

NOTE: This document was first circulated to stakeholders in November 2016 as part of the Trustee's preparations for the 2017 valuation.

In December 2017, a formal actuarial report was submitted to the Joint Negotiating Committee. This report set out the Trustee's conclusions on the scheme's funding position (based on the existing level of scheme benefits).

Proposed Approach to the Methodology for the 2017 Actuarial Valuation

Response to the Valuation Discussion Forum (VDF)

Summary

This paper addresses the methodology to be used in the 2017 actuarial valuation of USS. It is divided into two parts. Part 1 provides a relatively brief and high-level overview of the proposed approach to the methodology in its entirety. Part 2 addresses the specific questions about the valuation methodology posed by the Valuation Discussion Forum (VDF) in its letter of 16 May 2016 and further elaborated in various follow-up letters and discussions. Both of these parts of the document are presented at a level of detail agreed in discussions with the VDF co-chairs.

On the basis of these discussions this paper focuses on the methodology rather than the input assumptions. We do acknowledge that the distinction between methodology and input assumptions is not always clear, but we have endeavoured to provide a consistent differentiation between them. It is our intention to provide, at a later time, documentation of the input assumptions. The latest version of the [2017 Valuation Timetable \(USS, 2016b\)](#) provides an overview of how all the elements of the valuation fit together.

Note also that this paper focusses on the 2017 valuation, rather than the 2014 valuation (although the latter is referred to when it helps the explanation of certain points).

Contents

PART 1: Methodology Overview	3
PART 2: Reply to VDF Questions	8
References and Bibliography	25

PART 1: Methodology Overview

This section provides an overview of the methodology to be used for the 2017 actuarial valuation. The focus here is on the general framework for the valuation, not the detailed inputs and assumptions that will need to be made in terms of implementation.

The main purpose of the valuation can be summarised in two interrelated objectives:

- To establish whether the assets held by the scheme are sufficient to cover the benefits promised to members (as of the valuation date), given the employers' covenant.
- To determine the level of contributions required from employers and members to meet future benefits and clear any past service deficit.

There are four main elements to the valuation methodology:

1. Review of the employers' covenant;
2. Analysis of the scheme at the end of the covenant horizon and at the valuation date;
3. Setting a risk budget to maintain scheme risk within the covenant parameters;
4. Incorporating prudence.

It is important to note that the methodology not only relies on a degree of professional judgement, but is also somewhat iterative in its implementation and there is no unambiguous closed-form solution that tells us the valuation output. However, in what follows we have attempted to set out the steps involved in arriving at the valuation outputs.

1. The covenant is the foundation of the valuation

The valuation methodology is centred fundamentally on the covenant of the scheme's sponsoring employers. The covenant reflects the degree to which sponsor support can be relied upon to fund the scheme and pay future pensions. This support reflects the ability (capacity) and willingness (appetite) of sponsors to make contributions into the scheme not only on a regular, planned basis, but also on a contingent basis should certain adverse events materialise.

As a result the starting point for the valuation is a determination of the capacity of employers to support a certain level of remedial contributions in the event of an adverse outcome for the assumptions used to set the funding basis for the scheme. In particular, there needs to be an assessment of:

- The levels of contributions that could be sustained, to support accrued promises in the event of materially worse than anticipated outcomes;
- The covenant horizon, i.e., the future time period over which there is sufficient visibility of the continued strength of the covenant, to underpin planned contribution rates and contingent claims.

In practice there is no single answer to the first of these questions, because it involves a trade-off for both the sector and the scheme. As the contributions required of the employers increase, there comes a point beyond which their plans for future investment (in new facilities and infrastructure) must be curtailed and/or the financial flexibility necessary to help them weather a downturn in fortunes is significantly reduced. At this point the size of contributions begins to weaken the covenant, particularly if the covenant horizon is long. Of course there is not a sharp transition point at which this happens and therefore the answer to the question of what level of contributions can be sustained is

actually a range. At the start of the range there is little or no impact on future investment plans and at the end of the range all but the most important strategic projects may be curtailed.

It is important also to note that different employers have different risk absorbing capacities. The trustee's judgement is required to set an aggregate risk budget that is not determined by the risk capacity of the weakest employers, but on an aggregate assessment. This may lead to some weaker employers bearing (and benefiting from) more risk than they might be able to do independently. This is a key feature of the mutualisation of risk in the scheme.

Having addressed the capacity of employers to support the scheme, the valuation must next address the question as to how much of this capacity should be relied upon in the funding arrangements of the scheme. In other words: what is the degree of reliance that the scheme should reasonably place on the covenant, in order to facilitate the pursuit of uncertain investment returns, that can deliver cost effective pensions? This is a judgement made on the basis of a detailed review of the financial position of employers, and consultation with the sector as a whole.

The judgement that the trustee came to in the 2014 valuation was that the amount of reliance on the covenant in the existing funding arrangements was close to the maximum available, and should not be allowed to increase over time. This was enshrined as the first of the "guiding principles" adopted by the trustee at that time ([An integrated approach to scheme funding, USS 2014a](#)).

The adoption of this principle requires a projection of the demands the scheme might make of the covenant at the end of the covenant horizon, and ensuring these demands remain proportionate to the likely scale of covenant support available at the time. Given the length of time involved in these projections, the importance of ongoing monitoring and periodic adjustment to maintain the proportionality of the parameters is an important aspect of the funding plan.

2. (a) Evaluate the scheme at the end of the covenant horizon

To assess the position of the pension scheme at the end of the covenant horizon, the pension cash flows of the scheme (based on the current benefit structure) are projected forward. This requires a set of assumptions about salary growth and membership, which are documented, and are consistent with other valuation assumptions.

This understanding of the development of the scheme's liabilities is a critical requirement to setting a risk budget at the end of the covenant horizon that is defensible. The chain of reasoning is as follows:

- The scheme is fully-funded on a technical provisions basis at the end of the recovery period (by definition). Hence the level of assets is equal to the technical provisions at that time. It would be difficult to argue that the recovery period should be longer than the covenant horizon and, as a consequence, the scheme should also be fully funded on a technical provisions basis at the covenant horizon.
- The level of technical provisions at the covenant horizon can be determined from the principle that the amount of reliance on the covenant should be no greater than it is at the valuation date.¹ The reliance on the covenant can be measured as the difference between the projected

¹ In particular, note that principle one adopted by the trustee in 2014 states that 'Over the period for which there is visibility of the covenant (estimated to be 20 years) there should be no increase in USS's reliance on the covenant of the sector and, where opportunities arise, the reliance on the covenant will be reduced if possible.' Actually to be precise, this principle as stated only determines the *minimum* level of technical provisions and the *maximum* discount rate: a decision on the specific level of reliance must be made to fully determine these values.

value of assets at that time (which is equal to the technical provisions) and the value of assets required at that time for self-sufficiency (i.e., the value of assets such that the scheme is expected to have no need of further contributions to meet accrued pension benefits). Hence at that time the difference between the self-sufficiency measure of the liabilities and the technical provisions measure should be no larger than it is at the valuation date in real terms.

