus of the magnitude of the liquidity premium at a point in time. This modelling / calibration challenge will effectively require projecting the LP prospectively, which will be an order of difficulty greater.

[11] If the liabilities depend on the backing assets, then it will be necessary to project both backing assets and the matching asset portfolio.

the 'Predictability Ratio', PR where a value of the PR of 1 would indicate very predictable liabilities (from the firm's perspective) and 0 would represent unpredictable ones.

In practice, if firms are not interested in looking at liquidity by term, the simplifying assumption that the distribution of PR is the same at all terms could be made, which would allow the above routine to be trivially adjusted to calculate an aggregate (by term) PR.

The distribution of the PR over a number of stochastic trials will allow quantitative statements to be made about the predictability of the liabilities. An example would be that in 99 out of 100 times, a maximum of 10% of the matching portfolio had to be disinvested to fund unexpected cash flow shortfalls.

The above approach focuses on estimating the uncertainty in the asset liquidity event; the remaining challenge is how to set a consistent assumption for the proportion of the liquidity premium that should be allowed for in the valuation of a given set of liabilities. The following comments highlight why this is likely to be a very challenging question:

1) Original estimate of the liquidity premium. How much of the liquidity premium to allow for in the valuation of the liabilities will depend on the relative illiquidity of the assets used to estimate the liquidity premium in the first place. For example, if the estimate of the liquidity premium is derived from comparing a corporate bond plus credit default swap with a risk-free swap rate then it should be noted that these asset structures are not completely illiquid (though there will be considerable uncertainty in the pre-maturity realisable value). This suggests that a perfectly illiquid liability could arguably justify a higher liquidity premium than one based on a highly illiquid, but not entirely illiquid asset structure.

2) Converting cash flow uncertainty to allowable liquidity premium. There is also the very difficult theoretical challenge of how to convert a percentage liquidated number from the above analysis to a statement about the rational liquidity premium that could be allowed in the valuation of the liabilities.

We side step these issues and suggest a pragmatic approach. A simple starting point would be to base the proportion of the estimated liquidity premium on some statistic of the predictability ratio, which will lie in the region 0 to 1 by definition. If the statistic is close or equal to one, depending on how the PR is defined (mean, median, 5th percentile etc.), this would suggest that a significant proportion of the liquidity premium could be used for the valuation of these fixed cash flows. Then as the PR statistic tends down and away from 1 it will indicate that the liabilities are less and less illiquid, and hence that a smaller proportion of the liquidity premium could be used for the valuation of these fixed cash flows. While the absolute value of the statistic will be of interest, a comparison of the statistic across insurance products will help determine relative liquidity. In the next section we look at some simple worked examples.

4 Worked Examples

4.1 Simple Worked Example

Consider a simple example where the expected cash flows for a specific liability are a pay out of 5 at each of times 1 and 2. Further assume that there is some demographic/policyholder uncertainty associated with each of the cash flows such that the distribution of cash flows in a tranche of identical policies at time step t is: 5 + xt (where xt ~ N(0,σ)).

The matching portfolio for this simple liability consists of 2 risk-free zero coupon bonds, with terms of 1 and 2 years and nominal amounts of 5 each. Using Monte-Carlo simulation techniques we can calculate the aggregate Predictability Ratio (PR) in each of 2,000 trials. The summary statistics for the resulting distributions for PR for the 5 cases: σ = 0.1%, 0.5%, 1%, 2% and 3% are:

	Predictability Ratio, PR, for:				
Statistic	**σ = 0.1%**	**σ = 0.5%**	**σ = 1%**	**σ = 2%**	**σ = 3%**
Average	99.2%	95.8%	91.9%	83.3%	75.6%
StDev	1.2%	6.0%	12.0%	23.8%	32.4%
50th percentile	100.0%	99.8%	99.8%	99.0%	96.7%
25th percentile	98.6%	92.8%	87.2%	71.2%	56.2%
10th percentile	97.4%	86.8%	73.7%	47.8%	18.2%
5th percentile	96.7%	83.0%	65.0%	31.4%	0.0%
1st percentile	95.2%	76.5%	52.1%	5.2%	0.0%
0.5th percentile	94.9%	73.4%	46.2%	0.0%	0.0%

As expected, as sigma increases, the 'predictability' of the liabilities decreases. This is clearly illustrated by a decreasing Predictability Ratio, PR. For example, the average PR decreases from 99.2% when sigma is 0.1% to 75.6% when sigma is 3%.

