

ABSTRACT OF THE DISCUSSION

Dr J. M. Orszag (Institute Affiliate; introducing the paper): The financial economics discussions in the actuarial profession have largely been in the context of the classical economic model in which markets are perfect. This model makes a number of simplifying assumptions: that all contingent claims are tradable; and that property rights are well defined. However, these are assumptions, and they tend to break down over the long-term setting faced by pension schemes, and also by life insurance funds.

There are a number of ways which we can consider what to do when these assumptions, that all contingent claims are tradable and property rights are well defined, break down. We take the approach which is closest to the one which has been adopted in the mainstream economic literature, of explicitly modelling background risk in an expected utility framework. Background risks are unhedgeable risks which cannot be diversified away. They are not always financial, and they include regulatory and policy risks. For instance, changes in the legal requirements on pension liabilities associated with the Pension Protection Fund (PPF) or associated with changes in pension laws are things which cannot be hedged, cannot be diversified away easily, and would count as background risk, as would changes in the tax regime. Again, this is not easily hedgeable by pension funds. Changes in the legal definition of liabilities would also be a background risk.

There are labour market risks which are background risks, such as aggregate wage movements. Some people argue that part of long-term aggregate wage movements can be hedged, but the part which cannot be hedged is background risk. There are also other labour market risks, such as early retirement rates and changes in turnover which cannot be hedged by individual funds. Mortality is another risk which is unhedgeable to a large degree, and cannot be diversified away, although there are some instruments in the market which are trying to change this. Potentially, there are some long-term financial risks, particularly outside the United Kingdom, which cannot be hedged by long-term financial instruments.

So, there is a wide range of background risks which are quite important. You could argue that some of these could be hedged in the future through new financial instruments; and they can, but it is costly to introduce such instruments. There is literature about the introduction of financial instruments with a number of conclusions, but we start in this paper from the assumption that these markets do not exist, and we work through some of the conclusions, because these are risks which are endemic in pension schemes.

Section 2 has a review of the economic literature on background risk. This is a broad literature dating back to the 1960s. There has been a considerable effort made over the past 30 years in developing literature on what happens when markets are incomplete and imperfect.

Then the types of background risk affecting pension schemes are reviewed — for example mortality risk and wage rate risk. The point here is that people are focused on background risk being mortality, yet, in practice, there are many other important background risks to take into account, including regulatory and policy risks.

There follows a conceptual discussion concerning the applicability of the economic literature to pensions. We have to be quite careful here, because, although some of the constructs, like preferences in utility functions, are relevant in pensions, there are some subtle differences. In economics we talk about utility as something which depends on consumption and wealth; in pensions we are talking about the contributions made on funding and the overall funding position (that is the net asset position of the scheme), and the constraints are slightly different. Instead of it being a lifetime wealth constraint which faces an individual, or a resource constraint affecting a firm, it is an asset/liability constraint over the long run which affects a pension scheme.

For very small pension schemes we have to be careful about taking utility modelling considerations too far. However, from the point view of sponsors, utility and the formal modelling of objective functions are quite important.

The way in which pensions are structured is complicated — much more complicated than, say, the structure of a problem which faces an individual. For instance, we talk about risk aversion in the paper. For individuals, it is reasonable to assume that relative risk aversion, or the amount which people are willing to pay to avoid a percentage shock to their wealth, does not vary across their wealth levels, whereas absolute risk aversion, the amount which they are willing to pay for an absolute gamble, declines across wealth levels in a monotonic fashion.

For pension funds it could be more complicated. If we consider a pension fund which is not well funded, especially in the current environment with the PPF, the sponsor may be risk loving. It is not going to be in the situation where it will be highly risk averse and willing to put a lot of money into bonds. As the funding level improves, it will become more risk averse, and then at higher funding levels it could be less risk averse, so the pattern is not monotone, and it could have a very complicated shape and be scheme specific. So, we have been quite careful that objective functions which come from the economic literature may not be applicable to pensions.

We address three questions in the paper: “What does background risk do to asset allocation?” “What does background risk do to funding decisions?” and “What does background risk do to fair value calculations?”

I think that there are two dangerous myths in the current debate about background risk and unhedgeable risk and how they affect pensions. The first myth is: “Background risk should be ignored. Just invest in bonds and everything will be okay. This background risk is just noise on the side, and people should just invest in bonds.” I think that that is one thing which we address with this paper.

The second myth which needs to be avoided is: “Background risk should lead to more equity investment, because it does not matter that much if you put a little bit more money in equity, and there is actually no need to put a lot more money into pension plans because things could get better in the future.” This is a myth, because there are a number of considerations: background risk leads to a high level of risk, which may mean that you are more conservative; on the other hand, the argument is correct that the marginal cost of investing in equity is lower in the case where you have this background risk; and the two considerations can go in either direction.

Returning to the economic literature on background risk, saving because of background risk is called precautionary saving. In the paper we call extra funding due to background risk precautionary funding. This is driven by prudence. Actuaries use the word ‘prudence’, but economists use it in a different way, which is to tie it to the convexity of the marginal value of consumption in the background risk variable. This, in most economic applications — and we argue in the paper that this is not necessarily always applicable in the more complicated pension applications — relates to the third derivative of the utility function. Risk aversion is a second derivative property. Curvature of preferences is what determines asset allocation. The way in which the curvature changes, the way in which the risk aversion changes, is what determines prudence. If there is precautionary saving and if prudence is positive, it leads to a positive margin in the pricing of assets. This leads to a degree of prudence in assumptions, which is something which actuaries do intuitively, although I think that this could be made a bit more explicit.

How about asset allocation? Is it the case that, if we have background risk, we are going to put less money into equities or keep things the same? Admittedly, in some of the models which have been developed in the actuarial literature (for instance mean-variance models) background risk does not matter.

In Section 4.4 we consider a very simple model, altered a little from the text book model to give it more of a pensions context. For example, we have replaced the asset level with the funding level. It is reasonable, you would think, to assume that at low levels of funding the marginal value of extra funding is higher than at high levels of funding. So, with this particular objective function, you get twice as much benefit from an extra unit of funding when you underfund it than you do when you overfund it. So, again, in the objective function, if you are below 100% funding you get twice as much benefit than if you are above 100% funding — a very simple objective function with a kink at 100%, as shown in Figure D.1.

Now let us start with a model, Case 1, with no background risk, with four states of the world, as shown in Figure D.2. We begin with 101% funding, 1% over full funding, and we have

Slight Variation on a Textbook Example

- Let x be the level of funding (e.g. 100 = full funding)
- Assume after 100% funding, less important.




