

ABSTRACT OF THE DISCUSSION

Mr C. J. Exley, F.I.A. (introducing the paper): The paper begins by considering why the cost of capital matters and why financial firms are special in relation to the cost of capital. We then consider the advantages and disadvantages of some popular measures of the cost of capital, and propose ways of improving these existing approaches. Particularly, we point out the advantages of frameworks which take full account of a firm's market value.

We explore some reasons why the cost of capital is relevant, and I will focus on one in particular — how much capital one wants to hold. Many firms' capital calculations are based on the notion of economic capital needs. In other words, most economic capital models focus on the benefits of capital in terms of low ruin probability. However, capital decisions can be improved by also considering the cost of this capital at the same time. After all, for most economic goods, the amount which you choose to buy depends on how much it costs.

Most of the literature on the cost of capital considers typical firms, whose businesses are financed by equity and debt issued in the capital markets. For example, the typical firm has plant, machinery, intellectual capital, and the like. Modigliani & Miller (1958) famously explained why the value of a firm is unaffected by how it is financed — in other words, showing that the value of the firm is the value of its debt plus the value of its equity. The result works because non-financial firms can often rearrange their finances, for example the ratio of equity to debt, without any immediate impact on the underlying business.

Financial firms are substantially different, however. This is because their capital providers are also their customers. For a financial firm, reallocating capital between business units may translate physically into cutting off existing customer relationships and initiating other relationships with new customers. So, including the debt in the cost of capital calculation is troublesome. Because of the difficulty in measuring the total cost of capital for a financial firm, many firms have focused on other measures, for example the return on equity (ROE) or, more recently, the return on risk adjusted capital (RORAC). However, there are also problems with these approaches.

First, considering the ROE; does the return on equity predict the total returns? Is it a good measure of the required shareholder returns? In Figure 4 we compare non-life total shareholder returns with all-industries, and find that the ROE of non-life businesses was below all-industries for every year from 1990 all the way through to 2004. However, on a total return basis, that is looking at the return on the equities, non-life companies outperformed the all-industries over this period. So, there is a bias in the return on equity measure which is very important for financial firms.

Another popular alternative is the RORAC. This measure looks at how much capital is required to allocate to particular lines of business to lead to a particular probability of ruin for that business. In Table 4 we have examples for Fortis, where different lines of business are allocated different amounts of economic capital. The problem is that these amounts of economic capital provide the denominators in the capital of return on risk-adjusted capital calculation, which means that businesses which have small allocations to economic capital apparently have very large returns, for example as shown in investment services. Taking this at face value suggests that you could add value for the firm by closing down the merchant banking business and, instead, applying capital to the investment services business, but, of course, the problem is that economic capital is not the only measure of the shareholder investment in the business. Investment services have involved substantial investment, building up the brand name and building up customer relationships. You cannot simply measure the shareholder investment in a business by the economic capital which needs to be held.

There is a missing component in these previous calculations, both in the ROE calculation and in the RORAC. The important missing component is the return on franchise value.

In the paper we identify three different components of shareholder wealth. The first is the economic capital, which is the capital required to achieve a certain low probability of default,

the second is the free capital above that point, and the third is the franchise value. We regard the value of in force business as part of this franchise value; the franchise reflects the investment of the shareholder in the entire business.

Using this framework, we can look again at the biases in other measures. To investigate this we have two examples: an example of personal accident business; and an example of catastrophe reinsurance. In the case of personal accident business the economic capital requirements are quite small, whereas in the case of catastrophe reinsurance the economic capital required is extremely large. This reflects the fact that the losses on £100 of catastrophe reinsurance premiums is more volatile than the corresponding loss on personal accident insurance. You might imagine that you need to charge a high profit loading on catastrophe reinsurance on account of this.

However, returning to the issue of franchise value, the personal accident market shows greater customer loyalty, and is heavily dependent on distribution. This is because much of the business is cross sold to credit card holders, and, once a policy has been started, the premium is a small monthly deduction which the cardholder hardly notices. The different nature of these businesses is reflected in the market capitalisation. The high franchise value of the personal accident business should be a clear signal to the business owners of the margins required in the premiums charged.

Our two examples, therefore, have the same market capitalisation. If the businesses were of similar risk, then the shareholder required returns would be the same in pounds in both cases. In fact, in this case the personal accident is probably higher risk, since it holds lower free assets in relation to the economic capital requirements.

