

Evaluating the fairness of the Risk-Based Levy

A report prepared for the PPF by David McCarthy and Anthony Neuberger

July 2010

0. Table of Contents

Introduction	p. 3
The model	p. 8
Model calibration	p. 13
An absolutely fair levy	p. 17
The proposed levy formula	p. 21
Conclusion	p. 25
Tables	p. 27
Figures	p. 37

1. Introduction

We have been asked by the Pension Protection Fund to “support them in the development of a fairer risk-based levy” (RBL). In particular we have been asked:

- “to construct a theoretical measure of the cost of future pension scheme claims that will serve as a benchmark for measuring fairness and demonstrate that this theoretical levy is actuarially fair.”
- and “to provide a paper for inclusion in the consultation package setting out the case for using the theoretical approach identified, work done, and implications for the new formula.”

The PPF Steering Group¹ defines a fair levy as follows: “fairness is achieved if the bill of each scheme matches the cover it receives.” While the broad principle is clear, the detailed implementation is not straightforward.

One key issue is the timescale over which fairness is to be achieved. Fairness on an annual basis – with the annual levy bill for each scheme matching the insurance cover it receives over the following year – is conceptually straightforward but unworkable in practice. Claims are rarely completely unexpected; sponsors tend to fail after an extended period of financial weakness. If levies for each scheme were set equal to the chance of their failing over the next twelve months, most schemes would pay little as the chance of their sponsor failing in the short term would be very small. On the other hand, schemes with a substantial failure risk would be required to pay a lot. But these are precisely the schemes that cannot afford to pay. Putting the bulk of the burden of financing the Pension Protection Fund on the very weakest schemes makes no sense: it would help bring about those very scheme failures that the PPF is there to protect pensioners from, and it would not actually provide the resources the PPF needs to be solvent.

The levy is necessarily capped, and those schemes that are capped generally get more cover than they are paying for. Someone has to pay for this. With the PPF having to cover its costs, this means that other schemes have to bear the costs of providing cover to weaker

¹ “Future Development of the Pension Protection Levy - Feedback from the Steering Group”, March 2010, para 4.2.3.

schemes that cannot afford to cover themselves. We have given careful attention to the issue of how this burden can be fairly borne.

One possibility is to apportion the burden evenly across schemes through the scheme-based levy. This has a rough and ready attraction. All schemes benefit from insurance, and it is in every scheme's interests that levy payments are capped so that the levy does not itself precipitate scheme insolvency. No scheme can be sure that it will not, some time in the future, be the beneficiary of a cap. Credit ratings can deteriorate rapidly, and the sponsor who is investment grade today may find themselves in financial distress tomorrow. All should pay towards the cost of the cap from which all benefit.

However, such a solution involves some degree of unfairness. We will review the evidence, but whether one looks at actual default rates or the cost of credit in the Credit Default Swap market, the message is clear. High-rated borrowers are not only much less likely to default than low-rated borrowers, but those differences are highly persistent. The difference in buying credit protection for five years between an AA and a BB rated company has averaged about 2.5%. If you want to extend that protection on the same two borrowers for a further 5 years, the difference remains about the same. The AA rated sponsor is very much less likely to benefit from the cap in the foreseeable future than is the BB company.

We therefore take the view the fair levy should seek to match the present value of premiums and claims over the long term, not on a year-by-year basis. The uncapped scheme would pay a levy that reflects not only the claims it is likely to make in the short term, but also the possibility that its financial health may deteriorate in the future and that at that stage it may be capped, and pay less than it is getting in cover. With such a policy, in steady state the PPF will be making a notional profit on those who are uncapped which will offset the loss they expect to make on those who are capped.

In this paper we concentrate specifically on fairness, and more particularly on the fairness of the risk based levy. However, fairness is a loose and broad concept, and in order to provide a quantitative measure of fairness we need to define it narrowly and precisely. But it is important to recognise that fairness is not the only criterion by which the levy structure needs to be judged. There are practical concerns (the levy charged to a scheme must be based on scheme data that can be collected for all schemes at reasonable cost and are verifiable; it cannot be excessively complex; levy parameters should be stable over time; levy payments for individual schemes should be stable and schemes should be able to

predict their payments over time.) There are issues of incentives: schemes are likely to adjust their behaviour in response to the levy structure. Indeed, one of the arguments for fairness is that it forces schemes to internalise the costs of failure, and gives them an incentive to reduce them. But as we will see, fairness and proper incentives can conflict. There are issues of solvency: the levy must give the PPF sufficient revenue to pay the pensions of failed schemes, so it needs to ensure that the way the burden is spread is not only equitable but also sustainable. We have interpreted our remit as being restricted to defining and measuring fairness. While we note where there may be a conflict between fairness and other objectives we do not attempt to construct any kind of trade-off between them.

The definition of fairness we adopt is straightforward: in a perfectly fair scheme, the present value of levy payments that a group of schemes can expect to pay over a reasonably long period should equal the present value of the expected claims that they will impose on the PPF over that time. For various reasons we will discuss, we have chosen a period of 10 years. Turning this into “a benchmark for measuring fairness” is difficult and controversial. The main problems that arise come under the following heads:

- the definition of fairness;
- controllable and uncontrollable factors;
- chronically weak schemes;
- the levy inputs;
- modelling the sponsor’s financial strength;
- modelling the scheme’s funding level;
- accounting for risk.

The definition of fairness

The definition refers to groups of schemes, rather than individual schemes. No levy system can hope to reflect all individual circumstances. In our analysis we focus particularly on schemes classified by sponsor financial strength, current scheme surplus or deficit, and current investment mix.

The definition does not attempt to look at the PPF in aggregate. In particular it follows from our definition of a fair levy that two schemes with employers of similar strength and with similar funding levels would pay the same levy per pound of liabilities irrespective of the

fact that one is very much larger than another. The fact that the failure of the larger scheme could have material impact on the finances of the PPF is not taken into account given the definition of fairness that we have adopted.

We also do not take into account differences that are not observable. While this sounds like a truism, it is important in considering the levy paid by two schemes that start off from the same point but have very different funding strategies. Because we are interested in lifetime claims, the difference in strategies (how the sponsor will respond if its own financial strength declines, and how the investment strategy of the scheme will change) are very important in assessing likely future claims. But the levy can only take account of what actually happens rather than of the intentions of sponsors and trustees, so the ability of the levy to be fair as between two schemes with different strategies is very limited.

Controllable and uncontrollable factors

The risk-based levy is bound to depend on two things: the funding of the scheme and on the strength of the sponsor. But in comparing the fairness of levy structures in these two dimensions, it is important to be aware of the important differences between them. By and large a sponsor can choose to inject more funds into a scheme and increase its solvency. It is likely to be much harder to do much about its own financial strength. If the risk-based levy overcharges the more under-funded schemes, the sponsors can mitigate the effect by improving funding. If the levy overcharges the weaker sponsor, there is not much the employer can do about it. This suggests that fairness in the sponsor financial strength dimension is more important than fairness in the scheme funding dimension.

Chronically weak schemes

Some schemes are so weak that they cannot afford to pay a levy equal to their expected future claims. To reflect this limit on the ability to pay, there is a cap on the maximum levy rate per pound of liabilities. These capped schemes in effect benefit from a subsidy. There is no levy structure that can be fair between schemes that are capped and those that are not. But as we have argued, for the great majority of schemes that are uncapped, the existence of the cap has to be reflected in the levies they pay if the levy is to be fair. The premium they pay while uncapped must reflect the subsidy they will get if their finances deteriorate to such an extent that they are capped.

Levy inputs

In this paper we assume that the levy will be based on a few simple parameters for each scheme, notably: the value of liabilities and of assets in the scheme, a measure of the financial strength of the employer, and a measure of the riskiness of the asset mix. In principle there are other parameters that could be used (the time profile of liabilities, whether the scheme is open or closed, the size of scheme relative to the size of the sponsor) but they have not as yet been used in setting the levy. The approach we develop here could be extended to include them.