Figure D.1. Example from ¶4.4.1

Case 1: No Background Risk

- Risky asset: returns +190% or -100%
- 4 State Model
 - State 1: 101% solvency minus y
 - State 2: 101% solvency plus $y \cdot 190\%$.
 - State 3: 101% solvency minus y
 - State 4: 101% solvency plus $y \cdot 190\%$
- 1 Unit Gets Purchased

Figure D.2. Case 1, no background risk

a choice of investing in a risky asset with two very extreme returns. The returns are plus 190% or minus 100%. Minus 100% means that you lose all your money; plus 190% means that you get a big return. Over the long term, it is not unreasonable to have a high rate of return on a high level of risk. In States 1 and 3 you lose all your money invested in the risky asset. So, what you will get at the end of the period in States 1 and 3 is 101% of what you originally had minus anything invested in the risky asset.

In States 2 and 4 the risky asset does very well, and you get 101% plus 190% of whatever you invested. What is the investment in the risky asset? One unit of the risky asset gets purchased. This is exactly the conclusion that Bodie *et al.* (1992) would reach. You invest in bonds until you are fully funded, and above that you put money into equities. So, a very simple model leads to a very simple conclusion: you put everything in bonds and then, above that, if you have some funding cushion, you invest some in equities.

We now introduce background risk; a very simple type of background risk, something which is additive, is given in Case 2 and shown in Figure D.3. In States 1 and 2 it raises the funding ratio to 111% in the absence of investment in risky assets. Liabilities improve. There are changes in mortality; changes in labour; labour risk; changes in policy, whatever. Things improve. In States 3 and 4 it gets worse. The funding ratio goes down to 91%. These are random changes to

Case 2: Background Risk

- State 1: 111% solvency minus y
- State 2: 111% solvency plus $y \cdot 190\%$.
- State 3: 91% solvency minus y
- State 4: 91% solvency plus $y \cdot 190\%$

11 Units Get Purchased – More than a Tenfold increase in risk exposure!!

Figure D.3. Case 2, background risk

the solvency level outside the control of the person making decisions about the scheme. Again, we have in States 1 and 3 asset markets which have returns of minus 100%, and in States 2 and 4 we have positive returns. What happens? How much is invested in risky assets? Eleven units get purchased, which is more than a tenfold increase in the risk exposure with the introduction of background risk.

In the simplified setting where there is one source of risk, where the risk is additive and there is only one decision variable, the actual effective background risk depends on the convexity of risk aversion. In other words, just as precautionary saving is a property of a third derivative of the utility function or the rate at which the curvature of preferences changes, the effect on asset allocation depends on the curvature of risk aversion. That is complicated, and it is even more complicated to measure in practice. So, it is a very difficult thing for us to consider.

The reason why it depends on convexity is that we are considering shocks which perturb the value of a background risk variable. If the risk aversion level is convex, that means, on average, that risk aversion is going to be higher, and risk aversion, when it is high, means that you invest less in the risky asset.

In practice, matters are much more complicated. As we argue in the paper, in practical pension application there is a mixture of multiplicative and additive risk. There are multiple sources of risk and there are multiple decisions being made. Therefore, it is important for us to consider more complicated models.

In the paper, instead of having one really complicated model which included things like covenant risk explicitly and credit risk of the sponsor explicitly, we chose to present just three, simple models.

The results from these models do confirm what I said earlier, that the effect of background risk on asset allocation can go either way. Surprisingly, the effect on funding can also go either way in the models which we consider. We can consider more complicated models, but I think that it is a strength of the paper that we considered a variety of models rather than just one.

Mr M. Cardinale (Institute Affiliate; introducing the paper): In the paper we first considered a two-period model. We assumed that the pension scheme chooses an equity allocation for period 1 and for period 2 and chooses also a funding level for period 1. We implicitly assumed that the deficit at the end of period 2 is covered by a funding inflow at the end of the period. This scheme is 85% funded, which is broadly in line with recent U.K. experience. We considered the case of deterministic liabilities (no background risk) and the stochastic (with background risk) case. We assumed that background risk is uncorrelated with the equity returns. Figure D.4 shows one illustrative result. Essentially, what we show there is the optimal equity allocation when we keep the funding fixed. We also consider in the paper the case of optimal funding, with fixed equity allocation and the case of variable funding and asset allocation.

In this particular case, what we show is the optimal equity allocation implied by our simple two-period model. Figure D.4 shows how the optimal equity allocation changes for different values of what we call the period two funding aversion parameter. What this parameter

Optimal equity allocation with fixed funding for different values of period 2 funding aversion (Two period model)




Figure D.4. Two period results, optimal equity allocation with fixed funding

demonstrates is that the greater the value of the parameter the more the fund is averse to deficit at the end of the second period.

Essentially, what we show here is that the higher the value of the parameter the more we put in equities, because we assume that the scheme is only 85% funded at the outset. However, more importantly, what happens when we include background risk? Essentially, background risk in this particular case leads to a higher equity allocation. So, this is an example in a simplified setting, where adding unhedgeable background risk leads to higher equity exposure.

So, we have a conclusion regarding the risk budgeting process. We show that, even in our simple models, risk aversion changes with the level of funding. The shape of this relationship could well be scheme specific, because different schemes may have different preferences. So, the implication here is not just about risk tolerance — we have to have a more complex picture about how risk tolerance changes with, for instance, the level of funding.

I also consider making a joint decision as regards funding and asset allocation. As we demonstrate in the paper, if we keep the asset allocation fixed, we might get one result for funding, but if we consider funding and asset allocation together, we might get a different result. So, it is important to stress how asset allocation and funding should be considered as being part of the same problem.

Also note that background risk, even in a simplified setting, such as the one on which we focus, does seem to matter for fair value liability and funding calculations. If that is the case, then perhaps this should be explicitly taken into account by the valuation process. Of course, this might imply that a pension scheme more exposed to background risk might end up with higher liability estimates. There are many examples of this. An open versus a closed scheme may be one example of higher exposure to real wage growth, and there are many more examples.

Mr M. Weale, C.B.E., Hon F.I.A. (opening the discussion): This is an excellent paper which draws attention to what has to be a key issue for anyone involved in running or advising a pension fund. Background risk is always present. There have to be divergences between the sponsors and the trustees and, possibly, also between the trustees and the beneficiaries. At the

very least, I can imagine circumstances where there are different categories of beneficiaries with different interests, and the trustees have to do some sort of balancing exercise. That final aspect certainly was not discussed, and I did wonder how it was done. However, judging this paper against currently popular scenario analysis, I think that it is a great advance. It tries to say, not only what might happen, but how big the probability associated with particular things happening is, and, on that basis, it tries to advise on a response to the risk which is faced.