The calibration of the metric is a separate question. If we believe, for example, that in the case where sigma = 3% the liability should essentially be considered liquid (i.e. it is not appropriate to use a liquidity premium in the valuation of it) then this would suggest that the use of a lower percentile (i.e. 5th, 1st or 0.5th) rather than an average may be more appropriate.

4.2 Simple Fixed Annuity Case study

This section builds on the previous section and estimates the PR for a fixed annuity insurance product. This type of product is often thought of as the classic illiquid liability – even though there exists mortality uncertainty in the annuity payments. This may help calibrate our understanding of the typical values / distribution for the PR which would be associated with illiquid liabilities.

We look at two annuity model points using Barrie and Hibbert's integrated stochastic annuity model:

Model point A: Tranche of immediate annuities of 1,000 p.a. to males now aged 65

Model point B: Tranche of 15 year deferred annuities of 1,000 p.a. to males now aged 50

Now, trivially, if we assume that there is no mortality uncertainty then applying the methodology from Section 3.2 will give a PR = 1 for all cases, for both model points.

Let us now introduce mortality uncertainty using Barrie & Hibbert's stochastic mortality model[12.] Integrated into the Barrie & Hibbert's annuity model is a stochastic mortality model, the form of which is:

Experienced Mortality = Best estimate mortality × Stochastic mortality factor

Omitting all technical details of the model, the one year volatility of the mortality factor, as per Barrie & Hibbert's best estimate assumptions, is 4%. The impact of this stochastic mortality variation on the annuity cash flows is illustrated in Exhibit 1.

Before we can calculate the PR for these model points decisions need to be made with respect to the rules for disinvestment and reinvestment of net cash flows. In this and the following example we assume the following rules:

- Investment of surplus. Any investment income in excess of the liability cash flow at each point in time is assumed to be invested in cash.
- Disinvestment of assets. Where there is insufficient investment income to meet the liabilities cash flow at any point in time the funds are taken from:
    - i. First, any cash funds which may have been built up from previous surplus arising.
    - ii. Any remaining funds are taken from the matching (risk-free bond) portfolio. The zero coupon bonds are disinvested in proportion to their market value.

[12] For more detail see: B+H Technical Note 2000/17: A Stochastic Model for Mortality.

Using these assumptions, the annuity model is used to project the matching portfolio and annuity liabilities stochastically to allow the calculation of the PR in each simulation.

Exhibit 1:
Impact of stochastic mortality on annuity cash flows for the two annuity model points A and B.