Modelling the sponsor's financial strength

Once it is accepted that fairness is measured over the scheme's life, we need to model the ways in which the scheme may evolve over time. We describe below the model we actually use. But it is important to emphasise that the models used, however sophisticated, are crude representations of reality. We need to be able to describe how financial strength varies over time. There is data from the ratings agencies extending back over many years and over different markets and types of company. More recently, Credit Default Swaps (CDS's) have been created, giving a market-based measure of default risk. But the picture we can observe varies over time and across countries and industries. The future experience of the PPF will no doubt be different again. There are several specific reasons why one should treat the historical record with caution. The Dun & Bradstreet failure score does not map exactly on to the corporate credit ratings that are used to predict long term ratings transitions. The populations of the PPF universe and credit-rated corporations are not very similar. The procedures for reviewing, appealing and revising D&B scores are very different from anything that applies with the standard corporate credit ratings.

Modelling the scheme's funding level

We also model the way in which deficits grow or decline. While in the short term the behaviour depends upon asset returns and we have a moderately good understanding of the volatility of different asset classes (though even here our knowledge is not very certain), in the long term the behaviour will depend much more on investment strategy and on contributions policy. These are even harder to forecast and describe. They depend heavily on regulation and indeed on the design of the levy itself. We have no good theory to explain how well employers fund their schemes, or the investment strategy adopted by trustees. In the absence of good theory predictions of future behaviour are little more than educated guesses.

This does not make the model useless. It is better to have a rough measure of fairness than not to have one at all. Many of the conclusions that come out of the model are robust to changes in the input parameters. The model is also useful in exposing and clarifying issues. By making the definition of fairness explicit and by showing how fairness is assessed, we are creating an approach which others can modify or adapt, by redefining what is meant by fairness or by incorporating different modelling assumptions or parameters.

Accounting for risk

The projected life of a pension scheme is many decades. In computing the present value of claims and levy payments some account must be taken of risk. In particular claims on the PPF are most likely to occur in those states of the world – when assets prices slump and company failures abound – that are most expensive to insure against. We use modern financial theory to take account of this, discounting good states of the world more heavily than bad states of the world. This has the effect of increasing the present values of both levy payments and claims, both of which become larger when the world is in a bad state, but the effect on claims is much larger, precisely because many more claims occur in bad states of the world.

We now turn to explaining in detail the methodology we have chosen for measuring fairness.

2. The Model

As we have argued, any model of levies and claims is bound to be at best an approximate representation of reality. At least then the model should have the merit of simplicity so that it is transparent, and the number of parameters to be estimated should be kept small. But it also needs to model the key dynamics that determine the evolution of levies and claims.

It should model the broad features of both the evolution of pension scheme funding and of the financial strength of sponsors. In particular it needs to recognize that schemes have to make good deficits over a period of time, that the financial strength of sponsors varies both idiosyncratically, and also with the state of the economy, that there is substantial asset liability mismatch in most pension schemes, and that the size of pension scheme deficits is correlated with the probability of schemes failing.

It should capture the interaction between sponsor failure and scheme deficits. This is important because the correlation between the two greatly increases the expected claim size (in bad states of the economy, sponsors are more likely to fail, and the deficits in their schemes are likely to be unusually large). The correlation also increases the effective burden on the PPF because it increases the claims of the PPF in poor states, which are expensive to hedge, and decreases them in good states where hedging is cheap.

The main use of the model is to show the fairness of different levy structures. For any particular set of model parameters and premium design, and for any initial population, the present value of future claims can be compared with the present value of future risk-based levy income for particular groups of schemes. The resulting claim/levy ratio can be used to test the fairness of how the levy takes account of the strength of the sponsor, the solvency of the scheme and the scheme's asset mix. The model could also be used to see how differences between schemes that are not currently reflected in the levy structure (for example the maturity of the scheme, whether it is open or closed to new entrants, different funding strategies) lead to differences in the claim/levy ratio.

Our model differs from the Long Term Risk Model (LTRM) developed by the PPF in that it is not a simulation model which is designed to test the solvency of the Pension Protection Fund. It does not focus on the impact of the failure of very large schemes. It assumes that the levy formula is held constant over time, and ignores the effect of any counter-cyclicality, so PPF income rises as sponsors' finances deteriorate and deficits rise. It assumes that the level of protection provided by the PPF remains constant.

In fitting a highly stylised model, approximations are inevitable. One way of getting a sense of how important these are is to do a sensitivity analysis. To the extent that conclusions drawn from the model change if different plausible parameters are used, the conclusions should be treated cautiously. Conclusions that are robust to changes in parameters may still be unreliable if they depend on systematic features of the model that are incorrect.

The model does not take account of the way the scheme finances the PPF premium. Where the premium is small, or the sponsor is strong, this is not likely to be important. But if the scheme is large relative to the sponsor, and the premium is large, the payment of the premium itself may weaken the sponsor and either make a claim more likely, if the premium is picked up by the employer, or may reduce the funding of the scheme and so make any eventual claim larger. Thus the model will tend to understate the claims/levy ratio for the

weakest schemes. But in any case, as we have already argued, the question of fairness only arises with those schemes that are in a position to pay a premium equal to the value of their expected claims, so this is not likely to be an important limitation of the model.

We now discuss the main features of our chosen model, starting with the model we adopt for scheme sponsor solvency and pension scheme deficits, then present how we evaluate the present value of levies and claims and finally examine how the model is calibrated.

Modelling the failure of the sponsor

Regardless of the funding position of the scheme, claims only arise when the sponsor fails. So a critical part of the model is the modelling of sponsor failure. Essential features we need to be able to model are that:

- sponsor failure can either come out of the blue, or could follow a period of deteriorating financial strength;
- financial strength varies over time with some regression towards the mean;
- sponsor failure is correlated across sponsors and also tends to be correlated with the performance of asset markets (and hence pension plan deficits).

Our model is based loosely on a paper by Pierre Collin-Dufresne and Robert Goldstein.² We use an abstract measure of the financial condition of the sponsor, which we call financial strength, denoted by s . We assume that when s , which is positive for solvent firms, falls below zero, the company fails. To model both blue-sky and predictable defaults, we assume that s follows a jump-diffusion process with a mean-reverting diffusion and an exponentially distributed jump downwards. We assume that the diffusion component is correlated across firms, which can be conceived of as a common market component to the change in the value of the assets of the firm. The model can be expressed algebraically as:

$$ds = [\mu_s + \kappa(\bar{s} - s)]dt + \sigma_s dz_s - Jdw \quad (2.1)$$

where s is the financial strength of the sponsor, \bar{s} is the long term average financial strength, κ is the speed of mean reversion, σ_s is the volatility of the diffusion component of the change in financial strength, z_s is a standard Brownian process, J is a random variable which has a negative exponential distribution with mean value ξ , w is a Poisson process with constant rate λ , and μ_s is the drift in the financial strength.

² "Do Credit Spreads Reflect Stationary Leverage Ratios", *Journal of Finance*, Volume LVI, No 5, October 2001, 1929-1957.

On a risk adjusted basis and for a sponsor in steady state, we assume that the financial condition of the sponsor would on average remain unchanged, and so μ_s would equal $\xi\lambda$ to exactly offset the expected effect of jumps on the financial condition of the sponsor.

But on an unadjusted basis, μ_s would equal the risk premium on firm assets plus the offset due to jumps. We assume that investors do not observe the financial strength of sponsors continuously. This means that they are unable to differentiate between the change in financial strength due to jumps and that due to the diffusion. For this reason, we assume that investors expect the same reward for assuming jump risk as they do for diffusion risk, and that this reward is a linear function of the instantaneous variance of the combined process. Hence, the unadjusted drift in the financial condition of the sponsor is given by

$$\mu_s = \lambda\xi + \eta\rho\sqrt{\sigma_s^2 + 2\lambda\xi^2} \quad (2.2)$$

where η is the Sharpe ratio of risky assets in the economy (which we assume is the same for all firms), and ρ is the desired correlation between the financial strength of the sponsor and the market portfolio, which we cover in the next section.