- Having determined the technical provisions at the end of the covenant horizon, the discount rate at that time can be calculated from the projected pension payments. This discount rate applies only for the time period beyond the covenant horizon.
- Knowing the discount rate at the end of the covenant horizon means that we also know the required aggregate expected return on the portfolio of assets at that time (for an assumed level of prudence – see section 4.)
- Knowing the expected return on assets we can construct an “efficient” investment strategy in terms of risk vs. return for the period beyond the end of the covenant horizon.

2. (b) Evaluate the scheme from the valuation date

Having determined the discount rate and investment strategy relevant at, and beyond, the covenant horizon, the next step is to determine the combination of contributions and investment returns required between now and then to deliver an asset portfolio which is equal to the technical provisions at the end of the covenant horizon.

There is in general no single clear-cut solution to this problem. The principle remains that through this period the level of reliance on investment returns from return-seeking assets to pay pensions should be consistent with the ability of the sector to pay contingent contributions if investment outcomes are adverse.

If investment risk were to be too high, the correspondingly high discount rate for technical provisions would lead to too great a gap between technical provisions and self-sufficiency, and therefore too much reliance on the covenant. By ensuring that – at the beginning and through the covenant horizon – this reliance on investment returns is proportionate to the projected size of the employers’ future USS payroll, the sector’s ability to respond to adverse investment outcomes should be no worse than current levels.

This step in the valuation essentially involves determining a consistent combination of the following:

- Length of the deficit recovery period (which we assume to be no greater than the covenant horizon);
- Contribution level;
- Expected investment return between the valuation date and the end of the covenant horizon.

In practice this determination is iterative in nature and generally (although in principle it need not) seeks to achieve a smooth transition between the current investment strategy and the investment strategy at the end of the recovery period, and at the end of the covenant horizon.

3. Setting the risk budget

The risk budget of the scheme is set by assessing the gap between a self-sufficient funding level and the technical provisions, as outlined above. This is a proxy for maintaining an appropriate balance between investment return from return seeking assets (which require both economic growth and effective execution of the strategy) and cash contributions.

Should the trustee deem at any stage that anticipated returns are not likely to materialise as anticipated in the funding plan, then in order to pay previously accrued pensions, contributions would need to increase.

The trustee has deemed that the ability to increase contributions up to a level that would allow for a material further de-risking, and payoff of accrued pension benefits from available cash flows is the required 'buffer' contribution level.

For that reason, the trustee applies the test that over the covenant horizon, the sum of the 'contingent cash flows' available to the scheme should prescribe the limit for the gap between the technical provisions and the self sufficient funding measure, at the end of the covenant horizon, and the scheme's funding arrangements should be subject to this test.

4. Overlaying prudence

Decisions on how much prudence should be included in the valuation are decided by the trustee acting on advice. Prudence is required in order to ensure:

- An amount of adverse experience can be absorbed by the scheme without immediate recourse to the employers' covenant;
- Compliance with the Pensions Regulator's guidance.

The main areas where the best estimate assumption is likely to differ from that adopted in the calculation of the technical provisions include:

- Discount rate (this is the main area of prudence);
- Gap between RPI and CPI;
- Long-term trend in mortality improvements.

Prudence in the discount rate is reflected in a "spread" or "margin" between the discount rate used to estimate liabilities and the expected return on assets. In particular, the discount rate will be lower than the expected return so that assets are expected to grow at a slightly faster rate than that at which liabilities are discounted. This leads to a larger deficit than best estimate.

The combined effect of all the sources of prudence can be thought of as follows. Whereas the best estimate for the deficit at the end of covenant horizon corresponds to the median or 50th percentile outcome, a prudent estimate might correspond to an outcome, for example, at the 60th to 70th percentile.² This enables a certain amount of adverse experience to be absorbed without having to go back to the drawing board and redo the entire valuation, in order to increase contributions.

Note that this margin in the discount rate should be proportional to the amount of investment risk incorporated in the best estimate assumptions. In the 2014 valuation, for example, the level of prudence that was built into the technical provisions at the covenant horizon is substantially less

² This approach facilitates an overall perspective on the aggregate degree of prudence across different financial inputs. Note that the range quoted is merely an example.

(about half) that built into the technical provisions at the valuation date, reflecting the planned reduction in investment risk over the covenant horizon period.

Additionally it is not uncommon to use a lower margin for prudence in the determination of the deficit contributions than in the determination of the technical provisions. In fact, this was the case with the 2014 valuation in which the margin for prudence associated with the contributions was substantially less than that associated with the technical provisions. In particular, the calculation of the required deficit contributions (as opposed to contributions relating to future benefit accruals) in the 2014 valuation involved adding back half of the reduction for prudence that was made to the expected return in establishing the discount rate for the technical provisions calculation.

The prudence that is incorporated into the other assumptions is typically much smaller than that built into the discount rate. Furthermore, unlike the discount rate, the prudence that is built into other assumptions, e.g., mortality improvements, has minimal impact on the reliance on the covenant.

It is important to remember that the appropriate level of prudence is ultimately a matter of judgement, which the trustee makes with the aid of professional advice. Also, incorporating some degree of prudence is a part of the methodology as opposed to a driver of the methodology.

PART 2: Reply to VDF Questions

In the following pages we provide responses to the questions posed by the VDF on valuation methodology.

The VDF indicated that their members wished to understand the methodology being adopted for the 2017 valuation, as opposed to the actual inputs (i.e., assumptions) that will generate the output of the valuation. For the avoidance of doubt the key inputs include the following:

- The expected returns on different assets classes;
- The current and expected investment strategy;
- The level of prudence;
- The discount rate;
- The assumptions on future demographic experience (mortality, withdrawal, ill-health retirements, proportion married, etc.);
- The financial assumptions (including RPI, CPI, salary increases, pension increases);
- The length of the recovery period for any deficit.

To the extent that it is relevant to the discussion of methodology, our responses touch upon the approach to determining some of these inputs; however we have not detailed the analysis we would undertake and the full set of factors that we would consider in establishing the inputs.

The responses build upon the Methodology Overview (presented in the previous section). Following guidance from discussions with the VDF co-chairs, we have kept the answers relatively high level, but added more detail where requested.

It may be helpful to note that when we refer to “risk” in this document we are intending to reflect the following definition: “the potential for the scheme’s funding objectives not being achieved”. In the context of the financial performance of the scheme this risk is generally measured in an asset-liability context in terms of statistical metrics summarising downside scenarios for the deficit. The precise metric used in any situation will vary depending on the precise context and objective.^{3,4}

With this in mind, we now turn to the questions posed.