In looking at the paper, I found myself asking: "How do the models in it relate to the sorts of things that I am familiar with as an economist?" Economists, of course, have famously narrow minds. The structure which I want to consider is where people are worried about consumption in period 1 and period 2, and can allocate their salaries at the end of period 1 between a safe asset and a risky asset. Overall, the welfare depends on the utility: the welfare which we get from consumption in period 1; and the discounted expected welfare which we get in period 2. That is the decision problem.

Let us now suppose that the survival probability is not correlated with the return on the risky asset. Then, at least in some circumstances, the problem separates quite nicely: "How much do I spend in period 1 and how do I allocate my residual portfolio?" The specific circumstance in which this happens is when the utility function has constant curvature. This is a structure which is popular with economists, in some sense, because it has the property that the degree of relative risk aversion is the same everywhere. This is rather appealing, because why should a large pension fund have a different portfolio structure from a small pension fund if, in other respects, it is the same? I can see that, if you have a small number of beneficiaries with large pensions, at the very least the structure you might want will depend on their circumstances, and that might be different from one where you have many beneficiaries eating into a pot of the same size. However, in more general terms, why should the portfolio structure of a pension fund be size-dependent? You might say that the examples which Dr Orszag gave were normalised, so that they were not size-dependent, but, in what I regard as the standard framework, you can separate out the background risk from the investment decision.

Are there any examples where there might, plausibly, not be that separability? Dr Orszag described the case where the risk was additive. I would say that it being multiplicative seems inherently more plausible, because the background risk, which here is uncertainty about mortality rates, is not normally additive to uncertainties about the return on equities.

On the other hand, I can think of some examples where you would get an interaction. Take the Black Death. The Black Death was a shock to the mortality rates, but, at the same time, it sharply reduced the workforce. In the aftermath of the Black Death wages went up, rental rates on land fell, and, presumably, land prices also fell. Therefore, if a sort of Black Death type background risk was something about which an actuary was concerned, then she or he would also have to take account of the interaction between the mortality risk being different from what was expected and what was going to happen to the return on assets should that eventuality materialise.

In the four-state example attributed to Gollier (2001), the welfare function was kinked. You might say that, if the welfare function was kinked, then the utility function which is being analysed cannot be that of the beneficiaries, because it is difficult to imagine why beneficiaries should have a kinked welfare function. If my pension falls short by 1% of what I had expected, one might think that my reaction to that would not be very different from my reaction if it exceeded what I was expecting by 1%. As those shortfalls or excesses shrink towards zero, I would expect the divergence between the two to disappear; whereas, in the example quoted, the function was kinked and the structure of the optimum portfolio was essentially driven by the nature of the kink. Indeed, without that kink and because the function was linear, there was no risk aversion. So, it was something which may be a sensible approximation or teaching tool, but on its own it does not strike me as being a realistic representation of the situation which I might face as a beneficiary.

The sort of general thrust of the paper is, however, that the link between the background risk

and the portfolio risk is the norm rather than the exception. From my point of view, this is an issue: "Is the structure a useful and general one or is it a specific case which actuaries are unlikely to find themselves facing?"

Leaving aside the question of the Black Death, one can think of some more up-to-date reasons why the situation which I was describing might not arise. If you have a pension fund which is nearly insolvent, then the trustees might well take the view that they may as well invest the remaining assets on a racehorse: if it wins they will have done a good job and rescued the pension fund; and if it loses they will be bailed out by the new PPF anyway. This, of course, is how moral hazard can distort choices, and is why you might expect a departure from the structure which I have defined for a very poorly funded scheme. I am a little less clear why you would expect it from a very well funded scheme: "Why might you expect to have the kink at the point where the scheme is reported as being notionally solvent?"

Of course, the reason given is that having a scheme which is not fully funded is an embarrassment; having one which is fully funded is not an embarrassment, and that is worth something to you. I suppose that I can relate to that, because I am Director of a non-profit-making organisation, the National Institute of Economic and Social Research. I am embarrassed about turning in a small loss and I worry about what my trustees will say. If we turn in a small surplus I am perfectly happy, and so are they. So, I can understand this asymmetry, but I come back to the point that it must indicate a departure from the beneficiaries' interests.

Another issue about which I would like to hear more is the question of the sponsor's perspective. The model rightly suggested that, if we have background risk, then where people are risk averse in the framework considered in the paper they will raise funding. This is what you would expect to happen, but it raises the interesting issue which is discussed to some extent in the paper; the sponsor actually has a choice: "Should the sponsor agree to raise funding or should the sponsor try to keep reserves outside the pension scheme?" One can think of all sorts of reasons why sponsors might choose the latter.

The most obvious one is, in some sense, a fallacy. It is like the situation faced by students running a May Ball, who decided not to insure all the marquees against fire, because they were not all likely to burn down. So, they paid a partial premium, and were then somewhat surprised to discover that, when one marquee did burn down, the insurance company wanted to pay out only a part of the loss. In fact, the insurance company needed quite a lot of persuading to pay out on the whole of the loss and were persuaded to do so in the end by the college, not by the students.

The sponsor is likely to think: "There are a number of things which can go wrong. If I keep the reserves outside the system, then I can use those reserves to hedge against a number of risks. If they go into the system, then they are lost within the pension scheme." The fallacy with that is that, if the risks which the firm faces are independent of each other, it cannot use a single pool of money to address all of them. Equally, the trustees do have to think a bit further ahead than the two-period model which Dr Orszag described. There is a reasonable argument that, if funds have surpluses, then trustees and sponsors between them tend to reduce contributions and may also raise payouts to beneficiaries. We saw both of those going on. In the 1990s only, perhaps, the university superannuation scheme made the argument: "Yes, the fund is in surplus, but something could go wrong, and therefore we need to hang onto the surplus instead of dispersing it in some way."

So, another important message for actuaries not made in the paper is that, given background risk, you do not just raise your funding, you do not just see it on a two-period basis; instead you have to accept and defend a situation where funds are normally in surplus because they need to keep reserves to address background risk. This is the same way as banks need to keep reserves to deal with things going wrong, and the normal losses on loans which banks face should be dealt with by the normal risk premium which they charge. Reserves for banks and pension schemes are needed to deal with unusual or background risks. These background risks, by their nature, do not appear every year or every decade, and therefore the public, the sponsors, the trustees and the beneficiaries need educating about an appropriate response to them.

Mr A. D. Smith: I welcome this paper because the implications of non-tradable risks and of incomplete markets are becoming increasingly important, both for insurance and for pensions work. It is helpful to have an exposition of the precautionary saving phenomenon and how it relates to the third derivative of utility.