From a modelling perspective μ_s only enters into results in association with the target funding level; the key parameter to estimate is $\bar{s} + \mu_s / \kappa$, rather than \bar{s} and μ_s separately. Since we have defined the default threshold for financial strength to be 0, we define the initial financial strength of companies in different ratings bands and the long-run average financial strength, \bar{s} , relative to that point.

Modelling the pension plan deficit

There are key features of the behaviour of pension plan deficits that we need to model:

- deficits fluctuate over time, with some tendency to revert to a long term mean, as contribution rates are adjusted to move towards target funding levels;
- since pension fund assets and liabilities typically are not well matched, and funds tend to hold similar types of assets, there is a strong common factor in the movement of funding levels across schemes;
- the market factors that affect asset values also affect the financial strength of sponsors, so there is a correlation between scheme funding levels and sponsor financial strength.

The relevant measure of the funding level is the market value of scheme assets divided by the market value of PPF liabilities, which we denote by f . This is the measure that determines the size of the claim on the PPF if the sponsor fails. We assume that f follows this process:

$$df = [\mu_f + (\bar{f} - f)/T]dt + \sigma_f dz_f \quad (2.3)$$

Here \bar{f} is the target funding level. We assume that any difference between the target and the current funding level is amortised over T years. μ_f is the drift in the funding level. On a risk adjusted basis and for a scheme in steady state, the market assets and liabilities should grow at the same rate, so μ_f would be zero. But on an unadjusted basis, with liabilities being bond-like and expecting to grow at the interest rate, and assets being held as risk-bearing securities with a higher expected return, $\mu_f = \eta\sigma_f$, where η is the Sharpe ratio of risky assets in the economy and σ_f is the random variation in the funding level. which we assume is entirely due to the mismatch between assets and liabilities. σ_f will tend to be higher the more the scheme is invested in equities, but even a scheme heavily invested in bonds will have substantial mismatch due to longevity, real and nominal wage inflation and other factors.

The last model parameter that needs to be mentioned is ρ , the correlation between changes in a firm's financial strength and changes in the plan's funding ratio, which captures the common factors in sponsor financial strength and scheme funding. When the market goes up, asset values will rise, f will increase, and the sponsor's financial condition will strengthen, so ρ is likely to be positive. If we wish the overall correlation between changes in financial strength and changes in the funding ratio to be ρ , we need to assume that the correlation between dz_s and dz_f is given by

$$\rho^* = \rho\sqrt{\sigma_s^2 + 2\lambda\xi^2} / \sigma_s \quad (2.4)$$

to reflect the fact that the jump process and the distribution of jump size are both uncorrelated with either of the diffusions.

Finally, we can also readily accommodate within the model an upper limit on solvency. This reflects rules limiting overfunding. Equation (2.3) is modified by imposing a reflecting barrier at some suitably high level f^* to reflect the fact that highly-funded schemes are forced to reduce funding by HMRC. Although we have not done this, it would be feasible to use our framework to examine static investment strategies which change over time, such as

the situation where schemes which reach a very high level of solvency elect to buy out their entire liabilities with an insurance policy and therefore leave the PPF.

Modelling the present value of levies and claims

At any point in time, the state of each scheme can be measured by just two variables conditional on its future cash flows and the other model assumptions: the funding level of the scheme (f) and the strength of its associated employer (s). Since both levies and insurance claims are simply streams of payments whose size depends on the future path of a scheme's funding level and its sponsor's financial strength, they can be viewed – and valued – as a portfolio of options on two underlying securities, which we assume are both traded: the financial strength of the sponsor and the assets in the pension plan.

We use a recombining lattice (which is a standard technique used in options pricing), working backwards from the date the scheme terminates (because all liabilities have been paid out, or because the sponsor fails) or from a chosen horizon date, to compute the present value of expected future levy payments and expected future claims. Present values are calculated using either real world probabilities (to compare the expected claims and levy payments), or risk adjusted probabilities (to compare the market value of future claims and future levy payments).

3. Model calibration

The model is highly stylised, and calibration is critical. As described, the model applies to a single scheme with a specific sponsor. In modelling a set of schemes we must specify which of the parameters are common to all schemes and which are specific.

We assume that the following parameters are common to all schemes:

- the mapping between financial strength and credit rating; in particular we assume that the default level is the same for all sponsors;
- dynamics of leverage (so common values for $\mu_s, \kappa, \sigma_s, \xi$ and λ);
- dynamics of scheme funding (the amortisation period of scheme deficits and surpluses (T), the mean funding level (\bar{f}) and the correlation between the deficit and leverage (ρ));
- the time profile of their liabilities.

We assume however that schemes differ in the current funding level (f) and financial strength (s) of the sponsor and in the volatility of the funding ratio (σ_f). We assume that the expected return on assets is linearly related to the volatility of the funding ratio – schemes mismatch assets and liabilities deliberately in the search for additional return. This assumption affects the expected size of claims and payments, but does not affect the price of them, as the higher expected return on these assets is offset by the use of a higher discount rate.

Calibration of leverage assumptions

We calibrate our model of sponsor solvency to fit the historical data. There are, broadly speaking, two approaches to doing this. One is to match prices in the bond market, while the other approach is to fit historical default rates.

Corporate bond prices, and specifically the credit spread or difference between the yield on the corporate bond and the yield on government bonds, shows the price that the market charges to take the risk of corporate default. The credit spread is a forward looking measure, and reflects not only the market's assessment of the risk of default but also the premium that the market requires for taking that risk. It therefore provides a direct measure of the risk adjusted default probability. There are two issues that need to be addressed in going from credit spreads to default probabilities. First, the credit spread reflects not only the probability of a default, but also takes account of the recovery rate on the bond if the bond does default. To back out a default rate from the credit spread, it is necessary to assume a specific recovery rate. Second, there is substantial evidence³ that credit spreads, and particular those of higher rated bonds, also reflect the fact that corporate bonds are much less liquid than government bonds. In an attempt to strip out this liquidity effect, we use Credit Default Swap (CDS) rates rather than bond spreads.

The second approach we use for calibrating the model is to fit it to historical default rates. This is more appropriate when we are using the model to look at objective probabilities. In many ways this is simpler than using market prices since there is no concern about liquidity issues or recovery rates. On the other hand, data on default is necessarily backward looking. With defaults being rare events, it is difficult to ensure that the data the model is calibrated to are representative of the PPF population going forward.

In principal, we could calibrate the model either to prices, or to default rates, and then move between risk-adjusted and physical probabilities by using an estimated price of risk. We prefer to follow the other alternative which is to fit the model to prices when we are computing risk adjusted probabilities and to default rates when computing physical probabilities, leaving the price of risk to be determined implicitly from the difference between the two.

However the calibration is done, whether to prices or to default rates, a choice needs to be made about the period and population that we are trying to match. Since credit spreads and default rates do vary widely over time and between populations, the choice has a material impact on our results. We discuss below the particular data sets we have chosen.

The classification of sponsors for the levy and the calibration of the model to the data rest heavily on the assumption that credit ratings are a reliable measure of the financial strength of the sponsor that mean the same thing over time and across populations. That is an assumption that has been severely questioned in the recent financial crisis, though it is fair to say that much of the criticism has been levelled at the rating of structured financial products rather than of ordinary corporate credits.

