CORE RISK MODELLING FUNCTIONALITY

better data, better decisions

REDINGTON

1 EXECUTIVE SUMMARY

This document is designed to outline the basic tenets of our risk modelling. It is complemented by a variety of other, more detailed papers covering other asset classes. We chose this documentation structure to make it easier to update as we refine and improve our models. We separately produce an overview of the numbers produced by our model which is available at https://www.ajg.com/uk/benefits-legal-regulatory-information/asset-assumptions-summary.

Our model is largely intended to provide returns, and risk analysis at the 5th percentile level.

Where other papers are referred to, these are available to clients on request.

2 STOCHASTIC MODELLING OF KEY RISK FACTORS

Our stochastic model aims to quantify investment risk. This is the risk of assets or liabilities changing market value. We therefore include a variety of market risks detailed below. Other areas of risk such as governance, legal or political risk are not included directly in our model. Liquidity risk is allowed for in the risk/return assumptions of non-liquid assets but the primary tool for managing liquidity risk is our collateral and liquidity analysis and not the risk and return numbers. Demographic risk can be included in our projection of liabilities and details for our modelling of it are included but not all modelling includes it. Please find details of what we have modelled in the work produced or ask your consultant.

Our model is based around simulating different risk factors. We calibrate our model, every quarter end, for these key underlying risk factors using daily and monthly market data since 04 January 2006. Volatilities and correlations assumed are just those realised for each risk factor. We use this date as it is broadly when good quality inflation swap data became available. In general, we simulate instantaneous shocks calibrated to a 1y horizon.

We use normal and log-logistic distributions because they combine simplicity and tractability (especially minimizing the number of parameters needed) with a good fit to historical data, including allowing for fat-tails. However, this is not always appropriate, and we change when necessary. Some examples are described below.

Rates and Inflation

For interest and inflation rates we assume correlated normal distributions for each key tenor point. Gilt and swap rates are simulated separately.

Our model simulates nominal and inflation rates directly, and real rates as a consequence.

Equities, FX, Commodities and other Risk Premia

Equities, FX, commodities and other risk premia are simulated using a log-logistic distribution.

Brief outline of credit approach

For credit, we use a more complicated hybrid distribution, to capture the asymmetric behaviour of spreads when rising as against when falling. We also incorporate defaults and rating transitions. The full details are in a separate paper.

Exceptions

For some assets, using daily data is inappropriate (including credit spreads, as above) as the trading is not sufficiently frequent. In particular, property and hedge fund data are not readily available. We therefore use monthly data for the correlations of these asset classes.

For property, we also scale up the volatility to match the Solvency 2 stress of 25% at 99.5% probability (before converting to a log-logistic distribution). Using monthly data alone for illiquid assets will typically underestimate the downside risks.

For Insurance-Linked securities (such as catastrophe bonds) as well, typical distributions are clearly inappropriate, because the distribution of potential losses is highly asymmetric. We use key percentile figures and fit a Weibull distribution; we then assume the risk is uncorrelated with financial market factors (as has largely been true historically).

Principles we use for compound assets

Some assets, such as DGFs, are ultimately combinations of other risk factors. For such assets, we have specific modelling papers covering our approach. In principle though, we will typically attempt to map strategies onto existing risk factors unless we think there is something genuinely different.

Similarly, the value of an interest rate swap is a function of the relevant interest rate(s). Where assets can be valued using simple functions, we simulate the risk factors and re-price accordingly.

3 EXPECTED RETURN PRINCIPLES

Our core returns are derived every month end and mid-month if there has been a large market move. The methodology varies for different classes of assets depending on whether they are liquid or illiquid, and whether they are contractual or not contractual. We show this in a 2x2 chart below, with some example assets in each section.

We also use different assumption for different purposes:

- **Default Returns** – For many clients, especially those with asymmetric utility curves, it makes sense to use prudent assumptions.
- **Best Estimate Returns** – In other cases, it may be more coherent to use our best estimates of future returns.
- **Ultra Long Returns** – In some cases there is value in more stable assumptions with less dependency on current market conditions. For example, for end-game planning 10 years in the future, or for endowments, we may use our "ultra-long" (30 year) returns, which use largely the same approach but weighted towards longer-term averages. For example, in credit we use a spread input weighted much less towards current spreads and more towards long-term average spreads.