Paragraph 4.2.3 acknowledges the existence of several parties to a pension plan, in particular the members and the sponsor. The opener has already pointed out that the interests of trustees might differ from those of active members or from deferreds and pensioners. I would add distinctions between corporate managers and shareholders, as well as considering the incentives facing consultants and fund managers. As ¶4.5.1 indicates, and the opener confirmed, the classical approach is to ascribe utility functions to each player. You then look for an equilibrium in which no players can increase their utility, given their rational views about other players' strategies.

In contrast, the paper takes a novel approach to multi-party analysis. Paragraph 4.2.3 proposes a bargaining process, the result of which is apparently a single consensus utility parameter shared by all parties. However, I could not find the assumptions about bargaining power which have actually been made to derive the utility functions which the authors propose. As the results are sensitive to the choice of utility function, the empirical calibration is a significant omission from the paper. How might this calibration be carried out?

I found ¶4.2.1 rather odd. The paper appears to argue that utility functions are only relevant when markets are incomplete. I wonder where the authors gained this impression — regrettably no references are given to justify their view. In the same paragraph, I believe that the authors are mistaken in equating the corporate approach to the assumption of market completeness. The corporate approach is about recognising multiple stakeholders with multiple investment opportunities; it does not require all risks to be traded.

I did not see where the authors recognise the possibility that fund stakeholders may have risks and investments elsewhere. This point does not only relate to investment risks. For example, from a pension plan perspective, improving pensioner mortality is a bad thing because it creates deficits. From a personal perspective, most of us would be happy to discover that we are going to live longer. So, the sign which we attach to that risk could be different if we take into account the stake which I have in all sorts of other aspects of my life.

The choice of time horizon deserves more discussion. When I was training to be an actuary, I heard about the importance of the long term for pension funding and investment. Dr Orszag described how classical financial assumptions may break down for long-term financial institutions. The description of the model in ¶4.5.3 appears to be a multi-period model, but I am unclear as to how many periods have actually been used in the results presented. Section 5.4 suggests that the results are based on one or two-year horizons. Although I am no longer convinced by all the traditional long-term arguments, it does seem odd that a paper on longevity risk and equity investment should focus on one or two-year horizons.

It is slightly frustrating that the paper's findings are not conclusive. Dr Orszag already mentioned how allowance for background risk either reduces or increases the optimal level of funding. It also either increases or reduces the proportion of the fund to be invested in equities. Even in my most curmudgeonly moments, I struggle to disagree with all the conclusions in this paper simultaneously! I wonder whether the results tell us as much about the authors' interesting choice of utility function and time horizon as about background risk.

Professor J. Cocco (a visitor; Assistant Professor of Finance, London Business School): I shall consider concepts, and highlight the important things in this approach rather than try to attack each assumption used in the model. Every assumption has its own problems. One particular assumption which was previously discussed is the assumption of the kink in the preferences (see Figure D.1). What is the idea behind this? Increases in value above the kink are less important than decreases below it. In other words, going 1% up matters less than going 1% down. So, we are penalised by losses; we do not gain as much on the upside. That is the idea of the kink. Why is this important? Why do we think that sometimes there should be a kink in preferences? Think about the firm's perspective. There is a possibility that the firm will become bankrupt with

all the costs associated with that. Do you want to take the chance that the firm will go bankrupt? No! If we are close to bankruptcy, then we will lose more on the downside than we gain on the upside. So, there is going to be a kink there.

Think about my saving for my retirement. Throughout my lifetime I will get used to a certain lifestyle. I will get used to my car; I will get used to my house; I will get used to consuming goods; I will get used to going on holiday. Is that going to introduce a kink in my utility function? Yes, it will. I have a subsistence level and I want to be able to maintain my standard of living. Kinks in utility functions are important from the point of view of firms, and they are important from the point of view of individuals. If you think that people get used to their lifestyles, then kinks in utility functions are very important. This is the first reason why I think that it is important to analyse kinks.

What is the problem when we start analysing kinks? The first problem is: "Where is the kink?" Dr Orszag said that the kink is at 100%, meaning that we want 100% funding of my pension plan, meaning that above 100% I gain less than I would lose if I go below 100%. Who decides on this 100%? It is not clear. That is a difficulty with this kink utility function. Why is it 100%? Why is the slope on the upside one and on the downside 0.5? Who decides these things? Is it the company? Is it the trustees? If you are concerned about bankruptcy of the company, clearly the kink should not be at 100%.

Once we accept that we have a kink utility function, then it is very easy to see why the results can go in different directions. Suppose that we are just above the kink, that there is no background risk and that we are investing in equities. Because there is no background risk, there is no possibility that we will go below the kink.

If we now add background risk, it becomes possible that we may go below the kink. Will I want to change my asset allocation? Yes. This helps to illustrate that, if you accept that there can be a kink in preferences, the results can go in all sorts of directions. Does this mean that we should ignore them? No, it means that we should try to understand the reason why they go in any particular direction.

There are things which you can do about the kink. For example, you can buy options to avoid going below the kink. If you are also worried about the kink in preferences, you can invest in safe assets to avoid the background risk, together with the risk of equities.

One thing missing from the paper is any comment about real estate and real estate returns, but I have done some work in this area. If you have a company with a defined benefit final salary plan, one important background risk of these sorts of schemes is wage growth. If there is this source of background risk, having real estate as an asset can be very valuable, because the returns on real estate are very positively correlated with earnings growth. I would encourage the authors to expand their work to include wage growth as a background risk, and to incorporate within their model another ingredient, which is real estate returns. They could compute the correlation between the two. It is likely to be very large and positive, and could give some interesting implications for the proportion of pension fund assets which should be held in property.

Mr D. G. McCarthy, F.F.A.: I think that background risk is a really important issue, and a complex one for pension funds. I would not like to quibble with the assumptions of the model, but I agree with previous speakers that the results of the model should not be taken literally. The conclusion that background risk can either increase or decrease funding and equity investment in pension plans is probably correct. I expect that, in practice, background risk will probably increase funding and decrease equity investment, which is different from the authors' view. The reason for this is that the major source of background risk faced by most members is the failure of the corporate sponsor, which interacts in a very real way with the funding level of the plan. This interaction is not allowed for in the model which the authors have chosen. Given that equities are likely to perform poorly when corporate defaults are high and that corporate defaults are highly cyclical, this would further imply that the optimal investment in equities is reduced. However, that having been said, I think that the paper opens the study of a very interesting area for pension funds.