CDS calibration

We obtained CDS sector curves from Markit.com for the first trading day of each month between January 2001 and July 2010. Each curve contains average CDS rates, split by credit rating (AAA-CCC) and duration (6 months – 30 years). Figure 1 shows the time series of the 5-year CDS rate for each credit rating. There is substantial volatility, both in the level of each curve and the spread between different ratings, especially over the period since September 2007. In Figure 2, we focus on the CDS rates of investment-grade credits. Table 1 shows the average 5-year spread for each credit rating, for the entire sample and split by sub-period. It is unclear what the appropriate long-run average of these curves should be, and over what period it is appropriate to take an average.

Using an iterative approach (available from the authors) we extract risk-neutral survival probabilities from each curve, assuming that CDS premiums are paid annually in arrears, and all defaults occur at the end of the year. We assume that the yield curve was a flat 3% p.a. (results are insensitive to the particular yield curve assumed), and that the recovery rate on

default was 40% of the notional amount (the results are much more sensitive to this assumption; Markit.com provide an recovery rate for each sector curve; these average around 40% over the entire period). The resulting average survival probabilities are shown in Table 2 and Figure 3 for each credit rating for the entire period.

The data show some surprising features. The first is the lower survival of AAA-rated risks than AA-rated ones. Inspection of the underlying data reveals that this is entirely driven by the period 2007-2009, when spreads on the CDS's of AAA-rated bonds rose dramatically, possibly reflecting the liquidity crisis in the market for structured products. Markit was unable to provide a breakdown of the CDS market into structured products and other bonds.

The second surprising feature is the similar risk-neutral implied survival probabilities for CCC, BB and B-rated credits at very long time horizons. This may simply reflect a lack of available data for these types of contracts, although restricting the data sample to curves reported by Markit as being 'non-interpolated' does not change the results very much.

Finally, the results seem to suggest a large qualitative difference between the risk-adjusted survival probabilities of investment grade (BBB and above) and non-investment grade credits. It is not clear to what extent this difference reflects features of the CDS market, or genuine differences in survival probabilities. This suggests that care should be taken in applying these results directly to the population of risks in the PPF.

To calibrate our model, we chose common values for $\mu_s, \kappa, \sigma_s, \xi$ and λ , as well as the initial leverage ratios of each credit class, to ensure that the survival probabilities at one year and at 10 years matched as well as possible the CDS-implied risk neutral probabilities (RNP's). We choose 10 years because the CDS data for periods longer than 10 years appears to be unreliable, and because it is unreasonable to assume that our model parameters would be stable over periods longer than this. We exclude AAA-rated CDS rates, and focus on average survival probabilities over the whole sample period between one and ten years. The resulting fitted RNP's and CDS-implied RNP's are shown in Figure 4. Our parameter values are shown in Table 3.

Transition-matrix calibration

To estimate real-world survival probabilities, we choose our model parameters to ensure that the long-run real-world survival probabilities produced by our model match the long-

run real-world survival probabilities implied by a transition matrix approach. The transition matrix we chose was estimated by Moody's using global data from 1970-2008.⁴ We use a 5-year transition matrices to allow for the effect of ratings momentum, and matrices are adjusted for the presence of withdrawn ratings by making the assumption that ratings withdrawal is independent of ratings transitions.

Once again, we choose the common values for $\mu_s, \kappa, \sigma_s, \xi$ and λ , we also choose the initial leverage ratios of each credit class to ensure a good fit. The details of the particular parameter values chosen are presented in Table 3, and transition-matrix implied and calibrated survival probabilities are shown in Figure 5.

A comparison of the parameters for the two calibrations illustrates the broad features of the two transition calibrations. Financial strength under the risk-neutral basis reverts to a lower average value, and while the jump size is similar, the probability of a jump is much higher. The initial financial strength of companies of different ratings are also much lower than in the real-world basis. All of these imply much higher risk-neutral than real-world default probabilities, as we would expect.

Figure 6 shows the distribution of the credit ratings of a random sample of sponsors in the PPF on the 31st March 2009. Given the differences in the underlying population between the members of the PPF and companies with credit ratings, we do not expect that this closely represents the steady-state distribution of ratings implied by either of our transition matrices, although while there are few new entrants to the PPF, we would expect that the two distributions will converge over time.

Calibration of scheme funding

We make the assumption that in the long run, schemes would be required to target 110% of the s179 value of their liabilities, and that reversion to this mean would take around 8 years. We assume that the correlation between scheme deficits and financial strength was 0.6. Given the parameters of our calibrated process for sponsor financial strength, this implies an adjusted correlation between dz_s and dz_f of 0.9. Figure 7 shows the distribution of the s179 funding ratios of a random sample of 500 plans in the PPF as at 31st March 2009. It is likely that average plan funding will have improved significantly since then, owing to increased equity prices since that date.

⁴ From "Corporate Default Rates and Recovery Rates 1920-2008", Moody's Credit Research, February 2009, www.moodys.com.

4. An absolutely fair levy

Definition of the absolutely fair levy

Using our measure of the present value of the cost of PPF claims, it is possible to design a levy which is absolutely fair, in the sense that the (possibly risk-adjusted) expected present value of the levies to be paid by any uncapped scheme equals the expected present value of its future claims. (Those schemes which pay capped levies would have a present value of levy payments which was less than the present value of their future claims, and in this respect the levy design would not be absolutely fair.)

Before describing the absolutely fair levy in detail, it is important to understand that it is not, and cannot be an ideal levy structure. This is for two reasons: first, by making fairness the sole criterion, the ideal levy tends to perform badly or very badly on other criteria (predictability, stability, incentives to good behaviour). Second, the estimate of the fair levy is excessively sensitive to changes in model parameters that cannot be estimated with any precision. Despite these limitations, the absolutely fair levy is useful because it shows the direction of change that would make any proposed levy structure fairer.

We define the net cost of scheme i to the PPF as follows:

$$D_{i,0} = E[PV(C_{i,0} + C_{i,1} + \dots + C_{i,\omega}) - PV(L_{i,0} + L_{i,1} + \dots + L_{i,\omega})], \quad (4.1)$$

where $D_{i,t}$ is the expected cost of the claims of the scheme on the PPF net of expected future levy income from the scheme. $C_{i,t}$ denotes the claims paid by scheme i in year t , and $L_{i,t}$ denotes the levy paid by scheme i in year t . Since the future financial condition of the scheme and its sponsor are unknown at time 0, $C_{i,t}$ and $L_{i,t}$ are not known at the outset, so we have to compute their expected values.

We define the “absolutely fair” levy as one where the net cost of any uncapped scheme to the PPF is zero, so we set $D_{i,0} = 0$, and re-arrange (4.1), which yields

$$L_{i,0} = E[C_{i,0} + PV(D_{i,1})]. \quad (4.2)$$

The first component of the absolutely fair levy represents the present value of claims this year, while the second is the expected present value of the difference between future levies and future claims from next year on – which is only greater than zero should the scheme

become capped over the year. If the fair levy that a scheme should pay is greater than the cap, it only pays the cap, and hence carries a deficit to the PPF. Starting from the last period, ω , we are then able to calculate the entire fair levy schedule by working backwards. We use the absolutely fair levy as a basis of comparison for the proposed levy formula.

Results

Table 4 shows the annual absolutely fair levy, expressed as a percentage of scheme liabilities, assuming that the volatility of the scheme funding ratio is 15%. For both the risk-neutral and the real-world calibrations, the levy is heavily dependent on sponsor credit quality, and less so on scheme funding. For instance, for the real-world calibration, the annual fair levy, without a cap, varies by a factor of nearly 1000 between the worst credit score and the best for underfunded schemes. By comparison, for the strongest sponsors, a 30% reduction in scheme funding from fully funded to 70% funded raises the “absolutely fair” levy by a factor of only approximately 5. For the risk-neutral calibration, the variation is less extreme, but for underfunded schemes, the absolutely fair levy varies by a factor of 50 between the strongest credit quality and the weakest.