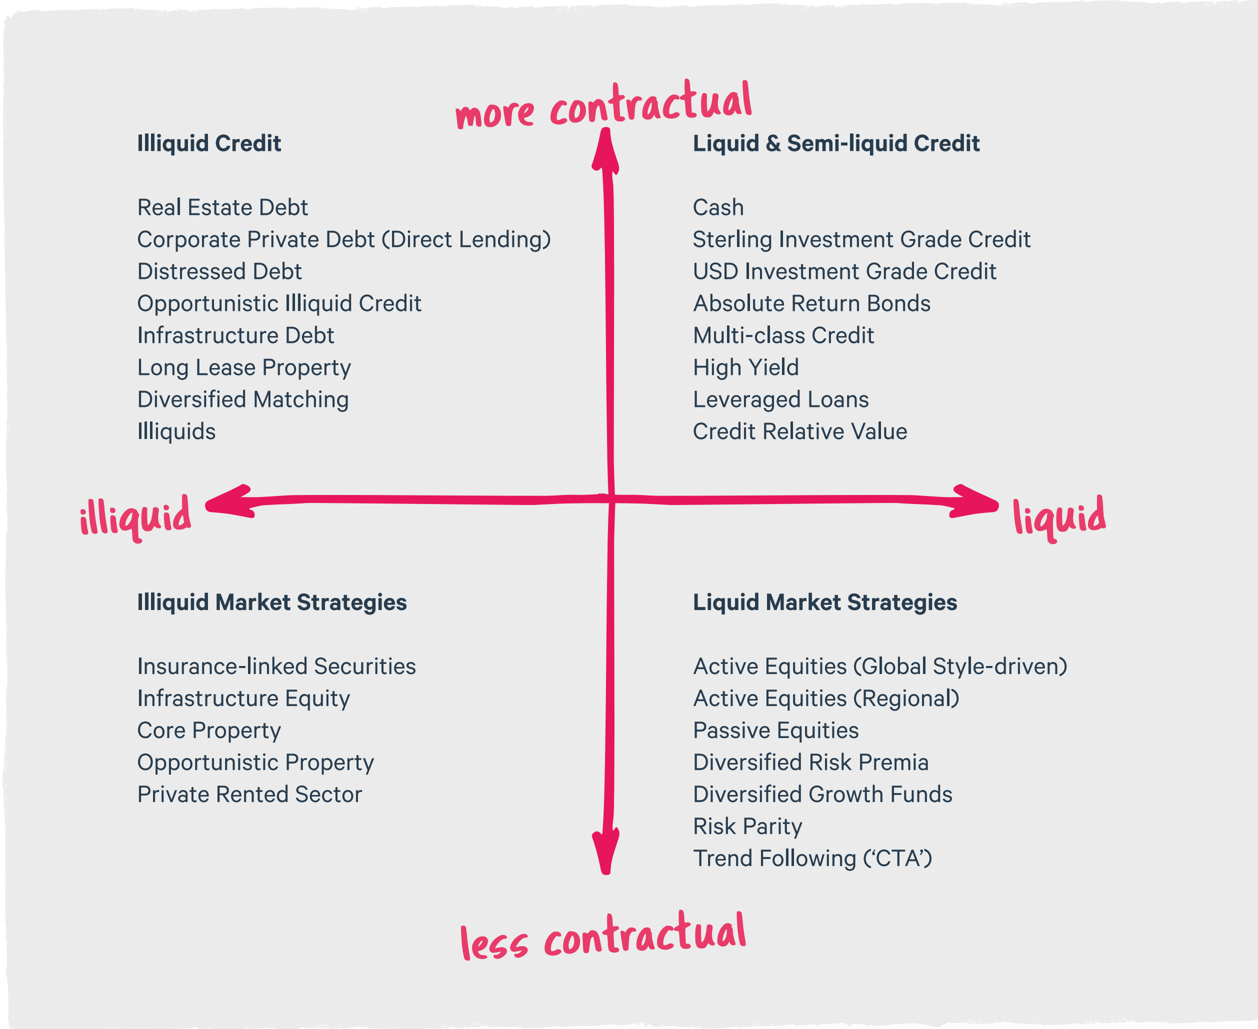

Our returns are available over gilts or swaps. We also offer them with and without term premium. While term premium is a core part of our return assumptions for clients investing to meet liability benchmarks any return on the asset side from term premium is offset by the liability side. For these clients we therefore exclude term premia.