³ It is important not to confuse the concept of risk with how it is measured. Risk may be measured in many different ways. The measurement of risk involves specifying (i) the variable of focus (e.g. the deficit, or the funding ratio, or the value of assets at a particular future time), and (ii) the statistical metric used to reflect the risk (e.g. standard deviation, percentile level, tail expectation, expected shortfall relative to a target, value-at-risk, etc.).

⁴ In addition we emphasise that all outputs of the valuation process are estimates. The accuracy of these estimates depends on the structure of the valuation model, as well as the uncertainty associated with the inputs. Whilst standard errors can be computed for certain components of this uncertainty, they do not capture the full range of estimation error. In particular, the errors associated with model misspecification are generally not probabilistic but systematic in nature and their impact is not always known.

1. Transparency

1.a. UCU and UUK would welcome an explanation of the current methodology, including a demonstration of how the evidence is used in the development and setting of the valuation methodology.

In undertaking the valuation the trustee is attempting to establish:

- Whether the assets currently held by the scheme are adequate to cover the benefits promised to the date of the valuation, given the covenant provided by the participating employers. For this purpose assets are taken at market value and the assumptions used for measuring liabilities are consistent with market conditions and the demographic profile of members prevailing at the valuation date.
- The level of contribution required from the employers and members to (i) finance benefits going forward, and (ii) clear any deficit identified, allowing for the employer covenant.

The term “technical provisions” is a statutory term which refers to a particular way of valuing the pension liabilities. It corresponds to the measure of liabilities used to set the level of assets which should be targeted for funding purposes. It is calculated by discounting the projected liability cash flows using an appropriate discount rate. This discount rate should be determined by the expected return associated with the investment strategy but with a margin of prudence.

In addition there are statutory requirements to calculate the liabilities in other ways, in particular:

- The position on a section 179 valuation (i.e. the value of benefits under the PPF);
- The wind-up position of the scheme (i.e. the buyout value of the liabilities).

The requirements for these statutory valuations are common to all pension schemes, but the trustee has considerable flexibility in the choice of methodology in respect of determining the required contribution rates.

i. Covenant assessment

The first part of the methodology for the funding valuation is to decide how much reliance can be placed on the covenant of the participating employers. Broadly speaking the stronger the employer covenant, the more risk can be taken in the funding of the scheme: for example in the investment strategy or the level at which technical provisions are set.

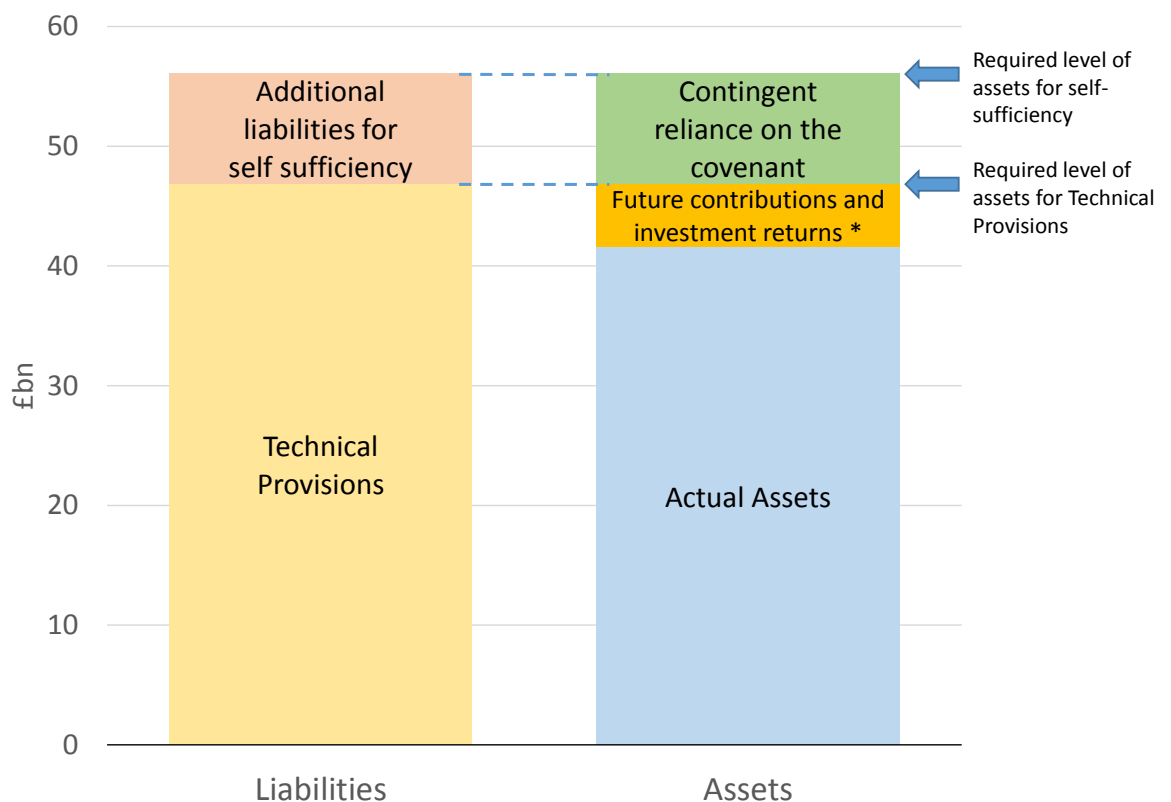
Simplistically, a scheme where there is no significant covenant (i.e. the sponsor cannot be relied upon to fund the deficit) would require sufficient funds to meet all future benefit payments with a high degree of certainty and low investment risk (we refer to this as the ‘self-sufficiency’ position).⁵ Where there is a substantial covenant, as is the case with the participating employers of USS, the assets held (or more exactly the assets targeted to be held equal to the technical provisions) can be lower than those needed to be self-sufficient.

⁵ The ‘self-sufficiency’ measure of the liability reflects the required level of assets to meet all future benefit payments to a very high probability without the need for additional contributions. It corresponds to a discount rate very close to gilts (i.e. very close to the yield on an appropriate portfolio of UK government bonds). The discount rate used by the trustee for the self-sufficiency liability is gilts + 0.5%.

ii. *Link between covenant and technical provisions*

From the above discussion it should be clear that the difference between the scheme’s technical provisions and self-sufficiency position represents a measure of the level of contingent reliance on the covenant, which is in addition to the reliance reflected in the agreed future contributions to fund the past service deficit over the recovery period. The diagram below shows how the reliance on the covenant is related to this difference. See also the discussion in the Methodology Overview section.

Figure 1: The relationship between the technical provisions and self-sufficiency liabilities, along with the actual and target (i.e., required) levels of assets as of 31 March 2014.



* This refers to future contributions to fund benefits accrued to date, not benefits that will be accrued in the future

A strong covenant allows a scheme to take more risk and hold less assets in the scheme, leading to a lower level of technical provisions than would otherwise be the case. The weaker the covenant the higher the technical provisions need to be, as the scheme needs to be closer to being self-sufficient.