Introducing a cap on the levy raises the annual levy payments for uncapped schemes. As we have discussed, this is because schemes must prepay not only their expected claims each year but also the present value of the difference between the present value of future claims and future levies should they become capped. The further the scheme from the capped region, the lower the probability that the scheme will become capped, and the closer the fair levy is to the uncapped case. For very strong sponsors, introducing a cap has only a small effect on the levy because the chance that they will become capped is low. As the cap falls, the levies rise correspondingly, with the greatest rise occurring in the categories adjacent to the capped region, where the probability of becoming capped in the next period is the greatest. Any scheme in the capped region will not be able to pay sufficient levies to match the present value of its claims to the PPF. This is true whatever the chosen design of the levy.

These results highlight the drawbacks of a levy structure which is required to be absolutely fair. Levies for weaker sponsors with underfunded schemes far exceed their ability to pay; there are large changes in the levy as sponsors change their credit rating, particularly for weaker firms; and the penalties attached to scheme underfunding and investment risk are

small for strong sponsors, arguably where these are most important to prevent new generations of PPF claimants.

We next examine the effect of different investment strategies followed by pension plans on the “absolutely fair” levy. In Table 5, we show the effect of increasing the standard deviation of the funding ratio from 15% p.a. to 20% p.a., and of reducing it to 10% p.a. Results are shown for both real world and risk-adjusted probabilities, for the risk-adjusted calibration. Investment risk is relatively unimportant for well funded schemes; the reason is that they are not likely to default for many years, so their funding level at default is much more dependent on the target funding level than on the volatility of the deficit along the way. However, for low rated schemes investment risk is much more important, and in particular for well-funded schemes, especially those schemes near the levy cap.

Partly to illustrate our methodology, but also as a check, we calculate the degree of fairness inherent in the “absolutely fair” levy. To test levy fairness, we consider the actual population of schemes in the PPF, using a random sample of 500 schemes. For each scheme, we had the value of assets and liabilities on 31st March 2009, a measure of the financial strength of the scheme sponsor as at that date, and a volatility of its funding ratio, based on its current investment strategy, provided to us by Redington. We test levy fairness using this sample in two ways. Firstly, we calculate the expected discounted risk-adjusted present value of claims and levies. Then, for each scheme, we calculate the ratio of the present value of levies to claims, and calculate the cumulative distribution function (CDF) of the liabilities of the 500 schemes w.r.t. this ratio. We graph these CDF’s for capped and uncapped liabilities in Figure 8. The proportion of the liabilities of the sample of schemes which have a ratio of the expected discounted present value of levies to claims that is less than a particular value can be read off the graph easily. We calculate this distribution only for the risk-adjusted values, and at a horizon of 10 years.

A second measure of inequality is the Gini coefficient, which is best understood graphically. If one were to sort schemes by the $PV(\text{Levies})/PV(\text{Claims})$ ratio, and then calculate and plot the cumulative proportion of total levies against the cumulative proportion of total claims, one would obtain a convex graph which connects the points (0,0) and (1,1). The area between this curve and the 45-degree line, expressed as a proportion of the total area under the 45-degree line, is a measure of the inequality of the distribution of levies w.r.t. claims and is called the Gini coefficient. For the perfectly fair levy without a cap, the Gini coefficient is 0, as we would expect, while for the perfectly fair levy with a cap of 1% p.a.,

the Gini coefficient is 17.75%. This represents the fairest attainable levy allocation available to the PPF, if it cannot expect to collect levies larger than 1% p.a. from any scheme.

5. The proposed levy formula

The present value of PPF claims

Before examining the proposed levy formula, we first use our model to examine the properties of the present value of PPF claims for schemes with different initial funding levels and different sponsor credit strengths. Results are shown in Table 6 for both risk-adjusted and real-world calibrations at a time horizon of ten years.

The financial strength of the sponsor is the main factor affecting the present value of claims. As the sponsor becomes weaker, the present value of claims rises for underfunded schemes. For overfunded schemes, the present value of claims rises and then, surprisingly, falls. This is because weaker sponsors have a much shorter expected life than stronger sponsors, and if the plan is significantly overfunded, it is quite likely that a very weak sponsor will have defaulted before the plan can become underfunded and claim on the PPF, while the financially stronger scheme is more likely to live on and default later when a deficit may well have developed.

Although in this paper we have focused on a 10-year time horizon, our model predicts a slow fall in the credit quality of pension fund sponsors. This is likely to be especially significant for stronger plan sponsors. Since most of the liabilities covered by the PPF are currently associated with higher credit quality sponsors, many PPF claims can be expected to occur after this 10-year time horizon. Although our model is by its nature imprecise, and only provides indicative values at long-run horizons, it is important to at least examine the predicted time pattern of claims on the PPF. Table 7 shows the risk-adjusted present value of claims which occur in each decade of the plan's life for three schemes of different credit quality and initial plan funding of 70% of liabilities. The risk-adjusted and real-world claims decline much more rapidly for weaker sponsors. Another consequence of this is that the initial funding ratio of the plan has a relatively small effect on the present value of claims for stronger sponsors, but is highly significant for weaker ones.

Table 6 also illustrates that the present value of claims is higher under risk-adjusted probabilities than under real world probabilities for almost all sponsor strengths and funding ratios, with the proportional difference being highest for very overfunded schemes with strong sponsors. The difference is also very large for underfunded schemes with strong sponsors, reflecting the high risk premium the market attaches to the defaults of strong sponsors. This difference appears to persist over time, as can be seen by comparing the claim profiles shown in Table 7, although we emphasise that given the extreme uncertainties in very long-run modelling, these results can only be taken as indicative.

The proposed levy

We now turn to the proposed levy, which is given by the following formula:

$$P_{it} = \min(\max(S\pi_{it}[(1+mL)L_{it} - (1-mA)A_{it}], 0), cL_{it}), \quad (5.1)$$

where P_{it} is the levy paid by pension fund i , S is a scaling factor to ensure that the aggregate levy collected matches the total PPF quantum, π_{it} is the default probability-based levy rate for pension fund i , which depends on the credit quality of the pension fund's sponsor, mL is the stress test applied to the pension fund's liabilities, L_{it} , and mA is the stress test applied to the pension fund's assets, A_{it} , all at time t . The annual levy is capped at a fraction c of scheme liabilities. The important parameters of this formula are π_{it} , mL , mA and c .

We first try to choose the parameters of (5.1) to match the "absolutely fair" levy rates we illustrate in Table 4. We approach the problem by ignoring the effects of asset mix and excluding very weak sponsors (so fitting only to the first three or four rows of each panel of Table 3). We focus only on risk-neutral probabilities, and set $S = 1$.

The results of this exercise are shown in Table 8 for different values of the levy cap. Overall, the parameters obtained are satisfactory: they are broadly consistent with the range of levy structures that have been under discussion, especially for the uncapped levy. In both cases, there is a large adjustment to the liabilities to reflect the fact that even relatively well-funded schemes should pay levies. In the case of the uncapped levy, the default probabilities look quite reasonable when compared with CDS default rates. If the levy is capped at 1% p.a., the probabilities rise for investment grade credits, especially for those nearest the cap. This illustrates the undesirable properties of the capped "absolutely fair" levy. Also, the default probabilities for Ba- and B-rated risks cannot be estimated because the cap binds at all funding levels.

Alternative levy designs

We now turn to measuring the extent to which the proposed levy formula is fair using alternative sets of parameters. The possibilities which we examine are presented in Table 9. The first column, A, contains a set of probabilities obtained from the PPF and designed to approximate the current levy basis with fewer rating categories. The second, B, contains the Basel-II probabilities used in the Steering Group report, and the third, C, contains risk-adjusted annual probabilities of default for each credit class, calculated as the average of the one-year risk-adjusted default probabilities over the entire period covered by our CDS rates. 'D' and 'E' contain the "best fit" parameters for the uncapped and capped levy respectively, taken from Table 8.