- For equities we use the 40th/45th percentile (default/best estimate) of long-term rolling 10 year geometric mean excess returns. We use US equity as this is the longest data set going back to the 19th century. We use a percentile slightly below the median to account for survivorship bias. We then adjust for valuation and trend-based timing signals. For the ultra-long returns, we use the 45th percentile and do not adjust for timing signals but do adjust for current vs long term rates.

- For evidenced risk premia (typically in liquid markets), we use an excess geometric return based on a Sharpe ratio principle. The Sharpe ratio depends on the strength of evidence and is 0.2 for strong evidence, 0.15 for moderate evidence and 0.1 for weak evidence. For our best estimate and ultra long returns these ratios are 12.5% higher, broadly in line with a 50bp pickup in equities.

- Alpha risks use the same approach with using the strong evidence ratios for high conviction managers and weak evidence ratios for other managers.

- For contractual based instruments (e.g. corporate bonds), we use the spread over the relevant risk-free rate adjusted for default losses, transitions, rolldown, carry and adjusted for cross-currency basis. We also adjust for re-investment – this is most pronounced when spreads are high enough to offer super-normal Sharpe ratios, we assume they will not persist – i.e. in such market conditions we also make conservative re-investment assumptions. For default/best estimate returns we do this over a 10 year time horizon but for ultra-long a 30 year horizon is used giving greater weight to re-investment.

- For illiquid credit instruments, we use the same approach, but accounting for any illiquidity premia observed in the spreads on assets.

- For illiquid markets, we typically map the asset class to risk factors and derive the return that way.

- For property assets, we use 70% of the spread from observed rental yields adjusted for inflation and costs and 30% long term average returns. For the default returns we adjust for a margin of prudence which is set equal to the equity margin. For the ultra-long returns, we just use the long-term average returns.

We have a separate paper documenting our expected returns, and individual papers for each asset class. These are available on request.

4 FUNCTIONS AND RISK METRICS

This section documents the key risk metrics and functions produced.

Value at Risk (VaR) and Funding Ratio at Risk (FRaR)

The first two risk metrics are VaR and FRaR. These are calculated by simulating every risk factor together and finding the xth percentile value of the change in balance sheet (VaR) or funding ratio (FRaR). The risk factor split can be calculated by simulating each risk factor separately.

Flight Plan

The flight plan projects assets growing at their expected return, receiving future contributions and paying out expected liabilities. This can be used to express any of the return needed, the date to reach full funding, or the contributions required, each as a consequence of the other two.