When forming a view of the covenant provided to the scheme by the participating institutions consideration is given to:

- The time horizon over which there is good visibility of the sector's ability to support the scheme. We call this the "covenant horizon";
- The ability and willingness (affordability) of the sector to make contributions to the scheme;
- Tests of the future financial performance, free cash flow and robustness of the sector;
- The assets available to the scheme held by the sector;
- How the covenant is expected to develop over time.

At the last valuation the trustee was advised by EY on the covenant. For the 2017 valuation the trustee appointed PwC as primary covenant advisor, with specialist support from EY Parthenon for their detailed knowledge of the higher education market.

iii. Principle that covenant reliance should not increase

A key principle that the trustee adopted at the last valuation (see question 1.b) was that over time, and in particular over the 20 years for which there was good visibility of the covenant, there should be no increase in the reliance (measured in real terms) on the covenant provided by the participating employers ([USS 2014a](#)). This is a principle we propose to maintain at the 2017 valuation (although the covenant horizon may be different from the 20 years obtained in 2014).

To determine the self-sufficiency position at the end of the covenant horizon, the scheme's cash flows are projected forward allowing for future benefit accrual up to the end of the covenant horizon and salary growth, as well as assumptions about changes in membership up to and beyond that point. The technical provisions at that time are then determined so that the expected difference between the two bases for the liabilities (i.e technical provisions vs. self-sufficiency) does not increase in real terms over the covenant horizon (this is the reliance on the covenant in Figure 1). Knowing the required technical provisions at that time allows the establishment of a suitable discount rate and investment strategy at that time. These results are then input into the determination of the investment strategy over the whole of the covenant horizon and the discount rate at the valuation date. This then leads to the determination of the technical provisions at the valuation date (incorporating an "appropriate" degree of prudence). See the earlier discussion in the Methodology Overview section for more details.

iv. Determining contributions

To the extent there is a deficit at the valuation date (i.e. the technical provisions are greater than the assets held), decisions need to be made about how the deficit will be made good. This frequently involves using a lower margin for prudence to determine the deficit contributions than is used to determine the technical provisions (see section 4 of the Methodology Overview). The reduction in prudence depends upon the level of additional expected out-performance of assets over the technical provisions discount rate and the length of the recovery period (used to establish the required contributions). Here again the covenant is important in deciding on (i) the length of the recovery period and (ii) the extent to which the out-performance is allowed for, as it reflects the participating employers' ability to make good the deficit if the additional return does not materialise.

At the last valuation when calculating the deficit recovery contributions the trustee adopted a 17-year recovery period (not to be confused with the 20-year covenant horizon), and allowed for a lower margin for prudence than that used in the technical provisions.

1.b. In particular, it would be helpful to understand how the three tests align with the current methodology.

The three tests are derived from the three guiding principles established by the trustee as part of the 2014 valuation ([USS 2014a](#)). These guiding principles are as follows.

Principle 1: “Reliance of the scheme on the sector. Over the period for which there is visibility of the covenant (estimated to be 20 years) there should be no increase in USS’s reliance on the covenant of the sector and, where opportunities arise, the reliance on the covenant should be reduced if possible. The reliance on the sector will be measured as the additional contributions which would be required if the trustee moved to a relatively low risk approach to investment strategy and therefore could not rely on the same level of investment returns which are anticipated under the current investment strategy.”

Principle 2: “Stability of contributions. There should be a high probability that the employer contribution rate will not exceed 18% of salaries over a three year period and there should be a very high probability that the employer contribution rate will not exceed 21% of salaries over the same period. In the longer term the stability of the contribution rate should be increased.”

Principle 3: “Investment risk and tail risk. The balance sheet of the scheme’s participating employers should be able to cover the impact which a rare set of adverse circumstances (tail risk) may have on the funding position of the scheme. This includes being able to cover both the level of any existing deficit, plus an allowance for a potential increase in this deficit over a one year period if an exceptional economic event were to occur with resulting adverse impacts on investment returns.”

Test 1: This test measures the reliance being placed on the covenant. Specifically it measures the difference between the technical provisions and the amount of assets which would be required for self-sufficiency (as shown in Figure 1). At the 2014 valuation the trustee decided that the reliance should not exceed the value of additional contributions (above the agreed 18% of salary) that they believed could be available from the participating employers over a 15-20 year period. Advice received by the trustee at that time indicated 25% of salary was an upper limit (i.e. a further 7%).

Test 1 is used to monitor consistency with the trustee’s principle that no greater reliance be placed on the participating institutions on an ongoing basis. The amount of reliance which can be placed on the employer covenant now and at the covenant horizon is used to determine the level of technical provisions.

Test 1 also has implications for the investment strategy. The stronger the employers’ covenant then the greater the capacity for taking investment risk. Similarly, the weaker the covenant then the lower the capacity for taking investment risk. Because of the relationship between investment risk and return, the long-term expected returns associated with the scheme’s investment strategy should be higher for a stronger covenant and lower for a weaker covenant.

Test 2: This test measures the likelihood (probability) of the need to increase contributions at the next valuation. This reflects the employers' preference for a stable contribution rate and desire for a low probability of contributions exceeding 18%, and a very low probability of exceeding 21%, in three years' time.

The test does not have a direct effect on the valuation methodology, but it does potentially influence the choice of benefit structures.

Monitoring Test 2, however, does not lead to direct action between valuations, but is reviewed by the USS executive with employer representatives on a regular basis.

Test 3: This test provides comfort to the trustee that in the event of an extreme tail risk event, the sector has sufficient aggregate balance sheet assets to cover the benefits promised to existing members. In other words, the aggregate balance sheet assets of the employers should be greater than the deficit in an extreme adverse scenario (e.g. in a one-in-20-year worst case outcome). Given the current position of assets and liabilities this test does not have a direct impact on the valuation methodology.

All three tests are monitored by the executive on a monthly basis and reported to the trustee at each board meeting.

1.c. How does the investment strategy link to the valuation methodology? In particular, the reference portfolio (given that USS defines the reference portfolio each year with a target return of gilts+)?

The reference portfolio is determined by the requirement to deliver the required (target) return determined in the valuation. However, once the target return has been agreed the composition of the reference portfolio will change if the expected returns (and/or covariances) of the various asset classes that make up the reference portfolio change. In other words the way of achieving the target return will change with changes to the long-term outlook for different asset classes. Hence it is important to monitor and review expected return assumptions. Indeed the long term expected return for each asset class is reviewed each year by the Investment Committee, and this together with advice from the scheme actuary is fed into the trustee's deliberations on the expected return appropriate to the valuation in valuation years.