Using these parameters as inputs to our model, we calculated the ratio of the risk-adjusted present value of levies to claims for each scheme. Following the approach introduced in section 4, we then calculated a Gini coefficient measuring the inequality of the distribution of levies relative to claims over the ten year period. This provides a measure of the inequality in the distribution of levies relative expected claims which takes account of the different size of schemes covered by the PPF. To get some idea of the distribution of levies across claims for an average scheme, we also calculated the standard deviation of the ratio of expected levies to claims, standardised by the mean ratio by scheme. This measure, the coefficient of variation, would be zero for a completely fair scheme.

Since the effect of the cap depends on the size of the levy collected (the higher the total amount of levy collected, the more important the cap becomes, and the less equal any levy structure can be), we also standardised each of the levy designs by altering the scaling factor so that the aggregate present value of levies collected equalled the present value of levies collected under levy schedule A – the best approximation to the current levy – with a multiplier of 3.4.

Table 10 shows our results. The first column contains the calculated scaling factors and the second and third columns contain the different values of mA and mL we examined. The next panel contains the mean ratio of the present value of levies to claims across all schemes in the sample, the standard deviation of the levy-claims ratio and its coefficient of variation across schemes. The final column shows the calculated Gini coefficients.

There is not a striking difference in the degree of unfairness inherent in the different levy designs. (While the "absolutely fair" levy performs best out of all the levy designs

examined, the results are not comparable with the other levies because the absolutely fair levy collects levies with a present value of 82% of the aggregate present value of claims; the other levies collect less, around 57% of the aggregate present value of claims.) The levy designs listed in Table 9 have coefficients of variation which vary between 45% for the “best fit” levy (D), with an allowance for investment risk, and 56% for the “best fit” to the capped levy (E). The Gini coefficients vary between 23.7% for the CDS-based probabilities (C) without any offsets, and 18.6% for the current levy probabilities (A).

These results seem to suggest that the particular choice of levy parameters has only small effects on fairness, a somewhat puzzling conclusion given the large apparent differences between the different designs in Table 9.

To investigate this, we calculate the “effective” probabilities in each levy design by multiplying each set of levy probabilities by the scaling factor we calculated in Table 10, in each case choosing the scaling factor with the smallest values of m_A and m_L and making no further allowance for the different values of m_A and m_L . The results are plotted in Figure 10, using a logarithmic scale on the vertical axis.

Figure 10 shows that the different levy designs are much more similar than first appears, once they have been adjusted by appropriate scaling factors. In particular, ‘A’ and ‘E’ are very similar for credit ratings below B, and ‘B’ and ‘C’ are quite similar to each other across all credit ratings. Levy ‘D’ is a compromise between levies ‘B’ and ‘C’, which are somewhat flatter, collecting more from investment-grade schemes, and levies ‘A’ and ‘E’, which are somewhat steeper, collecting more from weaker schemes. To some extent, the broad features of Figure 10 are also reflected in the results in Table 10: ‘A’ and ‘E’ look quite similar in both the scheme-based and value-based unfairness measures, while the same is true of ‘B’ and ‘C’ with zero values of m_A and m_L .

One important conclusion from Table 10 is that the levy designs that make allowance for investment risk are substantially fairer than those that do not, whether measured by the coefficient of variation or the Gini coefficient. This is because full funding does not fully shield the PPF from a future claim; if the sponsor’s financial strength deteriorates sharply, and then finally fails altogether after several years, the funding level at the time of failure may be much lower than it was initially, leading to a substantial claim on the PPF that will be only offset to a limited extent by levy payments in the intervening period.

In broad terms, we conclude that within the rather narrow range of levy structures that we have considered, differences in levy rates do not have a major impact on fairness, at least when measured over the sample of 500 schemes. However, levy designs which take account of investment risk are fairer than those which do not..

6. Conclusion

In this paper, we have established a definition of fairness of the risk-based levy. The definition is that the present value of the levy that a scheme can expect to pay looking forward should be proportional to the claims it expects to make on the PPF. We have discussed in some detail the difficulties associated with applying such a definition in practice, including the timescale over which it is applied (while recognising the limitations of long-term modelling, we favour a longer period over a shorter one) and what account, if any, it should take of the inability of the PPF to collect levies from weaker schemes (we have argued that any definition of fairness should make allowance for the inability of weaker schemes to pay levies). We have also argued that when computing present values for setting levies, it is appropriate to allow for a risk premium which reflects the economic circumstances in which defaults are likely to occur.

We have shown how to apply the definition of fairness through the construction of a model which values future claims as well as future levy payments. The model we have built incorporates the broad features of a sponsor and its pension fund, including the dynamic nature of scheme solvency, pension scheme funding which varies over time in response to asset returns and some relationship between the two. We calibrated our model to both real-world default probabilities, obtained from past default rates, and market-implied default probabilities, which incorporate a risk premium.

While the detailed results of our model are clearly sensitive to the particular calibration we have chosen and the assumptions we have made, the model provides a useful measure of the fairness of alternative levy designs.

We have used our model to show what a completely fair levy would look like. It produces a levy structure which is a useful benchmark, but has features that make it impracticable: there are steep discontinuities in levy rates for sponsors with different financial strength, and the levies charged are relatively insensitive to the funding status of pension schemes.

The perfectly fair levy demonstrates that there is a real conflict between fairness, stability and incentives when considering levy design. A perfectly fair levy is not stable or predictable for individual schemes and it does not provide schemes with strong incentives to encourage good behaviour.

This conflict aside, however, there is another reason why it may be sensible to deviate from our absolutely fair levy in practice: in deriving our levy, we have made strong assumptions which may not actually hold. Most importantly, we have assumed that the financial strength of sponsors can be observed with perfect accuracy, and that the parameters underlying the dynamics of sponsor solvency and scheme funding are known and do not change. The fact that these do not hold in practice suggests that while the absolutely fair levy may be useful as a guide to the appropriate level of the levy, it should not be applied in practice without some tempering.

We have also used our model to appraise the fairness of the proposed levy formula with different parameters. We have found that if the levy parameters are set to collect the same amount of levy in aggregate, there seem to be only small differences in fairness between the various parameterisations we have examined. Part of the reason for this is that the different levy designs actually look quite similar, once the levy cap has been allowed for and they have been appropriately scaled. However, levies which take account of the differences in the investment risk of different schemes allocate the levy more fairly than those which do not.

We have established a procedure for comparing different levy structures on the basis of fairness. We recognise that there are many other factors that need to be taken into account in determining the levy structure going forward, but the methodology we have developed provides a means for measuring the fairness of different proposed solutions, and for considering the trade-off between fairness and other criteria.

7. Tables

Table 1: Average 5-year CDS spread, b.p.'s, Markit.com sector indices, all industries, whole period and sub-periods

Rating	Period			All
	Jan 2001- Mar 2004	Apr 2004 - May 2007	Jun 2007 - Jul 2010	
AAA	31	16	151	63
AA	33	16	94	46
A	79	28	111	70
BBB	133	51	149	108
BB	335	173	365	286
B	341	356	781	479
CCC	706	667	1990	1081

Table 2: CDS-implied risk-neutral survival probabilities, based on Markit.com sector indices, all industries.