The effect can also be seen in the new business effect on shareholder value. The effect on the net assets is a profit at inception from writing new business. There will also be an effect on economic capital reflecting any additional risks assumed. The subtle part is the effect on franchise value. Several aspects come into play here. The first is the familiar value-in-force (VIF) effect. The market gives credit for future margins not immediately released from liabilities. However, rather than being expressed in absolute terms, this needs to be considered relative to the value which the market already gave to this business before it was incepted. More often than not, particularly if there is a profit at inception, the comparison of the value in force to the previous allocated franchise value results in a franchise value write down.

The third effect is even more subtle. Our extra block of business has made the firm slightly more risky. This increases the rate at which the market discounts projected profits. You might think that this is a second order effect, but, in fact, it matters, because the reduction applies to all other business too, not just to the block which we are valuing. The discounting effect gives a rationale for profit deductions expressed as a cost of economic capital, which we commonly see used in pricing practice.

We describe how the franchise value can be modelled by looking at asset swap margins and liability swap margins. Our approach represents a combination of three important, but separate, branches of research over the last few years. The three components are: the economic capital calculation, which turns out to be important in determining the risk of default of the firm; frictional costs, which have been a source of significant research in recent years; and the franchise value, which is related to asset swap margins and liability swap margins revealed by the securitisation of business. Our framework combines these into a model of cost of capital for financial firms.

Figure 12 illustrates how this model can be used to establish an optimal capital structure for a financial firm. There are two contrasting conclusions from previous approaches. First, looking at the embedded value approach, because this involves discounting future cash flows using a return which is greater than the returns on the assets, you find that the more assets which you hold the lower the value of the firm. So, that is the downward sloping line.

Second, if you look at the return on risk adjusted capital, this assumes that, as long as you can keep reallocating capital to a business which has a return above its risk adjusted cost of capital, the more you allocate to that business the more you increase its value. Those two lines do not, unfortunately, intersect at a point which gives you the optimal capital, but what we

propose is a framework which balances two different concepts. The first is that the lower the capital you hold the higher the ruin probability and the higher the risk of loss of franchise value. The second, at the other extreme, is that the more capital you hold the greater the losses due to frictional costs. We find that we do, indeed, arrive at an optimal capital structure between those two extreme points.

In summary, cost of capital is important, but difficult to measure, especially for financial firms. ROE and RORAC are not always good indicators of shareholder returns. The important point is that shareholders demand a return on market capitalisation, not just accounting assets or economic capital. Allocated economic capital turns out to be relevant, not as an entire solution, but because it measures the distance-to-default, which is an important input into the full model. Finally, an optimal capital structure is revealed which balances frictional costs of equity capital against the default effect.

Mr R. A. Rae, F.I.A. (opening the discussion): Whilst the cost of capital touches all of us, many might not count it as one of our core competencies, and that is not surprising, given that a paper on the subject is long overdue. Anything less challenging than this paper may not have satisfied the academic fraternity, and anything more theoretical may have deterred those of us with a more practical view on the subject. However, what we have is enough for everybody to discuss.

Sections 2 and 3 provide ample material in this respect. By contrast, Sections 4 and 5 extend the discussion and offer an interesting and compelling framework for optimising a company's capital base, and one which will appeal both to academics and to practitioners in the field.

I read the paper from the perspective of the financial management of a life insurance company, rather than of a general insurance company or a bank. As a generalist, I found that the paper left me wanting to deepen my understanding of the subject. As the authors point out in their introductory section, the cost of capital matters. It interacts with the return on capital. Once a company has capital, it can only borrow more if it provides a suitably high return on capital. If it fails to produce that return, then shareholders look to extract their capital as soon and as efficiently as possible, the one exception here being the mutual life office, where there is unlikely to be enough membership power to force appropriate membership action to take place.

In ¶1.1.6 the authors point out that some regulators are considering whether firms should be required to hold capital to meet the cost of any capital tied up for solvency purposes, as, traditionally, the regulator's ability to retain capital within a financial institution does not require it to consider the effect on the providers of capital (as the authors go on to observe in ¶1.2.5).

Section 1.3 establishes that, in 'traditional corporate finance parlance', bank depositors and insurance policyholders would not be considered as providers of capital. Capital for life companies is straightforward, in that there is equity and subordinated debt and, generally, no senior debt (at least not within an authorised life company entity). Life companies have new business strain, which is essentially a liquidity issue resulting from the investment in new business, and, within classical corporate finance parlance, this can be financed with short-term debt or working capital. Given that, increasingly, life companies are looking to finance this strain through new business financing arrangements, (either on a cash or a cashless basis), it is worth noting that rating agencies recognise working capital, and might exclude this as a debt item from their capital models, reflecting its short-term nature. Clearly, though, it is an additional liability on the balance sheet, and I would welcome the authors' views as to whether there is such a thing as working capital, and its status as capital or just another liability.