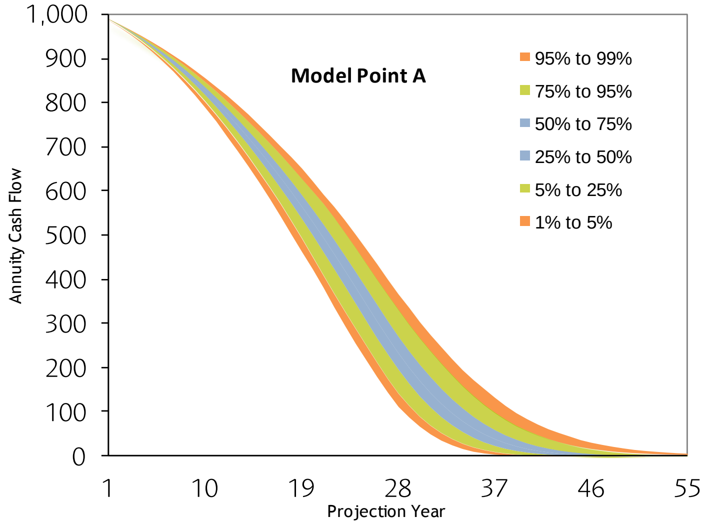

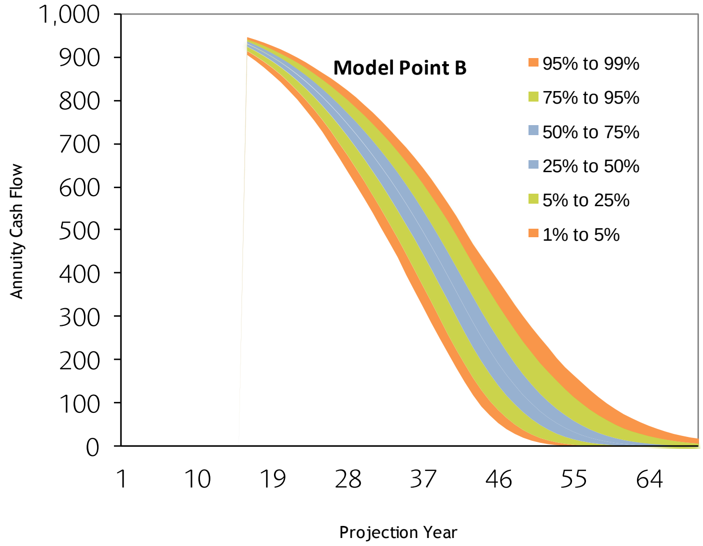

The summary statistics for the PR are given in the table below and their cumulative frequencies are plotted in Exhibit 2.

	Predictability Ratio, PR, for:	
Statistic	***Model point A***	***Model point B***
Mean	97.6%	98.2%
Std Dev	2.9%	2.5%
50th percentile	99.0%	99.6%
25th percentile	95.5%	97.0%
10th percentile	93.2%	94.4%
5th percentile	92.1%	92.8%
1st percentile	90.0%	89.8%
0.5th percentile	89.0%	89.1%

It is interesting to note that, while fixed annuity business is considered to be highly illiquid, on average 2% - 3% of the matching portfolio will need to be disinvested. In reality, this cash flow uncertainty could be met with a small holding of cash or other assets. However, in the more extreme scenarios (10 in 2000), just over 10% of the portfolio by market value was required to be liquidated due to policyholders living longer than initially forecast. This suggests that the proportion of the 'true' liquidity premium that could be justified in being associated with these annuity liabilities is in the range 89% to 98%.

Exhibit 2:
Cumulative frequencies for the predictability ratio for the two annuity model points A and B.

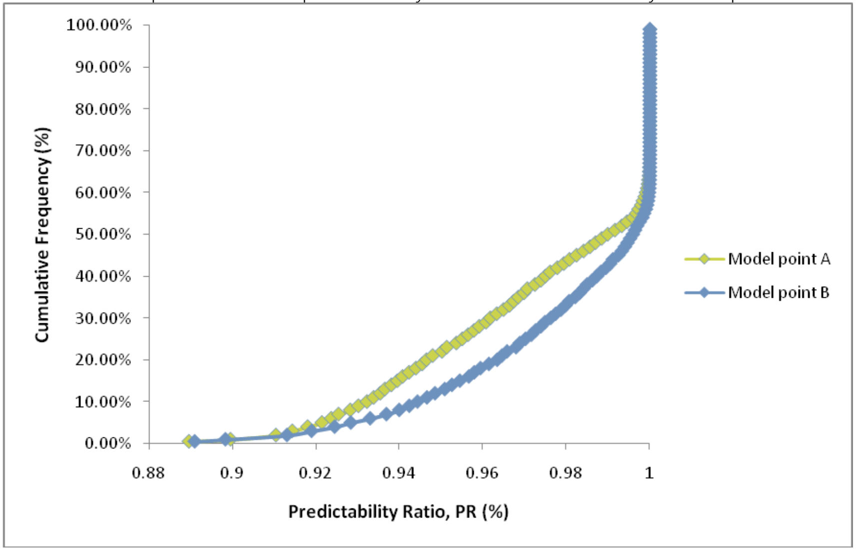

The next section looks at applying the same methodology to a simple endowment policy.

4.3 Simple Single Premium Endowment Case study

This section extends the analysis from the previous two sections to look at the following tranche of single premium endowment policies:

- Single premium policy for a fixed sum assured of 100,000 on maturity or earlier death
- Males aged 45 for 20 years
- Return of premium on lapse with a baseline lapse rate of 5% per annum.
- No allowance for expenses or commission.