Term	Rating						
	AAA	AA	A	BBB	BB	B	CCC
0	1.000	1.000	1.000	1.000	1.000	1.000	1.000
1	0.979	0.993	0.988	0.983	0.957	0.932	0.839
2	0.969	0.986	0.977	0.966	0.912	0.859	0.712
3	0.962	0.978	0.965	0.948	0.866	0.787	0.615
4	0.956	0.969	0.952	0.928	0.819	0.719	0.537
5	0.947	0.958	0.937	0.906	0.771	0.655	0.473
6	0.935	0.947	0.921	0.882	0.723	0.598	0.419
7	0.922	0.935	0.902	0.856	0.674	0.543	0.375
8	0.905	0.923	0.882	0.828	0.626	0.496	0.337
9	0.887	0.910	0.861	0.798	0.578	0.452	0.305
10	0.868	0.896	0.838	0.767	0.531	0.412	0.280
11	0.848	0.882	0.816	0.736	0.490	0.378	0.255
12	0.827	0.868	0.792	0.705	0.449	0.347	0.235
13	0.806	0.853	0.767	0.673	0.411	0.319	0.219
14	0.785	0.838	0.742	0.641	0.374	0.294	0.206
15	0.764	0.823	0.716	0.608	0.339	0.271	0.196
16	0.746	0.809	0.693	0.580	0.311	0.252	0.183
17	0.728	0.795	0.670	0.552	0.285	0.235	0.173
18	0.711	0.780	0.646	0.524	0.260	0.220	0.164
19	0.693	0.766	0.622	0.496	0.237	0.206	0.157
20	0.676	0.751	0.598	0.468	0.215	0.194	0.152
21	0.665	0.739	0.580	0.448	0.200	0.184	0.143
22	0.653	0.727	0.561	0.427	0.186	0.174	0.136
23	0.642	0.715	0.543	0.407	0.172	0.165	0.129
24	0.631	0.703	0.525	0.388	0.159	0.157	0.124
25	0.620	0.691	0.507	0.369	0.146	0.150	0.119
26	0.609	0.679	0.489	0.350	0.134	0.143	0.115
27	0.598	0.668	0.471	0.331	0.123	0.137	0.112
28	0.588	0.656	0.454	0.313	0.112	0.132	0.110
29	0.578	0.645	0.437	0.295	0.102	0.127	0.109
30	0.568	0.634	0.420	0.277	0.092	0.123	0.108

Table 3: Assumptions used in model of firm default.

	<i>CDS-based calibration</i>	<i>Moody's 1970-2008 calibration</i>
kappa	0.03	0.03
sigma-v	0.306495	0.230887
J	1.125599	1.517834
lambda	0.04634	0.014462
s-bar (real world)	-	0.609731
s-bar (risk neutral)	0.26503	-
rho*	0.9	0.9
<i>Initial financial strength</i>		
	<i>CDS calibration</i>	<i>Moody's 1970-2008 calibration</i>
Aaa and Aa	2.477503	4.341603
A	1.864526	3.091919
Baa	1.427386	1.843046
Ba	0.709605	0.720685
B	0.608703	0.274057
C	0.431586	0.106367

Table 4: Annual absolutely fair levy, real-world and risk-neutral calibration, various caps. Capped levies shown in bold. Volatility of funding ratio equals 0.15.

Real-world calibration								Risk-neutral calibration							
No cap								No cap							
Rating	Funding ratio							Rating	Funding ratio						
	0.7	0.8	0.9	1	1.1	1.2	1.3		0.7	0.8	0.9	1	1.1	1.2	1.3
Aaa and Aa	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	Aaa and Aa	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%	0.0%
A	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	A	0.3%	0.2%	0.1%	0.0%	0.0%	0.0%	0.0%
Baa	0.1%	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	Baa	0.4%	0.3%	0.2%	0.1%	0.0%	0.0%	0.0%
Ba	0.3%	0.2%	0.1%	0.0%	0.0%	0.0%	0.0%	Ba	1.2%	0.9%	0.6%	0.4%	0.2%	0.1%	0.1%
B	6.7%	5.3%	3.9%	2.4%	1.1%	0.2%	0.0%	B	1.8%	1.4%	1.0%	0.7%	0.4%	0.2%	0.1%
Caa and below	19.1%	13.8%	8.5%	3.3%	0.3%	0.0%	0.0%	Caa and below	5.2%	4.2%	3.1%	2.0%	1.1%	0.4%	0.0%
Cap = 1% p.a.								Cap = 1% p.a.							
Rating	Funding ratio							Rating	Funding ratio						
	0.7	0.8	0.9	1	1.1	1.2	1.3		0.7	0.8	0.9	1	1.1	1.2	1.3
Aaa and Aa	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	Aaa and Aa	0.3%	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%
A	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	A	0.7%	0.5%	0.3%	0.2%	0.1%	0.1%	0.0%
Baa	0.1%	0.1%	0.1%	0.0%	0.0%	0.0%	0.0%	Baa	1.0%	1.0%	1.0%	0.8%	0.5%	0.3%	0.2%
Ba	1.0%	1.0%	1.0%	1.0%	1.0%	0.7%	0.4%	Ba	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%
B	1.0%	1.0%	1.0%	1.0%	1.0%	0.4%	0.1%	B	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%	0.8%
Caa and below	1.0%	1.0%	1.0%	1.0%	0.4%	0.0%	0.0%	Caa and below	1.0%	1.0%	1.0%	1.0%	1.0%	0.9%	0.4%
Cap = 0.5% p.a.								Cap = 0.5% p.a.							
Rating	Funding ratio							Rating	Funding ratio						
	0.7	0.8	0.9	1	1.1	1.2	1.3		0.7	0.8	0.9	1	1.1	1.2	1.3
Aaa and Aa	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	Aaa and Aa	0.5%	0.3%	0.2%	0.1%	0.1%	0.1%	0.0%
A	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	A	0.5%	0.5%	0.5%	0.5%	0.4%	0.3%	0.2%
Baa	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%	0.0%	Baa	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%
Ba	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	Ba	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%
B	0.5%	0.5%	0.5%	0.5%	0.5%	0.4%	0.1%	B	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%

Caa and below **0.5%** **0.5%** **0.5%** **0.5%** 0.3% 0.1% 0.0% Caa and below **0.5%** **0.5%** **0.5%** **0.5%** **0.5%** **0.5%** **0.5%**

Table 5: Annual absolutely fair levy, risk-neutral calibration, various caps, different volatilities of funding ratios.. Capped levies shown in bold.

Volatility of funding ratio = 0.10 p.a.								Volatility of funding ratio = 0.20 p.a.							
No cap	Funding ratio							No cap	Funding ratio						
Rating	0.7	0.8	0.9	1	1.1	1.2	1.3	Rating	0.7	0.8	0.9	1	1.1	1.2	1.3
Aaa and Aa	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%	0.0%	Aaa and Aa	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%	0.0%
A	0.3%	0.2%	0.1%	0.0%	0.0%	0.0%	0.0%	A	0.3%	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%
Baa	0.4%	0.3%	0.1%	0.0%	0.0%	0.0%	0.0%	Baa	0.4%	0.3%	0.2%	0.1%	0.0%	0.0%	0.0%
Ba	1.2%	0.8%	0.5%	0.3%	0.1%	0.0%	0.0%	Ba	1.3%	1.0%	0.7%	0.5%	0.4%	0.3%	0.2%
B	1.7%	1.3%	0.8%	0.5%	0.2%	0.0%	0.0%	B	2.0%	1.6%	1.2%	0.9%	0.6%	0.4%	0.3%
Caa and below	4.7%	3.6%	2.5%	1.4%	0.4%	0.0%	0.0%	Caa and below	5.7%	4.7%	3.6%	2.7%	1.8%	1.0%	0.4%
Cap = 1% p.a.								Cap = 1% p.a.							
	Funding ratio								Funding ratio						
Rating	0.7	0.8	0.9	1	1.1	1.2	1.3	Rating	0.7	0.8	0.9	1	1.1	1.2	1.3
Aaa and Aa	0.2%	0.2%	0.1%	0.0%	0.0%	0.0%	0.0%	Aaa and Aa	0.3%	0.2%	0.2%	0.1%	0.1%	0.0%	0.0%
A	0.5%	0.3%	0.2%	0.1%	0.0%	0.0%	0.0%	A	1.0%	0.8%	0.6%	0.4%	0.3%	0.2%	0.2%
Baa	1.0%	0.9%	0.5%	0.3%	0.1%	0.0%	0.0%	Baa	1.0%	1.0%	1.0%	1.0%	1.0%	0.9%	0.7%
Ba	1.0%	1.0%	1.0%	1.0%	0.9%	0.3%	0.0%	Ba	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%
B	1.0%	1.0%	1.0%	1.0%	1.0%	0.2%	0.0%	B	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%
Caa and below	1.0%	1.0%	1.0%	1.0%	0.9%	0.1%	0.0%	Caa and below	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%	1.0%
Cap = 0.5% p.a.								Cap = 0.5% p.a.							
	Funding ratio								Funding ratio						
Rating	0.7	0.8	0.9	1	1.1	1.2	1.3	Rating	0.7	0.8	0.9	1	1.1	1.2	1.3
Aaa and Aa	0.3%	0.2%	0.1%	0.1%	0.0%	0.0%	0.0%	Aaa and Aa	0.5%	0.5%	0.4%	0.3%	0.2%	0.2%	0.1%
A	0.5%	0.5%	0.4%	0.2%	0.1%	0.0%	0.0%	A	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%
Baa	0.5%	0.5%	0.5%	0.5%	0.4%	0.2%	0.1%	Baa	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%
Ba	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.1%	Ba	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%
B	0.5%	0.5%	0.5%	0.5%	0.5%	0.4%	0.1%	B	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%	0.5%