The "reference portfolio" corresponds to a notional set of assets which is expected to deliver the targeted return within the agreed risk tolerance. It is constructed by portfolio optimisation to identify an appropriate portfolio which is risk-return efficient and has the expected return required by the valuation. The reference portfolio is used as the benchmark against which the trustee's investment manager (USS Investment Management Ltd) is measured, both in terms of achieved investment return and the amount of investment risk undertaken. Broadly speaking, USS Investment Management Ltd is charged with outperforming the reference portfolio at a similar, or lower, level of risk. ([Statement of Investment Principles \(USS 2015a\)](#), [Statement of Investment Principles: Supplementary Materials \(USS 2015b\)](#) and the [Statement of Investment Principles \(USS 2016a\)](#)).

The trustee does, as you indicate in your question, express the target return from the reference portfolio relative to gilts. This is an indication of the additional return the trustee expects to achieve relative to a risk free investment in UK government bonds, arising from the additional risk taken on by investing in other asset classes. The trustee chooses to express it in this way for convenience and in order to make explicit the risk premium (i.e. the expected excess return over gilts) targeted by the investment strategy.

1.d. How does USS use stochastic modelling in the valuation process?

USS undertakes stochastic modelling both as part of the triennial valuation process, and also as part of interim monitoring. The modelling involves carrying out a large number of future projections (or scenario paths) for the economic and market environment, as well as the assets, liabilities and contribution requirements of the scheme.

Stochastic modelling is used to consider a range of potential outcomes, summarised by multiple metrics over multiple time horizons, with risk measured in different ways. For example test 2 involves analysing the volatility in valuation outcomes by estimating the probabilities of contributions exceeding the desired levels over a three-year horizon. By contrast test 3 analyses tail risk by measuring the so-called Value-at-Risk (VaR), associated with the deficit and comparing it with the aggregate balance sheet assets of the employers.

Stochastic modelling is also used in monitoring the Financial Management Plan (FMP) where actual outcomes for the liabilities, assets and deficit are compared to the range of projected outcomes across thousands of different scenario paths.

Members of the USS Risk function have considerable experience with stochastic risk modelling and one member was involved in pioneering the use of stochastic risk models in finance in the early 1990s. USS has commissioned a new asset-liability model from ORTEC Finance to allow stochastic projections to be carried out in-house. This provides greater scope and flexibility to perform analyses for the 2017 valuation. For the 2014 valuation this was carried out by Mercer. ORTEC Finance is a Dutch firm that has its origins in Operational Research. The website (<http://www.ortec-finance.co.uk/>) contains a number of technical and applied papers on their work. For members of the VDF who are interested, USS is happy to provide some applied and technical papers on the ORTEC stochastic scenario model that we are using.

2. Prudence

2.a. How are decisions about prudence in specific assumptions decided upon (longevity, inflation, discount rate and so on)? Could information be provided on how the appropriate level of prudence is determined?

See the discussion of prudence in the Methodology Overview. Here we summarise some of the key points.

Decisions on the level of prudence are decided by the trustee acting on the advice of the scheme actuary. Prudence is required in order to ensure:

- An amount of adverse experience can be absorbed by the scheme without immediate recourse to the employers' covenant;
- Compliance with the Pensions Regulator's guidance.

The scheme actuary has identified a number of areas where a best estimate assumption could differ from that adopted in the calculation of the technical provisions, including:

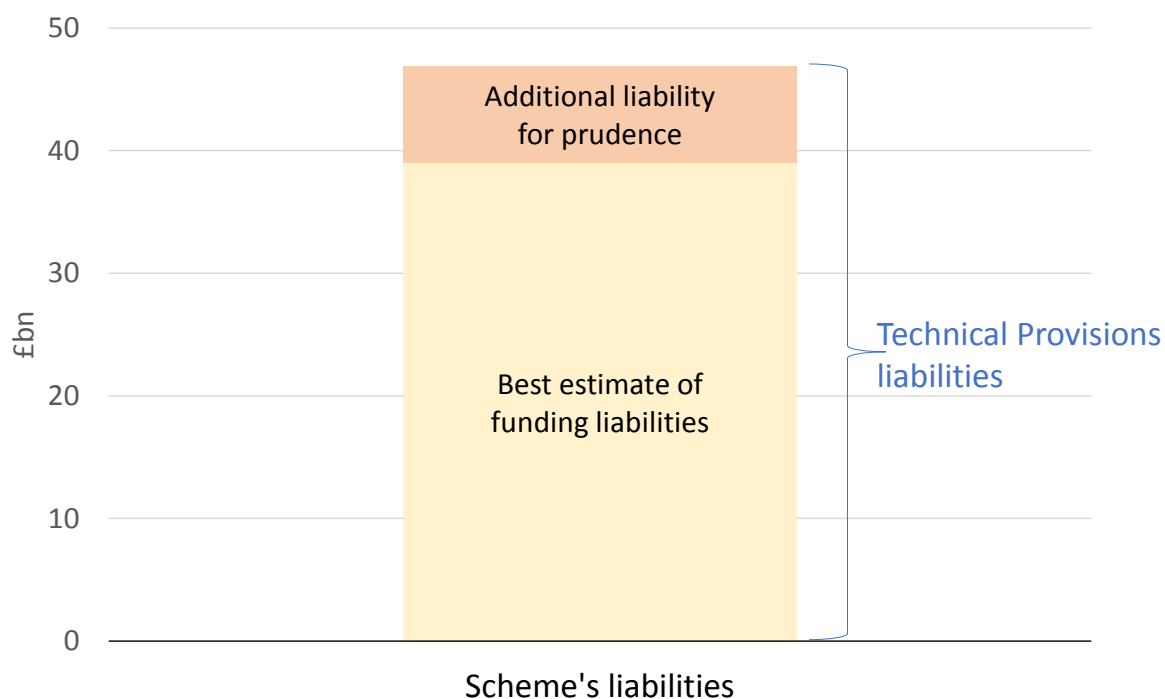
- Discount rate (this is the main area of prudence);
- Gap between RPI and CPI;
- Long-term trend in mortality improvements.

As discussed in the answer to question 1.a, the level of prudence incorporated into the determination of deficit contributions is often lower than that used in determining the technical provisions.

It is recognised that when making assumptions about the future there is one certainty: that the assumptions are more likely to turn out to be wrong than right. As a result the level of prudence within each assumption is generally decided upon by considering the range of potential outcomes and selecting the outcome with the desired level of probability of that outcome being realised in practice.

Figure 2 illustrates the relationship between technical provisions and best estimate for the value of funding liabilities.

Figure 2: Relationship between technical provisions and best estimate of funding liabilities as at 31 March 2014.



The trustee seeks to avoid being overly prudent by ensuring there is a full understanding of which assumptions potentially have prudence built in and the aggregate effect of moving from a prudent assumption to a best estimate. To this end the trustee will conduct a detailed analysis of the assumptions as well as:

- The impact of changing individual assumptions;
- The position of the scheme on a best estimate basis.

The combined effect of all the sources of prudence can be thought of as follows. Whereas the best estimate for the deficit at the end of covenant horizon corresponds to the median or 50th percentile outcome, a prudent estimate might correspond to an outcome at the 60th to 70th percentile, for example. This enables a certain amount of adverse experience to be absorbed without having to go back to the drawing board and redoing the entire valuation, in order to increase contributions. The idea behind this approach is that it should provide a clear perspective on the aggregate amount of prudence associated with financial assumptions.