The paper highlights the role which the cost of capital plays in traditional embedded value calculations, and notes how market consistent embedded values are now emerging as a preferable way of measuring the value of in-force. The development in financial theory has started to establish itself, and the authors explain, in Sections 4 and 5, that these developments extend these market consistent approaches to an alternative determination of the capital base and the quantification of 'franchise value'.

Section 2 questions the return which shareholders require and introduces CAPM. In

particular, ¶2.1.3 mentions that the weighted average cost of capital can be minimised by raising as much debt as possible. Raising more debt would increase the volatility of the equity return, and hence push up the cost of equity. This should give rise to a zero cost game, except for the arbitrage that debt is tax deductible. It is this tax relief which theory suggests makes it sensible for a non-financial institution to raise as much debt as possible. For non-financial institutions (i.e. 'corporates') there can be a point of inflection in the cost benefit equation, as one moves from being investment grade to becoming a sub-investment grade issuer of high yielding bonds. A major financial institution cannot entertain the latter, as its customer base relies on a strong investment grade capital base.

The reference to Hancock *et al.* (2001) describes the way in which replicating portfolios can be used to determine shareholder returns, and many United Kingdom actuaries will be familiar with this through their work on realistic balance sheets.

Franchise value is then introduced, together with the interesting concept of a franchise insurance premium (which is then carried forward to refine the capital model described in Section 4). The idea of attaching a cost to the loss of franchise value in the event of failure appears to have a sound logic to it.

Section 3 introduces the main theme of the paper, namely return biases and approaches to removing these biases. The idea that ROE ignores the implications of the capital sunk in a firm's intangible assets is intuitively obvious to life actuaries, given the importance of a life office's VIF in its embedded value. Rating agencies will recognise VIF as soft capital, and give, for example, 50% credit in its estimate of a firm's capital base.

Consequently, the inclusion of franchise value in the determination of total shareholder return is second nature to a life office actuary. Companies do set shareholder return targets, and there is a circularity which plays against management when it comes to life offices. Very often a life office's in-force portfolio is seen as a drag on shareholder returns. A traditional embedded value basis uses a shareholder discount rate to determine VIF. In the paper's terminology, a low discount rate will maximise franchise value and reduce a firm's total shareholder return. This gives rise to the question: "What value should you place on VIF to determine the franchise value and hence eliminate any accounting bias?" Adjusting a discount rate to achieve this is circular, making it clear how important it is to be able to calculate franchise value another way. Consequently, the development of market consistent valuations are to be welcomed.

RORAC can be a very useful tool in establishing the value, for example, of hedging a risk position against the alternative of backing it with capital. This gives rise to a discussion about how the capital released has been calculated and to what risk tolerance level, but it sets an objective target against which the costs of servicing this capital can be computed against the costs of a de-risking transaction.

The paper then moves on to asset swaps and liability swaps. It outlines how a life company could swap its asset cash flows for, say, LIBOR, and its liability cash flows for something less than LIBOR. I found Section 4 very compelling and, on the face of it, 'complete'. However, I would not consider myself accomplished enough to have the skills to critique its hypothesis rigorously. This methodology looks to have the potential to provide some valuable insights, such as indicating the optimal level of leverage.

Figure 12 shows franchise value increasing indefinitely with leverage under the RORAC model. Clearly this is the case if there are limitless opportunities to deploy it and earn the same RORAC. However, I see additional capital being deployed in such a way that it provides a progressively decreasing return on capital. There is a strong case that a company should keep on raising capital whilst its cost of capital is less than the incremental RORAC. Hence, a life company might experience a decreasing increase in franchise value as it borrows more, so as to include additional, lower yielding investments. Hence, in my model of the world the dotted line in Figure 12 would have a convex shape.

In terms of the formulae which the authors develop, I wonder if there would be 'irrecoverable losses of tax credits' — I am not sure that in insolvency you would necessarily lose all of this asset. A life company is likely to carry on in some form rather than be formally wound up, and, as such, some of these assets ought to continue to have value.

From the practical perspective of 'doing the sums', the credit default swap rate would need to relate to a senior creditor, and many life companies have not borrowed on such a basis. Clearly, some have borrowed on a subordinated basis, and the credit spread on senior debt could be imputed from this. However, most subordinated debt has a coupon structure such that it is expected to redeem on its first call date, typically ten years from its issue, and this is reflected in its market price. Consequently, you would need to extrapolate credit spreads over the very long term. This is also likely to be true for aspects of the liability swap and possibly even the asset swap.