For the above tranche, the expected initial cash flows are illustrated in Exhibit 3, below.

Exhibit 3:

Initial expected cash flows for a simple endowment model point.

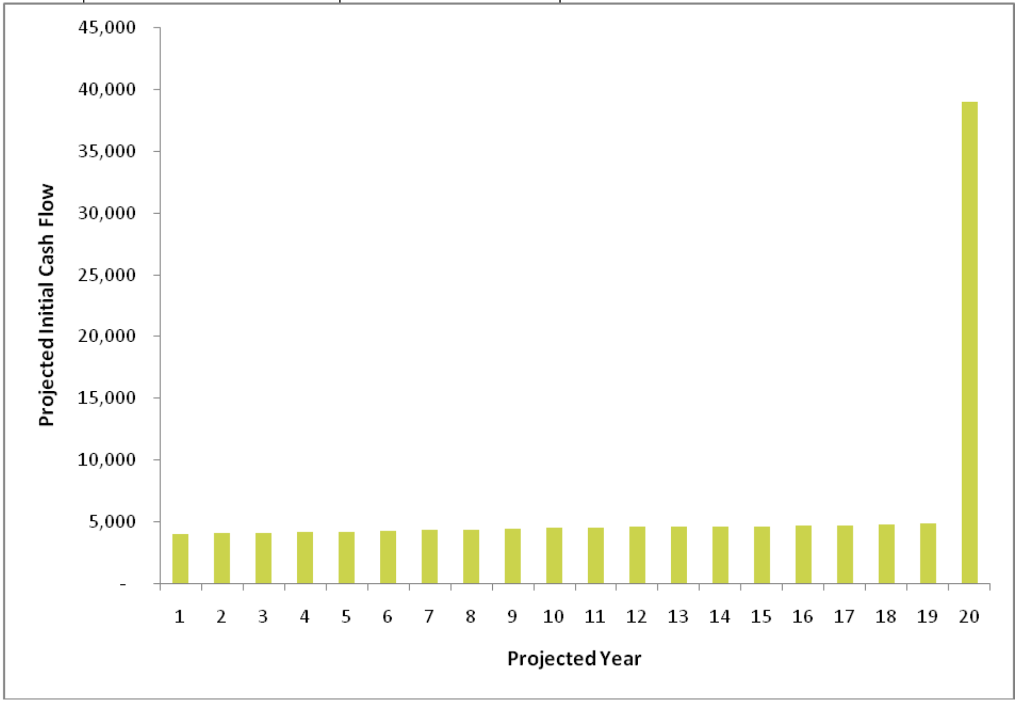

The main driver of uncertainty for such business will not be from mortality uncertainty but from policyholder behaviour i.e. policyholders choosing to exercise their option to lapse their policy and receive their initial premium back. We make a very simple assumption to allow for lapse uncertainty – we assume each year that the actual number of lapses follows a log-normal distribution with a mean of 5% and we compare four values for the standard deviation: 1%, 2%, 3% & 5%.

We also make the same assumptions as in the previous section regarding the investment / disinvestment of surplus / deficit during the stochastic runs. The summary statistics for the predictability ratios (PR) are given in the table below:

	Predictability Ratio, PR, with Lapses Dist'd LN with Mean = 5%			
Statistic	***St Dev = 1%***	***St Dev = 2%***	***St Dev = 3%***	***St Dev = 5%***
Mean	98.5%	97.3%	96.1%	93.6%
Std Dev	1.4%	2.9%	4.5%	8.1%
50th percentile	98.7%	98.1%	97.4%	96.3%
25th percentile	97.7%	95.7%	93.9%	90.7%
10th percentile	96.5%	93.3%	90.0%	82.8%
5th percentile	95.8%	91.3%	87.1%	77.1%
1st percentile	94.3%	88.0%	80.9%	65.1%
0.5th percentile	93.8%	86.9%	78.6%	59.7%

Comparing these statistics with those from the previous section we make the following observations:

- For the case where the standard deviation (SD) of lapses is assumed to be 1% or 2%, the PR statistics are very similar to the results for the annuity model points. This may suggest that in these cases the level of illiquidity of these liabilities is comparable to the annuity liabilities and hence one may argue that a similar level of liquidity premium could be used in their valuation (i.e. of the order 90%).