We refer to the 60th to 70th percentiles merely as an example of what might correspond to an appropriate level of aggregate prudence. The reasoning behind this percentile range is qualitative in nature and is as follows.⁶ First, for a prudent estimate the percentile level should be comfortably more than the 50th percentile, so that there is a lower probability of an adverse outcome than a favourable outcome. By the same token the percentile level should not be so great that the chance of an adverse outcome is tiny, as one might expect in an insurance context, because this would imply an unreasonably large level of prudence in the context of funding a pension scheme.

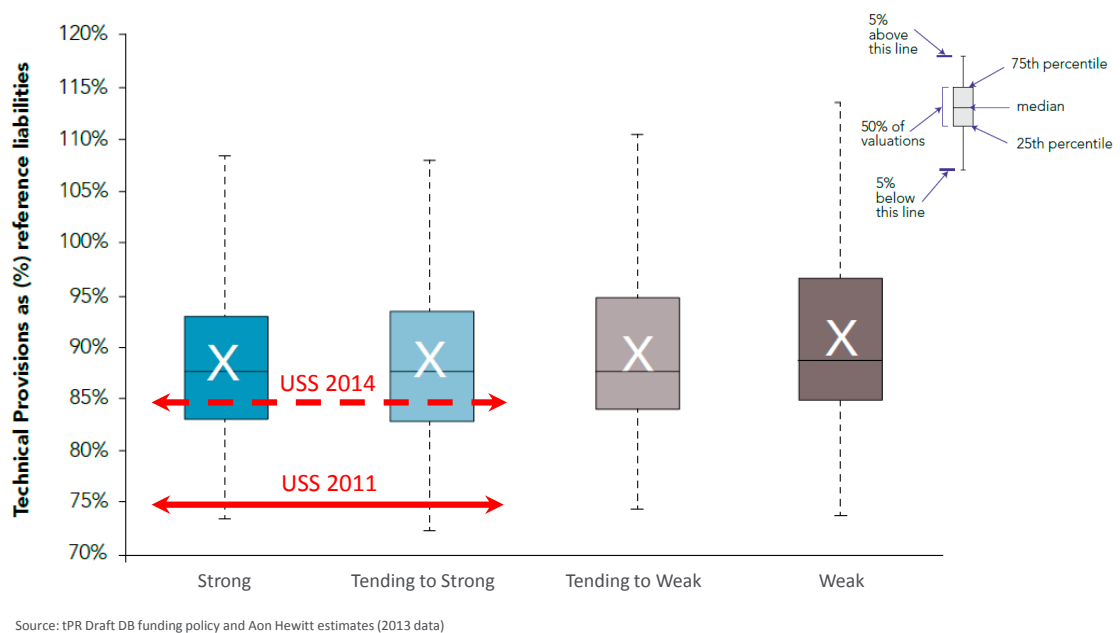
Secondly, this percentile range seems to be consistent with the degree of prudence associated with UK pension funds that follow this approach. While the majority of schemes do not explicitly set the financial assumptions entirely by reference to a specific confidence level, there is a general expectation, which is often articulated in the market, that the discount rate, for example, would correspond broadly to a “two-thirds probability of success”. Thirdly the diagram in Figure 3 which was put together by Aon-Hewitt as part of their review of the 2014 valuation, shows that the level of prudence incorporated into the 2014 valuation is consistent with other pension schemes with strong covenants, lying in the second quartile in terms of the amount of prudence.

The appropriate level of prudence is ultimately a matter of judgement which the trustee makes with the aid of professional advice.

Incorporating some degree of prudence is part of the methodology as opposed to a driver of the methodology. The amount of prudence influences the inputs into the methodology.

⁶ For the avoidance of doubt, an adverse outcome means a requirement to increase contributions, modify benefits, take on an excessive amount of investment risk and/or increase the reliance on the covenant beyond what is acceptable.

Figure 3. The strength of USS technical provisions relative to other UK pension schemes. The x-axis refers to the strength of the covenant. (Taken from an Aon-Hewitt slide).



2.b. What is the approach used by the trustee when deciding on the discount rate to apply to the liabilities and how does that approach relate to market conditions and asset valuation approaches?

The discount rate which is used to calculate the scheme’s liabilities is a reflection of the future expected investment strategy of the scheme, the expected investment return for each asset category that will form part of that strategy and an adjustment for prudence. In developing the scheme’s investment strategy, it is a fundamental requirement to understand the expected return, together with the risk characteristics of that return, for each asset class.

i. The approach to determining expected returns

Recognising that all forecasts are subject to considerable uncertainty, assumptions about expected returns on assets are developed in a process that uses multiple approaches and different perspectives to “triangulate” a self-consistent set of best-estimate forecasts. The approaches include going back to first principles and looking at historical data, the long-term economic outlook, fundamental drivers of returns, different models for future asset valuation and the expected return forecasts developed by other major market participants. As part of this process the USS in-house investment team develops forecasts for returns based on a long-term equilibrium (LTE) model, a central reversion scenario (CRS) model and an implied returns (IR) model.

The process by which the trustee decides on a preferred set of expected returns involves both detailed in-house analysis by experienced investment professionals and a review of the expected returns developed externally by major international consultants and asset managers. This process combines the historical perspective of the risk premia of different asset classes with a forward-looking assessment that takes account of the current economic environment and how it may evolve in the future. Historical data are an important input into

the final decision on expected returns, but the precise connection between the two is indirect and as much a matter of professional judgement as it is economic and statistical analysis. In developing the capital markets assumptions, care is taken to establish a self-consistent set of risk premia for all asset classes (including covariance assumptions), which are also consistent with the outlook for economic growth and inflation. Comparison with the assumptions developed by consultants and asset managers helps check consistency and ensure that the output is reasonable.

The Investment Committee keeps expected returns under review and once a year, as part of their annual process, formally considers whether the reference portfolio remains appropriate in view of the required return target. The scheme actuary also provides a view to the trustee as part of the valuation process. As indicated in the reply to question 1.c above the trustee chooses to express the return relative to that on gilts (i.e. gilts + a risk premium) for reasons of convenience and understanding.

The "+" part of the methodology refers to two things: (i) it is used to express the relative return on assets relative to gilts (as described in 1.c) and (ii) it is used to express the discount rate for liabilities.

The "+" part of the discount rate, as detailed above, is derived by considering the following:

- The current and future investment strategy (i.e. the asset allocation);
- The expected return on this investment strategy;
- An adjustment for prudence.

ii. *Triennial valuation methodology vs. monitoring methodology*

At each valuation the trustee undertakes a full review of all the inputs upon which the valuation is based. This includes development of an up-to-date set of expected returns for each asset class, starting from first principles and taking account of the prevailing economic situation and outlook. These expected returns are expressed in terms of risk premia over a risk-free investment (i.e. gilts). The inputs adopted reflect the trustee's view at that point of time. As such, the "+" should be seen as variable at each valuation. This reflects both market expectations of returns on assets and investment strategy (and in respect of the discount rate, current judgement on prudence).