The suggested methodology is not going to be a precise science, but one which could potentially produce valuable insights into a company's value. I would have liked to have seen a worked example included within an appendix, and would encourage the authors to provide one in any future literature which they may produce on the subject.

Mr K. Foroughi, F.I.A.: The paper emphasises the need to understand the context when using the phrase 'cost of capital'. I have come across three distinct uses of this phrase:

- (1) This is when setting top-down hurdle or discount rates to value a stream of future profits, in particular to value assets where there is no market price. Depending on the situation, this will include concepts such as the cost of equity, possibly set using the CAPM, as well as the weighted average cost of capital (WACC) approach. This definition is the classic corporate finance definition of cost of capital, the return which investors require on their investment, given the risks involved. When setting the WACC, the authors point out some potential pitfalls for management. For example, in ¶2.6.2 there is the use of book values rather than market values as weights to set the discount rate. The insurers with whom I have worked have not fallen into such traps when setting the WACC.
- (2) This is in the embedded value context. Here the cost of capital is an additional accounting provision applied as a discount to the capital required over and above the basic regulatory reserves. This somewhat artificial split of required capital is simply the result of life insurance regulation. Terminology starts to confuse, because one way of allowing for risk in an embedded value is to use the top-down WACC approach, using a cost of capital approach to set the discount rate, and then applying it to calculate a different cost of capital provision.
- (3) This relates to frictional costs. These are a component of both the authors' total company performance measurement model, described in Sections 4 and 5, and also of market consistent embedded values. With both these models the allowance for market risk is made automatically using a bottom-up approach cash flow by cash flow, and the phrase 'cost of capital' is a reference to frictional costs. Which frictional costs to include depend on the purpose of the valuation, and is a matter for debate. I would always argue for taxation and investment management expenses as a minimum. Regarding agency costs, I would agree with the authors' comments, in ¶5.8.4, that some allowance for agency costs is already made by setting expense and other assumptions with regard to recent experience. Setting appropriate best estimate assumptions, allowing for the full impact on shareholder value of fluctuations in experience, is my preferred way of allowing for non-market risk in a valuation. This may avoid or minimise the need for an additional agency cost factor. Regarding financial distress costs, the biggest sign of financial distress for an insurer is a sudden inability to write profitable new business. That is why, in making commercial decisions, you should always consider the impact on the embedded value and also the impact on the value of future new business. By doing this, you move away from the simplified straight line in Figure 12 to a curve much closer to the authors' model.

I am glad to see that the authors recognise that market consistent embedded values are one of the big ideas to emerge in insurance management in recent years. This is a key building block for their corporate model, breaking down their model into manageable component parts. This breakdown, looking at both the embedded value and the value of new business, enables firms to make sensible commercial decisions.

Professor P. Klumpes (a visitor; Professor of Accounting, Tanaka Business School, Imperial College, London): I am an academic with an Australian background, so you will hear my views on this topic slightly at odds with some of the others.

Obviously, investors, regulators, fund managers and employees are concerned about their performance and face strong incentives to assess the cost of capital for valuation, budgeting, compensation and fund management decisions. However, there is little consensus as to how the cost of capital should be calculated for financial firms. In the paper the authors provide some fascinating new insights into the cost of capital issue.

My comments will focus mostly on some technical methodological issues. There are three approaches to the cost of capital which I shall consider, and I shall fit the authors' views into these three. The first is what you might call the standard corporate finance textbook approach, where you are concerned about valuing the firm in terms of calculating beta and the expected return on capital using a standard asset pricing model such as the CAPM or the Fama & French free factor model. The authors dismiss this approach, and move on beyond it. They criticise beta because it is only a measure of systematic risk, and, for an aggregated entity, you see the financial statements, and you are not going to realise that you really want to calculate the beta for various different business segments. However, cost of capital researchers have done this, and they have got around the problem by using what is called a full information industry beta approach, where you can observe the market value beta as the weighted average of some unobservable betas of the firm's underlying business segments. Robert Merton and others have derived an asset beta for the firm's operating asset as a function of both the equity beta, a pension asset beta, and debt and pension liability betas. They then used that approach to gauge the efficiency of the stock market in pricing pension risk, which is surely an issue of interest in the U.K.