- For the case where the SD of the lapses is assumed to be 3%, the PR statistics are slightly lower. For example the mean is just above 96% (c.f. about 98% for the annuity model points) and the 0.5th percentile is below 80% (c.f. about 90% for the annuity model points). This may suggest that the level of illiquidity for these liabilities is less than for the annuity model points. If we assume illiquidity is assessed allowing for extreme cases (say at the 0.5th percentile or more) then an estimate of the proportion of the LP would be below 80%.
- For the case where the SD of the lapses is assumed to be 5% the PR statistics are, as expected, lower once again. The same reasoning as above suggests a proportion of the LP should be in the range 60% to 70%.

While clearly not an exact science, hopefully we have shown that this methodology can provide objective insights in to the predictability of liabilities that can be used as an input in to the assessment of the liquidity of liabilities. Further, as more liability types are studied, the more information available to compare the PR statistics should help achieve consistency in the assessment of liquidity between product types.

4.4 Extensions

The examples chosen were deliberately simple to help illustrate the proposed methodology. There are a number of possible complications that would need to be addressed depending on the particular application. For example:

1) Analysis by business line, rather than by policy. Depending of the fungibility of the assets supporting a business line, allowing for cross subsidy between policies will likely increase the assessed illiquidity of the business line.
2) Impact of renewal premiums. If renewal premiums are considered as an illiquid loan to the policyholder, this will likely reduce the assessed illiquidity of the business line.
3) Allowing for new business. What impact will an allowance for new business have on the assessment?
4) Impact of aggregating multiple business lines. Similar to 1, but across business lines.
5) Modelling products where the liabilities are in some way dependent on the underlying assets (e.g. with-profit business). This will introduce significant complications for the modelling work. In particular two sets of assets will need to be modelled – those backing the liabilities and the replicating portfolio. Also there will be complications around the replicating portfolio as its composition will expand beyond only nominal fixed income.

There are more factors that may force liquidation of assets than policyholder behaviour and variances in demographic variables, which this approach will not capture. For example, an important consideration may be the extent to which a regulator (or counterparty) could force the liquidation of assets under certain market circumstances.

The benefit of having a basic framework is that it lays the foundation for thinking about and analysing the more complex cases – such as those above, even if the framework itself will need to evolve in the process.

5 Summary

Given the recent momentum that has been gained on the use of a liquidity premium (LP) in addition to the risk-free rate for the valuation of insurance liabilities, there are some hard questions to be answered regarding the practical implementation of the underlying theory. Without a standard methodology for the estimation or application of the LP, there may be a danger that the inclusion of an LP in the valuation of liabilities becomes considered by external users of insurance companies' financial reports as a subjective 'fudge' factor, effectively reducing the comparability of results across companies and ultimately the credibility of the approach.

The aim of this note is to add to the discussion of how an objective, quantitative measure for the liquidity of specific classes of insurance liabilities might be developed, by tentatively putting forward a framework for assessing the predictability of insurance cash flows. While not directly answering the problem, this approach is able to provide insights in to the relative predictability of insurance liabilities – which in turn, we suggest, can be used (with some heroic assumptions) as a basis for measuring the relative liquidity of insurance liabilities.

We illustrated how the framework could be applied in practice by applying it to some simple insurance liability types. We showed that for the archetypal illiquid fixed annuity liability, in extreme circumstances, up to 10% by market value of a perfectly cash flow matched portfolio could need to be liquidated prior to maturity to meet unexpected cash outflows arising from annuitants living longer than expected. This in turn, would suggest that is may be imprudent to allow more than about 90% of a liquidity premium to be used to value such business. This would need to be balanced against the fact that the annuity cash flows may well be more illiquid than the market instrument from which the liquidity premium was estimated – which would support an argument that more than 90% of the LP could be used. It seems unlikely that all subjectivity can be removed from this estimation process.

We finished by briefly discussing some of the possible complications that will need to be considered when applying the framework more generally to insurance business.

Disclaimer