Professor W. Perraudin (a visitor; Head of Finance and Accounting Group, Tanaka Business School, Imperial College, London): I have two points to make. The first is that I think that it is very important to consider carefully what the relative magnitudes of the different risks are. If we have hedgeable risks on the one hand and what we are calling background risk as non-hedgeable risk on the other, then I think that it is really important to reflect on their relative magnitude. It is difficult to do that in a convincing, statistical way; but, at the very least, we can look back on what has happened in the last few years. After all, everyone here will be very interested in the issues of asset allocation and funding levels for pension schemes. This is particularly so because of the problems which pension funds have experienced with deficits in the past few years.

I was, at one time, a special adviser to the Bank of England. We did a careful breakdown of the relative contributions to deficits in major schemes coming from changes in tax policy, changes in demographics, perhaps unexpected shocks to demographics, and changes to market prices. Our estimates, I think in common with some other estimates which have been produced by the industry, suggested that the big bulk of the problem in the last few years was the double whammy of a fall in the equity market coinciding with a fall in interest rates. Normally, you would expect a fall in interest rates to be accompanied by a rise in the equity market.

That double whammy of hedgeable market risk shocks has led to the problems which we currently face. So, on the one hand I am interested in the point made in the paper that background risk should be taken into account in designing optimal funding and optimal asset allocation decisions, but, on the other hand, I am somewhat concerned that the paper might be obscuring the fact that many of the current problems stem, perhaps, from a misunderstanding of the hedgeable risks which were faced by some schemes.

My second point is to say that this distinction between background and hedgeable risks is helpful, in the sense that it pushes us all to reflect on why background risks are 'background'. Why are they non-hedgeable? In a sense, these risks are not non-hedgeable, because insurance companies in the past have been willing to provide bulk annuity buyouts of different sorts, not necessarily on deferred liabilities, but at least on some of the liabilities. Why has that access to hedging not grown? Why is it that the insurance companies and then the reinsurance markets have not been able to provide a conduit for pension schemes to hedge these risks? We have to worry, to some extent, about the regulatory environment which insurance funds face. We also need to worry about the way in which actuaries value some of these pension liabilities.

When I have talked to firms which have been considering whether they should try to go for an insurance route to reduce some of their liabilities, they are impressed with the fact that often the estimates of the values of the liabilities held within the scheme are much less than the price which the insurance companies will put on them. It may be that the insurance companies, because of capital constraints and tight regulation, have to price these very high, or it may be that the schemes themselves should really be advised that these liabilities are very expensive, and hence the cost of getting insurance should be accepted.

I think that the way in which we think about the valuation of liabilities is all part of the same issue of why these liabilities are not marketed more.

Dr S. Satchell, Hon F.I.A.: Regarding background risks, Professor Perraudin addressed the issue as one of definition. I think that this is an interesting point, one which needs to be clarified. It is quite clear that any background risk which can be quantified can be insured or a derivative contract can be written on it. It is not really that which is the issue. I think that the important thing is whether the underlying entity represents something which is tradable. That is the way in which I have tended to analyse it in problems.

Background risks have been, in this context, either actuarially fair in terms of the utility analysis (meaning that, on average, their mean is zero), and, because of risk aversion, they are 'bads', but one can think of background risks which are 'goods'. So, for example, I would regard consumer confidence as a background risk. You cannot trade it. Nevertheless, when it is up that is generally good for utility, and it is not necessarily mean zero either.

How should we model it? One of the criticisms made against background risk, which is one which I share, is the fact that, when you read papers on it, such as this one, the utility functions tend to be a bit bespoke. You have a problem, and you build a utility function which addresses that problem. There is not really a general framework, and one is always worried how sensitive the results are to the way in which utility has been specified. That is, of course, a concern. So, one wants to look at a more general framework to address this.

How should we deal with all the uncertainties? There are many others than the ones mentioned. We could regard background risk, to some extent, as involving things like model uncertainty and parameter uncertainty. These are all things which we cannot trade, and they all affect the expected utility decision. The popular resolution of this, certainly in the fund management sector, is to look at things like robust optimisation. This is because asset allocation is, in some sense, the solution to an optimisation problem. Robust techniques give some promise of being able to bring into play some of the uncertainties associated with their inputs.

Mr C. A. Speed, F.F.A.: I agree with previous speakers who talked about the relative materiality of background risk and tradable risk and also about the importance of convergence between actuarial prices and market prices where we do have tradable risks.

Concerning background risk, I want to consider the materiality which it has to pensions. To my mind, the liability for a pension scheme is the accrued liability based on current salary. In other words, wage and labour market risk does not appear. In the past it has been argued that using projected salaries reflects how work is carried out in practice, but consider the current environment. We know that the debt on the employer, if it walks away from the scheme, is the accrued liability. We also have the pensions regulator telling us that the triggers for funding targets are going to be based on accrued liabilities — that is not considering future salaries. So, I think that there is a strong theoretical and practical argument for looking at liabilities to date without allowing for future earnings growth.

The paper is left to address demographic risk, in particular, mortality. Even here we have a partial buy-out market, and there is the possibility of reinsurance and maybe even capital market products. The one big background risk which I would have liked to have seen more about is the background risk which employees have of their company going bust, which they cannot hedge.

In Section 4 we find the justification for the approach to all the subsequent modelling. While much of the financial literature applied to pension schemes recognises the corporate structure, where the trustees and the corporate managers are agents acting for the principals (shareholders and members), this framework is rejected. Instead, we have a game theoretic approach, where objective functions are used and applied to the pension fund.

Thinking about this, we are using objective functions, not only for groups of individuals who may have very different financial backgrounds, but also for different groups of individuals (shareholders and members), who may, at times, have divergent interests. This is a hard task for an objective function.

The particular objective function used considers the deficit and the inflows to the pension scheme where large inflows are seen as undesirable. In Section 4.2 we have the assertion that trustees do not want to have sudden jumps in the financing required by the plan. My reaction to that is: “Why not?” A sudden jump from 60% to 90% funded seems totally consistent with the fiduciary duty of the trustees to pay benefits and to deliver greater security to members.

The approach also has the potential for giving rise to somewhat confusing results, so, when we look at ¶6.2.3, we are told that the optimal funding goes up from 51.71% to 52.56%. I am left wondering: “Is this really optimal?” Again, think from the point of view of a member. We have been told that a scheme is 51% funded, therefore the remaining 49% of the pension promise is essentially invested in a non-diversified, unsecured long-term loan to the sponsor. There are few or no loan covenants and the repayment schedule is ill-defined. Is this really optimal? I struggle to believe that it is.