For tracking and monitoring purposes the "+" part, i.e. the risk premium over gilts adjusted for prudence, is generally kept fixed between valuations. Typically it is only as part of the triennial valuation process that the risk premium over gilts for each asset class would be adjusted. However, as indicated above, the trustee does monitor expected returns on assets, and with the Investment Committee, undertakes a formal review once a year. The trustee could decide to adjust risk premia between valuations for monitoring purposes. Moreover, the trustee has the ability to bring forward a valuation if it is believed that the investment outlook has changed substantially.

It is important to emphasise the distinction made above between the methodology for *monitoring* (in which certain input assumptions generally remain fixed) and that for the *valuation and setting contributions* (which involves a comprehensive reassessment of *all* the input assumptions from first principles).

Furthermore, it should be noted that the purpose of monitoring the scheme between valuations is to reflect progress against the scheme's FMP and indicate where attention may be required at the next valuation. No decisions are ever made to change contribution rates or benefits on the basis of the monitoring data alone. Such decisions will always be based on a full valuation in which all inputs are comprehensively reviewed.

iii. Example of estimating expected returns

Without wishing to put undue weight on any one approach, we can illustrate the types of models used in the triangulation process for expected returns with a particular example. Equity returns can, for instance, be forecast based on a number of fundamental drivers, including economic growth, earnings growth and particularly that part accruing to shareholders (having accounted for dilution via share issuance and differential earnings growth of listed vs. unlisted or pre-listed companies) and taxation. An alternative approach, which generally gives a similar result involves basing the forecast on dividend yield, real dividend growth, inflation, the inflation risk premium and the dilution effect. Obtaining similar results from different approaches is a crucial element of the process of triangulation.

For fixed income assets (e.g., gilts and corporate bonds) the process is different. Expected returns on fixed income assets can be estimated from forecasts of future yield levels. The expected return on a long-maturity corporate bond, for example, will be determined by the aggregate impact of three components. The first is the regular income coming from the receipt of the bond's coupon, which is determined by the market yield at the time the bond is issued. The second component is the capital appreciation or depreciation coming from the change in yield since issuance, taking into account the reduction in the bond's maturity with the passing of time. The final component is a downward adjustment for the expected loss due to default, reflecting the credit quality (default probability) of the issuer. The key element in this process is the forecast of the future yield level, which can be performed in different ways.

As part of the 2017 valuation process the expected return assumptions and the absolute level of the discount rate will be discussed with stakeholders.

2.c. What is the gilts part of the methodology representing and how is it set? For example is quantitative easing taken into account? Is an average duration set and a single point on the yield curve used, or are liabilities valued by matching annual cash flows to points on the yield curve? How is the approach selected to give what the trustee considers an appropriate figure?

Current gilt yields are an input into the current value of technical provisions at any point in time. But how the technical provisions evolve in the future depends on the future path of gilt yields. Both are important in determining the required contributions and the amount of investment risk that can be taken in the scheme.

The current gilt yields used in the methodology are those observed in the market. These yields are not adjusted for events or programmes such as quantitative easing. This is because these events and programmes are already reflected in prevailing market yields and prices. However, such actions are reflected in the differences in the expected returns on other asset classes.

Current gilt yields reflect current market prices for financial instruments that are traded between willing counterparties: they are a matter of fact. By contrast, future gilt yields –

whether they are derived from market-implied breakeven forward yields or otherwise – are forecasts that are not guaranteed to be realised. The future path of yields is uncertain and must be forecasted in a way consistent with the future evolution of the economy and other asset returns (see the answer to question 2.b.).

Moreover, the expected return on a portfolio of gilts is a deterministic function of the initial (current) gilt yield, the expected future path of gilt yields and the portfolio's rebalancing strategy.

In terms of estimating future gilt yields, we need to consider the possibility that yields may remain low for some considerable time. In particular, the valuation should certainly not rely too heavily on yields quickly "reverting" to the high levels seen in previous decades.

The 2014 valuation expressed the discount rate by reference to a "single equivalent" average gilt yield, which was based on the gilt yield curve weighted by the projected cash-flow profile of the scheme's liabilities. For the 2017 valuation our proposal would be to use the full gilt yield curve rather than the single equivalent yield. A single equivalent yield may still be evaluated for convenience of communication, but we would envisage using the full yield curve term structure for the calculations.

3. Consistency

3.a. The Forum would like to understand further the consistency between the methodology and assumptions used in the valuation.

The methodology sets out:

- What factors are considered in the valuation;
- How inputs are set;
- The calculations that need to be undertaken.

The implementation of the methodology requires inputs (so-called assumptions) which:

- Are based on best estimates;
- Are internally consistent;
- When considered overall have an appropriate level of aggregate prudence.

These inputs are based on data, but informed by judgement to take account of the fact that the past is not necessarily a guide to the future.

The trustee looks to its professional advisors for confirmation that the assumptions they adopt fulfil these requirements. The full set of assumptions will be published along with both the methodology and the results of valuation.

3.b. As noted above how does the reference portfolio align with the current gilts + approach and the investment strategy more widely?

As outlined above in the response to 1.c. the reference portfolio is the investment strategy that is expected to deliver the targeted investment return within acceptable risk parameters. The trustee chooses to express this target return in the form of gilts + for the reasons described above.

The reference portfolio is used to monitor the performance of the scheme's investment manager (USS Investment Management Ltd) and to constrain the amount of investment risk that the investment manager can take.

The reference portfolio at any one time reflects a benchmark investment strategy designed to deliver the target return, which in turn allows the discount rate to be established. The fact that the target return is expressed as gilts + as opposed to say a total percentage return number does not alter the process required to establish the target return. Furthermore, with the expected return and discount rate both expressed in this way, it makes it very straightforward to compare the two.

4. Volatility

4.a. It is important to understand where volatility affects the valuation outcomes over consecutive valuations. Does the current methodology add to volatility?

Volatility in the outcome of consecutive valuations, and in particular the volatility in the required contribution rate emerges from the entirely necessary updating of inputs and realised performance on the basis of new information. There are, specifically, two types of sources of this volatility:

- The difference between the realised experience and the forecast experience between the two valuations;
- The changes in prevailing market conditions and future assumptions between the two valuations.

The first of these includes the difference over the intervaluation period between expected and realised outcomes for: investment returns, inflation, mortality, membership profile, etc. The second includes changes in forecasts for future investment returns, inflation, mortality, etc., as well as the change in the discount rate. Note that the discount rate can have a very large impact, but it is not always the most significant source of volatility between valuations, as its impact can be offset by changes in inflation. As a result, it is changes in the real discount rate that often have the greatest impact.