FLIGHT PLAN TO FULL FUNDING BY 2034 (CURRENT CONTRIBUTION SCHEDULE) 31 MARCH 2017

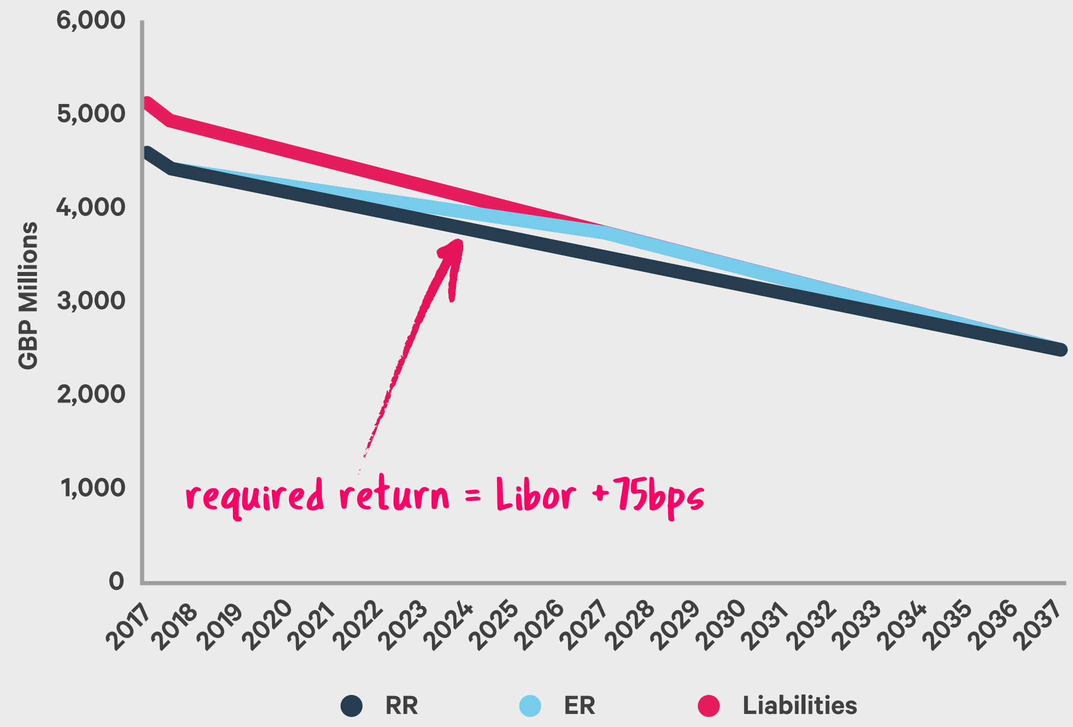

PV01 and IE01

We also calculate PV01s and IE01s. These are changes in the value of a portfolio (or a liability present value) arising from a 1 basis point move in interest rates or inflation. We calculate them by bumping the underlying curves half a bp up and half a bp down and taking the difference.

As a default, we calculate a curve structure of PV01s and IE01s using annual tenor points out to 100 years.

Stress Tests (single factor and historical)

Similar to the PV01s, we also calculate stress tests. These can be single factor (e.g., interest rates rise 1% and nothing else changes), multi-factor, or historical. For historical stress tests, we interpret the moves in line with our model. So, for example, an equity index falling from 100 points to 60 points would be interpreted as a 40% fall, rather than a 40 point fall.

Collateral and Illiquid Asset Budget

We calculate collateral statistics to show whether a scheme has sufficient collateral.

Total Collateral – This is the available LDI collateral assets (Gilts/Cash) and liquid growth assets which can be rebalanced into the LDI portfolio within 12 days.

Target collateral – Collateral required under a prudent stress of the derivative instruments and eligible collateral assets. Our prudent stress is broadly equivalent to a 4% rise in real gilt yields.

Net collateral surplus – The difference between total collateral assets (what's available) and target collateral assets (what's required)

Illiquid asset budget – This is the maximum amount of illiquid assets (excluding buy-ins) we recommend DB schemes hold such that after an 8% movement in real yields (4% move + 4% collateral buffer) you would not have more than 50% of the scheme in illiquids. It will depend on the yield sensitivity of the assets held and the duration of the scheme. This approach can also be used with other investors but is most pronounced for DB schemes and other investors with long-duration liabilities

5 LIABILITIES

Cashflow approach

Our liability modelling is based on inflating and discounting projected cashflows sourced from the scheme actuary. We typically simplify these into annual cashflows, although we will often split out the first three years into monthly cashflows where timing is an issue.

We take liability data in one of two forms – 2D and 3D. For 3D cashflows, we have active and deferred cashflows split both by year of payment and year of retirement (which then tells us when pre-retirement indexation switches to post-retirement indexation). For 2D cashflows we just have the cashflows by payment date and rely on the actuary to provide the average switching date. Pensioner data is identical in both forms.

We can also model simplified fixed and real benchmarks, which avoids confusion over minor or technical valuation differences.

Before we release outputs, we reconcile our liability modelling with the actuary's values.