We are also shown, in Figure 7.1, that, for many values of the risk aversion parameter γ , 100% equities is optimal even without background risk, despite the acknowledgement that the pension fund liabilities are bond-like. I am left with the question: “What happens if,

in this game, there is an important player which has been missed?" If we have this important player, there will be an optimal strategy where the two key players which we are considering could work together at the expense of the third player. That seems to me like the situation which we actually have. The tax arbitrage approach, which has been advocated, is totally missing from the framework, so, when we remove background risk we do not get back to the 100% bond strategy which has been advocated as optimal in much of the financial literature.

This is a great paper for bringing background risks to the attention of actuaries. However, I think that, in current times, the interaction of capital structure and the tax effects are more important. Whereas I welcome bringing background risk to the fore, I still have concerns about the particular framework which the authors have chosen to present in this paper.

Mr P. N. Thornton, F.I.A.: It strikes me that such a huge amount has happened in the evolution of pensions in the past ten or 15 years that this area deserves much more research and exploration.

For example, if you went back to the 1980s, pension funds were usually in surplus. The utility function was then completely different. It is true that employers took advantage of contribution holidays, but normally the trade-off for that was benefit improvements for the members, and the trustees often regarded themselves as having a duty to negotiate benefit improvements for the members in return for the contribution holidays. So, the interests of the different parties were very different.

They were also very different in regard to the impact of potential deficiencies, not that many funds had them in the 1980s. The accounting treatment was different, and employers were in a position to fund deficiencies in a way which did not cause enormous volatility in their financial reporting. So, when you go back in time, the utility of surplus or deficiency was very different from what it is currently. When you also go forwards in time, everyone at the moment is trying to assess what impact the new regulations will have, and every trustee and every sponsoring employer is trying to decide where the balance is going to be as to what the funding strategy should be. I think that this is going to change the utility of surplus and deficiency again from where it has been more recently.

One of the interesting issues for me in how pensions have evolved in this country is how we have got ourselves into a position where, ideally, employers should now be able to renegotiate the liabilities and then fund them on a completely different strategy from that which they embarked on historically. In other words, when all of these final salary liabilities originated, the funding strategies which were acceptable were very different from those now. We have been moved into a different era in terms of what is regarded as appropriate, not just because of accounting standards, but, perhaps, because of greater awareness of some of the risks which are involved.

The Chancellor of the Exchequer, at the CBI conference held earlier today, said that we should avoid gold-plating regulations. I think that, without realising it, we did gold plate pension fund regulations when we introduced the pension increase legislation.

When you compare Britain now with all the other European countries, there is nowhere else which has the same level of guarantee attached to pensions in payment in terms of inflation protection, and it is that which is causing us so much difficulty. One of the nearest examples of countries with big pension funds is the Netherlands. There, if the funding will not support full inflation-proofing of the pensions, then a business plan is agreed with the regulator for restoration of full funding, and until it is restored the pension increases are cut back, so that the pensioners share some of the pain. If we were able to rewrite history, that is the one piece of legislation which I would revisit, because I think that it is the one which has really tied us up in knots.

I have a question for the authors: "I was not left very clear by the paper as to what conclusions this led, and whether you felt that you needed to do more work before reaching conclusions on the extent to which it is worth hedging hedgeable risks in the presence of unhedgeable risks?" I also think that we should rename this paper 'Unhedgeable Risk' rather

than 'Background Risk', because, clearly, mortality is more than a background risk, and the real issue is whether or not you can hedge all of the risks involved in pensions.

Ms C. E. Hobro, F.I.A.: Concerning the kink in the utility function, I disagree with the opener and agree with other speakers that, from an individual point of view, I am sure that there is a kink. I heard a proof of that when I went to a lecture on happiness at the Royal Institution. Many economists, social scientists and psychologists have been working together to look at this. Essentially, you get used to a certain way of living, and you get very unhappy if something happens to pull you down from that level. You get slightly more happy when things improve, but then you rapidly get used to that new level, which no longer makes you happier.

As an example, the lecturer said that, when you get married there is a peak on your wedding day when you are really happy. However, about two years afterwards you are maybe slightly happier than if you were just single, but that there was not really that much in it! The lecturer said that economists do not tend to recognise this effect, which is commonsense to everyone else!

I now consider companies and how they feel about their pension scheme deficits. The paper did not mention that there is a connection between how the trustees and the sponsor feel about their deficit and the average level of scheme deficits in the rest of U.K. PLC. I have been to several presentations where the trustees and/or scheme members, although they have a scheme deficit, are definitely not as unhappy as they would be if they had a deficit and everyone else had a surplus. The company comforts them by saying: "Well, our scheme is not much worse than everyone else's. Everyone has had the same problems", and they manage to make their members a bit happier than they otherwise would have been when hearing of the deficit. So, it is certainly not a case that we are advising a pension scheme which is viewed completely in isolation from everything else. I am not quite sure what that means in terms of the advice which we give. Maybe it just means that our clients are not so unhappy with us, as advising actuaries, when their schemes go into deficit despite our advice, as long as other pension schemes also go into deficit.

Dr D. C. Bowie, F.F.A. (closing the discussion): I thank the authors for a very good paper, which has clearly involved much detailed work. They have summarised their intentions as being threefold. The first was to put background risks more firmly in our thinking of the risk management of pension funds. The second was to exhort us to consider funding and investment strategies jointly. The third was to encourage us to spend more time thinking about decision-making frameworks for pension schemes and for the stakeholders of pension schemes.

Most of the comments and the discussion have focused on the third comment, that is: "How does and should the decision-making process for pension scheme funding unfold?"

Much of the discussion stems back to the way in which thinking about pensions within the profession has evolved over the past several years. The discussion divides into the usual two camps of the more scheme-centric 'equity equals lower cost' camp and, on the other side, the 'FE (financial economic) approach' (as termed by the authors).

In particular, the authors claim that the FE approach is flawed because of the existence of background risks, which make the market incomplete. The incompleteness means that the FE approach does not yield a unique set of (arbitrage-free) prices. The authors choose to deal with the incompleteness introduced by the background risks which, by specifying, as many speakers have pointed out, a single bargaining objective function for the pension scheme single objective function approach. That beds the authors' approach largely in the scheme-centric type of mould.

Several speakers contrasted the idea of having many different stakeholders, each with its own utility functions, negotiating to arrive at some sort of outcome, with the authors' approach of specifying a single objective function. Several speakers also discussed quite extensively the features and the consequences of decision making using the kinked utility function which the authors used in their introductory example.

However one categorises this paper, it has applications under both the scheme-centric and an

FE approach. However, it appears mildly contradictory that, in the detailed examples which the authors constructed in the paper, they did not end up with the all-bonds and full-funding answer which they had initially explained to be a consequence of the assumptions made in the FE approach. This seems to indicate that it is the decision-making framework rather than the 'background risks' which are changing the conclusion.