The above sources of volatility manifest themselves in terms of:

- Changes in the technical provisions (accrued liabilities). (This is a function of both experience during the intervaluation period and changes in inputs, i.e. assumptions, between the valuations, the most important of which is the real discount rate).
- Changes in the cost of future benefits. (This is driven in the main by the changes in assumptions between the two valuation dates and, to a lesser extent, changes in the

member population. Again the most important driver is the change in the real discount rate).

- Changes in asset value. (This is driven by investment returns over the intervaluation period).

Collectively, these changes translate into changes in the required contributions. However, the link with contributions involves a complex non-linear relationship (see below).

i. Changes in technical provisions

The two major sources of volatility in technical provisions are changes in assumptions, and the impact of actual members' experience relative to what was assumed at the previous valuation. We consider each of these in turn.

Sources of volatility arising from members' experience during the intervaluation period can arise from:

- Salary increases;
- New members joining the scheme;
- Retirement patterns;
- Mortality experience;
- Withdrawal experience (including transfers out);
- Benefit increases (over and above promised pension increases).

Historically the largest source of experience-related volatility has come from salary increases due to the former final salary nature of the scheme. The changes to scheme benefits implemented on 1 April 2016 significantly reduces the salary increase effect. Volatility from the other experience-related sources tends to emerge gradually over time and these effects are typically picked up in the modification of assumptions from one valuation to the next. For example, realised mortality rates are typically very close to assumed mortality rates, as longevity risk is a trend risk that emerges over longer periods of time.

Changes in inputs from one valuation to the next can include both financial and demographic variables. When monitoring the position between valuations, typically only financial inputs (and in particular the real discount rate) are updated for prevailing market conditions, with all other assumptions kept at the values set at the previous valuation. As part of each valuation, however, all inputs and the principles underlying them are reassessed, as described above in the answers to questions 1 and 2. The change in the real discount rate – reflecting changes in market conditions, investment strategy and expected returns for different asset classes – typically gives rise to the majority of the volatility in technical provisions.

ii. Changes in the cost of future benefits

Changes in the cost of future benefits from one valuation to the next are driven largely by the changes in assumptions reflecting the changes in outlook between the two dates, with the change in real discount rate generally being the most significant.

At each valuation, the cost of future benefits, like the technical provisions, is dependent on the assumptions agreed together with the prevailing profile of the scheme's active members, in particular their age, gender and salary distribution. At each valuation assumptions must be made on how stable the member profile is (including assumptions around the levels of leavers

and new entrants). If these demographic assumptions change from one valuation to the next, this introduces additional volatility in the cost of future benefits. For the purpose of monitoring the cost of future benefits during the intervaluation period, however, adjustments are typically made only for changes in financial conditions, not demographics.

Changes in benefit structure also clearly affect the level of technical provisions and the cost of future benefits.

iii. Changes in asset value

The valuation methodology uses long-term expected returns on assets as an input. To the extent that the actual intervaluation experience differs from the long-term assumptions, there will be variance between the expected value of the assets and the actual assets held, which will translate into volatility in the deficit recovery component of the contribution rate.

iv. Impact on contribution volatility

The contributions determined at each valuation can be decomposed into two distinct parts: a component for deficit recovery and a component to cover future benefits. Changes in the technical provisions, the cost of future benefits and asset values discussed above will clearly lead to changes in the contributions required from one valuation to the next. But there is one further source of volatility in the contributions, namely, changes to the length of the recovery period. Since both the deficit and the cost of future benefits are met by a combination of contributions and investment returns over the recovery period, the annual contribution rate is highly dependent on its length.

In conclusion, because both the methodology and assumptions are fully reassessed as part of each valuation on the basis of up-to-date data, the trustee does not take the view that the volatility from valuation to valuation is “unnecessary”. Certainly the trustee does not believe that a desire to avoid such volatility should be a driver of the valuation methodology.

It is certainly possible to remove the majority of volatility from the valuation without changing the valuation methodology, but only with a significant increase in the cost of running the pension scheme, as this would involve shifting to a fully gilt-based investment strategy.

4.b. How does the trustee manage volatility of market based inputs? Considering the relationship between gilt yields and market movement, what is the short and long term view of USS towards the volatility within the approach, particularly given the current market conditions?

Market movements affect the financial position of the scheme and, in line with the Pensions Regulator’s requirements, these movements are reflected in the results of each valuation and the interim monitoring. As noted in 2.c. above, the trustee does not take a different view on gilt yields from that observed in the market. As such the trustee’s view is that changes to market based inputs translate into changes in our estimate of the cost of providing pensions.

Note that changes to non-market based inputs to the valuation, such as assumptions about long-term expected asset returns, outperformance in the recovery plan and the length of the recovery plan, involve explicit trustee judgment to overlay market-based changes and smooth contributions.

For example, the trustee reassesses its expected return assumptions at each valuation and does not believe the relationship between the expected returns on gilts and the expected returns on other asset classes will necessarily remain fixed. Between 2011 and 2014 these and other changes were taken into account in the construction of the 2014 reference portfolio, which differed from its 2011 predecessor.

Because any change in deficit between valuations is effectively spread over the deficit recovery period and the length of this recovery period along with the assumptions used in that period will also be subject to review, *there is no simple direct relationship between volatility in the deficit and volatility in the contribution requirements.*

The trustee's current FMP following from the 2014 valuation includes a trigger-based derisking framework (called the Journey Plan), which will allow it to take advantage of opportunities to derisk the scheme's investment strategy (towards its agreed 20-year target investment strategy) without increasing contribution requirements. This relatively modest amount of derisking was built into the 2014 valuation so that the level of risk in the scheme should not increase over time, a requirement for consistency with Principle 1, that reliance on the covenant not increase (see answer to question 1.b.). The derisking is defined in terms of gradually reducing the risk associated with the scheme's deficit and manifests itself in terms of gradual adjustments to the investment mix, and hence the expected return of the reference portfolio.

The trustee's investment strategy also incorporates diversification and hedging of certain unrewarded risks, in order to reduce the impact of market volatility on the scheme's funding position. For more discussion of these issues see [USS 2015a](#), [USS 2015b](#) and [USS 2016a](#).

We hope that the above provides the required clarification on the approach to the valuation methodology.

References and Bibliography

The following documents can be found on the USS website:

[USS \(2014a\), An Integrated Approach to Scheme Funding \(July\).](#)

[USS \(2015a\), Statement of Investment Principles \(July\).](#)

[USS \(2015b\), Statement of Investment Principles: Supplementary Material for Institutions \(July\).](#)

[USS \(2015c\), Schedule of Contributions \(July\).](#)

[USS \(2015d\), Actuarial Valuation as at 31 March 2014: Statement of Funding Principles \(July\).](#)

[USS \(2016a\), Statement of Investment Principles \(April\)](#)

[USS \(2016b\), 2017 Valuation Timetable \(November\).](#)