Caa and below **0.5% 0.5% 0.5% 0.5% 0.5%** 0.2% 0.0% Caa and below **0.5% 0.5% 0.5% 0.5% 0.5% 0.5%**

Table 6: PV of future PPF claims as a proportion of scheme liabilities, real-world and risk-adjusted calibration, different initial sponsor financial strengths and plan funding ratios, time horizon 10 years.

		<i>PV (Claims), real-world probabilities</i>						
		<i>Initial scheme funding ratio</i>						
		0.7	0.8	0.9	1	1.1	1.2	1.3
Aa		0.2%	0.2%	0.1%	0.1%	0.1%	0.1%	0.0%
A		0.4%	0.3%	0.2%	0.2%	0.1%	0.1%	0.1%
Baa		1.0%	0.8%	0.7%	0.6%	0.5%	0.4%	0.4%
Ba		7.0%	6.2%	5.4%	4.7%	4.0%	3.4%	2.8%
B		17.3%	13.9%	10.5%	7.1%	4.1%	1.8%	0.9%
C		24.3%	17.8%	11.2%	4.9%	1.1%	0.4%	0.2%

		<i>PV (Claims), risk-adjusted probabilities</i>						
		<i>Initial scheme funding ratio</i>						
		0.7	0.8	0.9	1	1.1	1.2	1.3
Aa		2.4%	2.1%	1.7%	1.5%	1.3%	1.1%	0.9%
A		4.0%	3.5%	3.0%	2.5%	2.2%	1.9%	1.7%
Baa		6.1%	5.3%	4.6%	3.9%	3.4%	3.0%	2.7%
Ba		13.5%	11.6%	9.8%	8.1%	6.6%	5.3%	4.1%
B		15.2%	12.9%	10.7%	8.7%	6.9%	5.2%	3.7%
C		18.8%	15.5%	12.3%	9.3%	6.5%	4.0%	2.5%

Table 7: Time profile of risk-adjusted and real-world present value of claims as a proportion of initial scheme liabilities for a plan which is initially 70% funded, for different sponsor credit qualities.

Plan decade	Aa-rated		Baa-rated		B-rated	
	Risk-adjusted	Real-world	Risk-adjusted	Real-world	Risk-adjusted	Real-world
1	2.42%	0.19%	4.02%	1.02%	15.29%	17.32%
2	2.20%	0.17%	2.66%	1.06%	2.01%	0.87%
3	1.09%	0.17%	1.07%	0.58%	0.61%	0.26%
4	0.40%	0.13%	0.37%	0.23%	0.19%	0.09%
5	0.11%	0.05%	0.09%	0.07%	0.05%	0.02%
6	0.02%	0.01%	0.02%	0.01%	0.01%	0.00%
7	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
8	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
9	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
10	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
TOTAL	6.24%	0.73%	8.23%	2.97%	18.15%	18.57%

Table 8: Parameters to obtain best fit to “absolutely fair” levy

	<i>Levy uncapped</i>	<i>Levy cap = 1.0% p.a.</i>
<i>Default probabilities (b.p.'s)</i>		
Aaa and Aa	26	32
A	48	86
Baa	75	248
Ba	239	-
B	357	-
C	1043	857
<i>Other parameters</i>		
<i>mL</i>	0.1991	0.3095
<i>mA</i>	-	-

Table 9: Alternative levy rates tested

Rating category	A Current levy schedule	B Basel-based schedule	C One-year risk-neutral default probabilities (CDS calibration)	D Uncapped levy "best fit" parameters	E Capped levy "best fit" parameters
Aaa and Aa	6	20	67	26	32
A	18	50	117	48	86
Baa	50	110	173	75	248
Ba	135	160	432	239	857
B	450	400	677	357	857
C	2150	400	1614	1043	857

Table 10: Measures of fairness of PPF liabilities using various levy parameters, risk-adjusted probabilities

Probability schedule	Levy basis			Levy cap 1% p.a.			Gini coefficient (by value)
	S	mA	mL	PV(L)/PV(CI) (by scheme)			
				Mean	Std dev	Co-eff. of var	
Current (A)	3.400	0	0.2	57%	31%	54.2%	18.6%
Basel (B)	2.952	0	0	54%	27%	49.7%	22.0%
Basel (B)	1.499	sig	0.05	54%	24%	45.4%	20.3%
CDS (C)	1.298	0	0	54%	28%	52.0%	23.7%
CDS (C)	0.680	sig	0.05	54%	25%	46.9%	19.9%
Best fit (D)	1.528	0	0.1991	56%	30%	53.7%	18.7%
Best fit (D)	1.673	0	sig	53%	24%	45.0%	19.2%
Best fit (capped) (E)	0.543	0	0.3095	57%	32%	56.1%	18.7%
Absolutely fair (capped)	-	-	-	75%	30%	40.5%	17.8%

Note: sig refers to the scheme-specific volatility of the funding ratio based on each scheme's investment mix and liability profile as obtained from Redington. For each levy design, the scaling factor has been adjusted to ensure that the aggregate risk-adjusted present value of levies collected over the 10 years for all schemes in the sample equals the aggregate risk-adjusted present value of the current levy (A), with a scaling factor of 3.4, which is 57% of risk-adjusted expected present value of claims over this period. The "absolutely fair" levy is not comparable with the others as it collects 82% of the present value of claims.

8. Figures

Figure 1: CDS spreads (b.p.'s), Markit.com sector indices, broad rating category, all industries, January 2001-July 2010.




Figure 2: CDS spreads (b.p.'s), Markit.com sector indices, broad rating category, investment grades only, all industries, January 2001-July 2010.




Figure 3: CDS-implied risk-neutral survival probabilities, based on Markit.com sector indices, all industries, January 2001-July 2010.





Figure 4: 10-year CDS-implied and calibrated model-derived risk-neutral survival probabilities



(Dotted lines indicate CDS-implied survival probabilities)

Figure 5: Transition-matrix-implied and calibrated model-derived risk-neutral survival probabilities




Figure 6: Distribution of schemes by sponsor credit rating




Figure 7: Distribution of schemes by s179 scheme funding, 31st March 2009.




Figure 8: Distribution of PPF liabilities by risk-adjusted PV(Levies)/PV(Claims) ratio, 10-year horizon, CDS-implied model-derived survival probabilities




Figure 9: Time profile of the risk-adjusted and real-world present value of claims as a proportion of scheme liabilities by scheme decade, Aa-rated sponsor, initial funding ratio = 70%, volatility of funding ratio = 0.15 p.a.




Figure 10: Levy probability schedule for each levy design in Table 9, multiplied by the scaling factors calculated in Table 10. Logarithmic scale.