Many of the speakers have applauded the authors for bringing background risks to the fore. The paper adds considerably to the other papers which have already been published within the actuarial literature on second order effects, such as tax, default insurance, sponsor covenant and risk sharing between the different stakeholders.

One of the key things which the paper does for us is to draw together some results from a literature with which we might not all be *au fait*, namely the economic literature on background risks. The discussion of the economic literature gives us an alternative theoretical context, known as precautionary saving, for a popular actuarial reaction to unknown risks which can be summarised as 'hold a reserve'.

Mr Weale and Mr Smith both suggested that there were other risks which should be taken into account, which can be classed generally as agency risks. These represent the disconnect between trustees and members, between management and shareholders, etc.

Mr Speed, Dr McCarthy and others have also mentioned that some of the most important non-hedgeable (background) risks were tax and sponsor covenant. These have not been tackled explicitly in the paper. It might be that the authors envisage that these should be built in indirectly through the parameters in their objective function. Mr Speed specifically suggested removing wage risk as a background risk, because it was not relevant to the value of the pensions being considered.

Perhaps a more fundamental issue which cannot be dealt with very easily within the framework is the treatment of funding. I took from the paper that the model placed a lot of focus on the timing of contributions into the scheme rather than on the timing of the accrual of the benefits. Of course, both of them do feed through to the objective function, in the sense that both the deficit and the contributions appear on it. In ¶4.5.10 the authors specifically suggest that it is the timing of cash flows into the scheme rather than the accrual of benefits in the scheme which changes the economic position of the sponsor.

Clearly putting money into the pension fund does cost the sponsor something, as the opener, for example, concluded, but I think that it is different from the cost of the benefits accruing which equates to most of the economic cost.

During the discussion mention was made about the sponsor losing control of, or flexibility in, the use of the money once it is paid into the scheme. Some of these costs include the ability of the sponsor to default on any deficit should the company enter bankruptcy. Also, there may be other more subtle influences on the costs related to frictional elements, such as acquiring liquidity quickly enough to avoid bankruptcy or to take on new business opportunities. On the other hand, Mr Speed mentioned the fact that there may be huge tax advantages to having a fully funded scheme, particularly for strong sponsors. Having a fully funded scheme may also alleviate some other frictional costs, such as having to negotiate with trustees or the regulator when undertaking corporate activities which might affect member security. The Cobb-Douglas model, as used by the authors in this context, strikes one as more of a descriptive model than a normative one. In other words, it seems to describe more the outcome of the negotiation rather than the way in which it should be undertaken.

With that in mind, I thought that this paper could probably have said a great deal more, in Section 8, on establishing the fair value, which is the outcome of a negotiation far more, perhaps, than on how a pension fund should be funded and invested.

Mr Speed also mentioned that underfunding is the equivalent of a concentrated equity investment. Seen in that way, reducing funding can be seen as increasing equity risk, and this perspective does not seem to be captured in the objective functions specified. In terms of the details of the model, Mr Smith and others made comments about how unstable the results seemed to be. Indeed, the authors bring up that point by suggesting that the calibration and the choice of model are absolutely critical to any conclusions reached. The inference that may be

drawn from this is that simple models do not suffice in extracting clear insights into how pension funds should deal with background risks.

Professor Perraudin, Mr Speed and Professor Satchell all mentioned that a great deal more research and understanding need to go into the nature of the different types of risks. Background risks are even more taxing, because it seems that you may not actually know that they have happened until some time after the event.

As a profession we need to be careful — and Professor Perraudin made this point particularly — not to be seen to be trying to sneak in or justify equity investment under the fluffy duvet of background risks. We have to acknowledge the fact that there were, and are, hedgeable risks out there, and at least we should have a full understanding of those and what can be done about them before we spend too much time or get too misdirected with the unhedgeable risks. It is a very tricky communication issue, when taken from a public interest perspective, to try to explain that the existence of background risks means that risks should also be taken elsewhere within the fund.

The opener and Professor Perraudin suggested that holding reserves, and thereby acknowledging the difference between the insurance value and the pensions value on all the liabilities, is a possible way forward.

We could do with much more investigation on background risks. Perhaps we should think about dealing with them more in terms of risk management than using them to justify taking other risks. Professor Cocco suggested that we could explore other asset classes to try to develop hedges, or at least partial hedges, against some of these issues. To the extent that this paper has brought all these issues to a head and enabled us to discuss them, it has made an important contribution to our understanding on pension funds.

Dr J. M. Orszag (Institute Affiliate; replying): I shall focus on three points, of which the first is kinks. Kinks are only referred to in the paper in ¶¶4.4.1 and 4.4.2. I have been quite surprised at the interest in them. It was just an example, but there are regulatory roles in the United States of America on funding which could induce kinks in objective functions. 100% funding is, perhaps, a reasonable kink for those who are very close to retirement, because they are less likely to get benefits which are above 100% — but you can take issue with that. It was just an example to motivate the idea that asset allocation is affected by background risk.

Next, there have been a number of comments on models, and one reason why we presented different models was to try to show what could happen in different settings. In terms of the issue about why, without background risk, there was not a 100% bond investment, there are two answers to that:

- (1) You could think of these models as models with a bit of background risk and then models with a lot more background risk.
- (2) It is not clear in the basic financial theory that you should be 100% in bonds or 100% fully funded until you take into account other considerations, such as tax.

If you allow for tax considerations and the tax shield, you could be 100% in bonds, but with PPF type liabilities you could be 100% in equities, and there is a knife edge between the two. In competitive financial theory there is no reason to prefer asset allocation. This goes back to Exley, Mehta & Smith (1997). It is considerations like tax which lead you to 100% bonds. Considerations, like member protection, lead you to 100% bonds, and there are economic models which justify under funding. For instance, firms do not want to have fully funded pensions because they can lead to workers demanding too high wage increases. There is a large literature in this area. In terms of the comment on one to two-year horizons, it is really a two-period time horizon. The periods could be quite long.

Many other detailed comments were made. The important point is that we considered a number of different models. One of them did have a game theory approach which is commonly used in labour economics, which is my background, but you could consider much more complicated models and you could have dynamic models. Paragraphs 4.8.1 to 4.8.4 consider, briefly, the implications of dynamic models, relating decisions about pensions to individual

decisions about asset allocation, given the human capital. I think that the points made about considering things outside the pension plan are very important ones.

The magnitude of risk has varied over time, and it could become more important because of the more mature nature of pension schemes. It is clearly something which needs more measurement and careful consideration.