SABR, Black-Scholes and Delta one approaches

Pension increases are generally capped and floored in retirement. This motivates the use of an option-pricing model.

At present, we can run the model using 3 different options. Our preferred approach is to use either a SABR model calibrated to market prices, or a Black-Scholes fixed volatility approach; we can accommodate a delta-one approach (where pension inflation moves in line with RPI right up to the caps or floors), though we think this approach has meaningful drawbacks and recommend using an option model.

We have a separate paper outlining the advantages and disadvantages of each approach.

Longevity

Sometimes, we also analyse the impact of longevity risks. If so, we will state this in our papers.

We carry out both deterministic and stochastic longevity analysis. Base case cashflows are reconciled with the cashflows provided by the Actuary.

For deterministic longevity analysis, we compare the impact on liability cashflows by changing mortality assumptions defined in the actuarial valuation report with other deterministic mortality scenarios.

For our centralised stochastic longevity modelling, we simulate future mortality rates using the Cairns–Blake–Dowd (CBD) 2 factor model calibrated against England and Wales mortality data over a 20-year period. Future pension liability cashflows are calculated under each mortality simulation. These results are summarised and applied to individual schemes based on their demographic make up.

Consistent with the simulation of our other market risk factors, longevity simulations are instantaneous shocks calibrated to a 1y horizon. We also extend the model to a multi-year, path-dependent mortality projection model, to simulate long term mortality paths.

6 CLIMATE AND CARBON METRICS

Emissions accounting

We calculate the emissions for liquid funds using underlying holdings data, where available, in line with PCAF's Global GHG Accounting and Reporting Standard for the Financial Industry. Aggregated metrics are calculated on the portion of holdings that has ESG data coverage, with the remaining holdings proxied using the covered portion of the fund. Proxying uncovered portions of a fund in this way enables us to provide a more complete strategic overview of the portfolio's position. For regulatory reporting we can show only the reported figures.

For strategic decisions, our preference is to treat short positions as negatives of long positions, because this is the only mathematically consistent approach.

Where ESG data is not available, ESG metrics are proxied using our centrally modelled, generic asset class universe. The modelling of this generic universe is reviewed on an annual basis. Where the asset class proxy is unavailable, a judgement has been made based on nature of the fund on a best endeavours basis.

Portfolio Alignment Metrics

We calculate three different portfolio alignment metrics to help our clients meet new TCFD reporting requirements introduced in 2022. These are:

- SBTi score and TPI scores: we calculate the proportion of a portfolio (amount invested) that has Paris-aligned decarbonisation pathways that are approved by the Science Based Transition initiative (SBTi) or the Transition Pathway Initiative.

- Implied Temperature Rise: expressed in degrees Celsius (°C), the Implied Temperature Rise metric estimates the global implied temperature rise (in the year 2100 or later) if the whole economy had the same carbon budget over-/undershoot level as the company (or portfolio) in question.

We only calculate SBTi and TPI scores where the underlying holdings data is available and do not produce proxies for these metrics. We ignore short positions because there is little industry-wide consensus on how these should be treated. We use a similar approach to proxying portfolio ITRs as we do when proxying portfolio emissions.

Climate risk and opportunity

Climate scenarios have been calculated for our generic asset class universe. These climate-related risk and opportunity stresses show the impact on asset values in today's terms as an instantaneous shock. Three main types of scenarios considered: orderly transitions, disorderly transitions, and hot house world (limited/insufficient transition). These stresses are reviewed on an annual basis.

1 DATA SOURCES

Our primary market data sources are Thomson Reuters, Bloomberg and Merrill Lynch. We also get RPI and LPI swap quotes from a variety of banks and ESG data from MSCI and TPI. This data has historically been very reliable and is widely used with in the market. We therefore preform high level checks on this data only.

We also rely on scheme actuaries for liability information. Asset data is sourced from either managers or custodians. Our primary check on this data is to compare the data received to the previous data allowing for changes in time and market conditions. This highlights any large unexplained moves which we can then query.