We have to be careful to avoid overemphasising risk, as I think that the biggest issue is underfunding. The biggest loss to members' welfare is the fact that pension schemes are not fully funded. If you take an underfunded scheme and you invest entirely in bonds, you do improve the members' welfare, but do you improve it that much compared to a situation where you put more money into the scheme? I would argue that, yes, there is an issue about the magnitude of risk, but the mere existence of background risk suggests to me, at least, that we need to pay more attention to funding.

The Senior Vice-President (Mr P. W. Wright, F.I.A.): The subject of non-hedgeable risks, as Mr Smith pointed out, has been much discussed over the past five or six years in the insurance area. For example, earlier this year we considered the appropriate allowance for such risks in the specific context of a market consistent embedded value (O'Keeffe *et al.*, 2005). We have discussed the International Accounting Standards Board's insurance project, where, for both life and general insurance, we have discussed the appropriateness or otherwise of including market value margins for non-hedgeable risk. For general insurance, some years ago we had a discussion on reinsurance-to-close at Lloyd's, where again one of the biggest issues was the treatment of non-hedgeable risk.

I agree with Mr Thornton that we should not try to introduce new terminology for what is essentially the same thing. I hope that we can move away from talking about background risk.

We have been talking here largely about funding pension schemes. As a non-pensions practitioner, I can see that, maybe, if the accounting bodies are so keen on risk margins in the insurance field, there is no real logic for them not wanting to see such margins in accounting liabilities for defined benefit pension schemes.

There is also the thorny question of cash equivalents, and whether they should be increased, or possibly even reduced, in the presence of non-hedgeable risks.

All this has led to a lively discussion. I should like to thank the authors, the opener and the closer, and all those who took part in the discussion.

REFERENCES

- O'KEEFFE, P.J.L., DESAI, A.J., FOROUGHI, K., HIBBETT, G.J., MAXWELL, A.F., SHARP, A.C., TAVERNER, N.H., WARD, M.B. & WILLIS, F.J.P. (2005). Current developments in embedded value reporting. *British Actuarial Journal*, **11**, 407-496.

WRITTEN CONTRIBUTIONS

Mr M. G. White, F.I.A. (who spoke briefly at the meeting, and subsequently wrote as follows): Professor Perraudin mentioned that insurers were apparently not accepting some risks which were, in theory, hedgeable — and suggested that this might indicate that the liabilities and the risks are simply being undervalued by pension funds.

I am inclined to agree with this analysis, as in this situation there is no natural acceptor of the hedge. We can, of course, look through insurers to the owners of their capital. Most capital is held, directly or indirectly, on behalf of future and current retirees — precisely those people who are already exposed to mortality risk of their own, and would not naturally wish to increase their exposure to improving mortality. When the mortality risk of defined benefit pensions (whether deferred or in payment) is hedged, then, to a certain extent, it is being accepted by the owners of defined contribution benefits — and we can think of all savings wealth as being

‘defined contribution’ in nature, whether held within a pension fund or not. This argues for a very substantial risk premium for the hedge; indeed, the authors refer to mortality risk as being unhedgeable, which suggests that the willing buyer’s price falls short of that of the willing seller.

The authors subsequently wrote: Our paper on background risk was intended to show that background risk does matter — it affects asset allocation, optimal funding and risk-neutral pricing. The comments do not take issue with this so much as with some of the simplifications in our models. It is useful to respond briefly on some of these points:

- *Credit risk.* This was something which respondents thought that we should have included as a background risk. We modelled background risk generically, and there is no reason why our analysis and certainly our framework do not apply in the case of the credit risk of the sponsor. To a degree, credit risks can also be hedged by credit default swaps, but we accept that credit risk is an important background risk.
- *Bargaining.* We used a blended objective function, as is common in the labour economics literature, to avoid writing down a complex bargaining model. We experimented with different assumptions, but there is no doubt that we could have considered more complex models. However, all that we were trying to show was that background risk does not matter, and that this conclusion is unlikely to be different when additional complexity is added in.
- *Kinked objective function.* There was a lot of discussion around the appropriateness of our kinked objective function and where the kink should be. We accept this discussion, but the central result of the paper — background risk matters — does not depend on whether or not there is a kink in the objective function.
- *Other risks.* We acknowledge that stakeholders have other risks in their portfolios, but we do not see how that affects the essence of our results and the implication that background risk matters.
- *Other assets.* Professor Cocco mentioned including real estate as a hedging asset. We think that this is interesting, but again, while it might affect specific examples, it does not affect the fact that, generically, background risk is meaningful.

There were a number of comments about our baseline and why it was not calibrated to 100% bonds (Professor Cocco suggested this). We do not think that calibrating to 100% bonds is necessarily correct, both because of background risk in the baseline and also because optimal asset allocation need not be 100% bonds in perfect markets. However, we cannot see how this would possibly change our underlying conclusions either. Let us return to the example with the kink, and put the kink at 100% funding instead of 101%. Here, without background risk, the optimal asset allocation is 100% bonds. If we move to a background risk of ten, optimal asset allocation moves ten units of equity. Background risk matters just as much here as in the example in the paper.

Mr Smith raised a few issues about our comments in ¶4.2.1 on utility functions in perfect markets. In particular, he said: “The paper appears to argue that utility functions are only relevant when markets are incomplete. I wonder where the authors gained this impression — regrettably no references are given to justify their view.” While not pertinent to the conclusions of the paper, we were simply acknowledging what is behind Mr Smith’s own Modigliani-Miller analysis in Exley, Mehta & Smith (1997) and elsewhere. Where markets are complete, the corporate approach is above questioning. Where markets are incomplete, the standard corporate approach needs tweaking. That is all we were saying, and it does not affect our conclusions in any way.

Mr Speed and Professor Perraudin questioned the materiality of background risks. Unanticipated policy/regulatory shifts, such as the introduction of the PPF and the change in the dividend tax treatment in 1997, strike us as among the more material events which have affected U.K. pension funds in recent times. Mortality movements are also material in the U.K. One point raised was that wage risk does not matter so much because schemes can be shut for new accruals. The latter is just one possible approach and, to date, there are not many examples

of this (perhaps due to reputational risk). Another point is that some of these risks might be securitised in the future. However, they are largely not securitised at present, and it is unclear whether things will improve — risks which may be hedgeable now may become unhedgeable due to supply/demand imbalances (e.g. is it worth hedging inflation risk at 0.5% real rates or is the risk of selling these bonds at a capital loss larger than the inflation risk which we are supposed to hedge?). Ultimately, traded financial instruments are needed to cover risks, and, with a few exceptions, these are not there and relative prevalence may not improve.

We welcome this excellent discussion. Background risk is important for pension schemes, and this is an area worthy of further attention from the actuarial profession.