Our asset/liability model is dependent on data from third parties including, historical market data and the asset and liability data above. Please note while we carry out some high-level checks on the data, we take all information we receive “as is” and we will not be responsible for any errors or omissions in data pro-vided by you or any third party. The outputs of our model will ultimately depend on the accuracy of any data we receive.

8 LIMITATIONS

All models are less reliable in the extreme tails because there is less information against which to calibrate them.

In between valuations, we have to either make adjustments or simply roll forward the cashflows to update them to the date of the analysis. There is a risk that the profile changes meaningfully, and we are reliant on the actuary to tell us if it does.

We are implicitly assuming that the history since 2006 will be reflective of the future. We are therefore at risk of a regime change.

The model also takes no account of developments since the date it was calibrated.

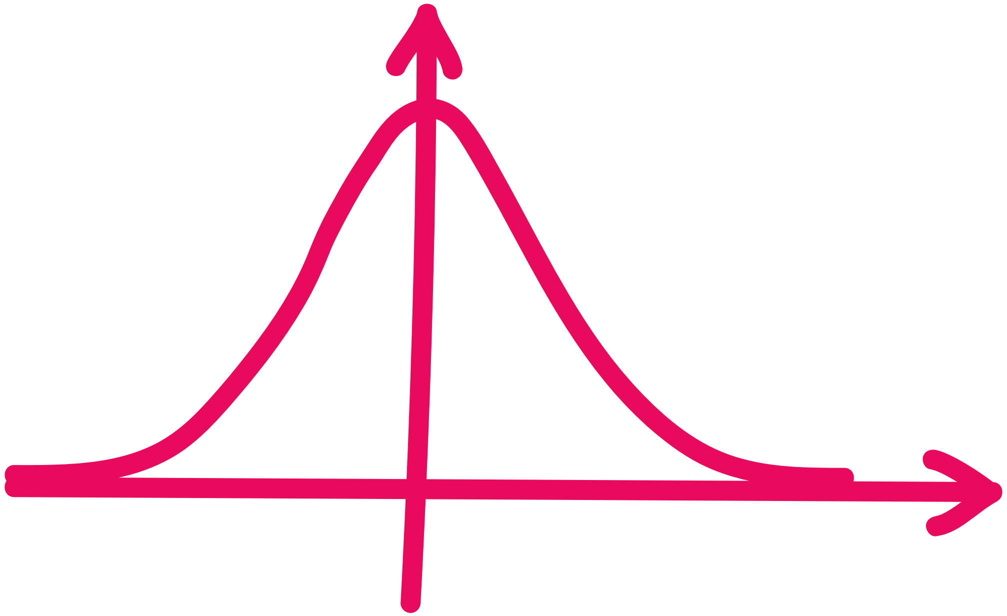

9 TECHNICAL APPENDIX

Our standard risk settings

- Historical daily data since Jan 2006 (monthly for less liquid data, especially credit spreads)
- 365 day annualization factor
- Decay factor of 1 (i.e. all data equally weighted)
- 1 day return horizon

Curve construction details

We use a cubic spline to fit a curve around the key tenors.

Curve and PV01 types

PV01s are calculated based on the curve used – so a par curve will create a par PV01, and so on. We typically use annual curves for key curve construction, though this can vary depending on the convention in different jurisdictions.

Correlation approach

We use Cholesky matrices to generate correlated random numbers. We generate correlated normal distributions then percentile map to logistic distributions for logarithmic returns.

Risk Premia

We classify risk premia based on the statistical probability of being random noise (based on historical data) providing these fits with an economic / behavioural rationale. If the probability is below 5%, we classify the premium as “Strong Evidence”, 10% as “Moderate” and 20% as “Weak”. Otherwise, we classify the risk premia as Unrewarded and do not give any expect return (over gilts) for that risk.

WANT TO DISCUSS FURTHER?

get in touch

Alexander White
Managing Director, Co-Head of ALM

ISO/IEC 27001 Information Security Management

QAS Institute and Faculty of Actuaries Quality Assurance Scheme

Floor 6 One Angel Court
London EC2R 7HJ
+44 (0)20 7250 3331

REDINGTON