The authors are much closer to what I would call an accounting approach, because they introduce the concept of franchise value, which accounting centres would regard as a kind of intangible asset which you would then add on to the net tangible assets about which auditors worry. Strangely, that concept is not really fully defined, even though its concept is quite crucial to accounting. I assume, in actuarial parlance, that this concept is what you might call an 'appraisal value', the sum of embedded value, VIF plus new business. In terms of the accounting link, this appraisal value is seen in the non-financial firm setting as being what we might call inherent goodwill, which is then feeding into accounting base models of the cost of capital. We just add the future residual income or EVA under various assumptions to the book value of equity.

This is where we have a second problem, because franchise value is being used in actuarial science quite a lot nowadays, and is arguably part of a fair valuation replicating portfolio. Unfortunately, however, the concept is not so well defined for other kinds of financial institutions, such as banks and fund managers. Moreover, many non-financial firms exclude franchise value. Instead, it is just treated as an intangible, so that we might come up with some kind of real option in order to help us to explain the difference between market and book values. Interestingly, the market to book ratio is much smaller for financial firms than it is for non-financial firms. Franchise value is also going to be highly affected by leverage and solvency, but banks and insurers have at their disposal various tricks of value illusion and other financial engineering devices to help them manage their regulated leverage.

By contrast, when we talk about leverage for non-financial firms, the theory says that these firms are subject to onerous restrictive covenants. The literature will use an unlevered beta to do the calculations. Again, that procedure is less likely to be acceptable for financial firms, whose assets and liabilities will be measured using a mixed attribute market based, or fair value, system, than for non-financial firms, which would use a more conservative basis for measurement.

It is also not clear from the paper what the accounting principles or the GAAP are which underlie the cost of capital. Are we going to use the U.K. GAAP? Are we using IFRS? Are we using some form of regulatory GAAP, or are we using United States GAAP? Clearly, many of these are now being reported together by firms in their financial statements. GAAP quality can have a significant effect on the measured reliability of the assets.

I now consider the final approach to valuation, which I think is what the authors are also using, which is what I call the financial economics approach. This approach is slightly different, because it is actually modelling the assets and the liabilities of various kinds, including the franchise value, explicitly. The difference with the accounting approach, which is implicit, is that in this case we are developing the concept of the securitised portfolio. The authors used that approach to help them develop a hedge portfolio of assets and liabilities. Clearly, that is then supplemented with this notion of a credit default swap derivative to help them price default risk, and then come up with various scenarios. My main criticism here, as an accountant, is that for a financial firm there is a full complexity of various bundled or unbundled insurance and investment types of contracts, which firms manage and which are now being disintermediated for new financial products, and these are subject to evolving accounting rules on financial instruments and insurance contracts. So, there is more complexity here which we are not modelling.

If we want to move forward, we probably want to have a combination of both the asset/liability approach and the accounting cash flow approach as a more modern risk based approach to measure financial performance and to develop better concepts of cost of capital. So, we can think of a broad financial economics perspective, where we have three elements: franchise value; net tangible assets (which is just the net market value of assets minus the present value for liabilities); and a third variable, which I call the default put option. The main advantage of this third approach over the authors' approach is that you are seeing that the net tangible value is completely independent of what kinds of assets a firm has, and is independent of insolvency risk. Moreover, the default put option will separate out all the credit risk issues from the hedge portfolio, and it is treated as a put option game play with regulators which we model separately.

So, we need to make further progress in this area, and the authors have made a great start. We need a more consistent approach to reach consensus about what reporting standards we should set to model the hedge portfolio: "What is the mathematical context, how do you model uncertainty, what is the probability structure, the solution structure, required to implement the valuation model?" Also, we need standards to consider the implications of this model for decomposing various sources of earnings, so that we can think of earnings as being within a more comprehensive performance framework for financial firms. It is here where we need more in the way of rules and guidance for developing more consistent auditable and reliable estimates which will help us come up with cost of capital. This paper is a helpful and timely contribution to this end.

Mr I. B. Farr, F.I.A.: Considering Section 3, the paper points out the basic, but often overlooked, fact that the return on book equity of any company is fundamentally dependent on the proportion of its total value which is recorded as its book equity. A fund manager with very low book equity, but high franchise value, provides a high return on equity. A life insurer reporting on an embedded value basis, where a large proportion of total value is recorded as book equity, will return a relatively low return on that equity.

Another way of looking at this bias is that part of the return which is being divided by the equity is not actually being generated by that equity at all, but by the portion of total value which has been omitted. This differentiation of returns by source is easier to apply in some industries than in others. However, one clear example is for a life insurance company. Again, using an embedded value basis, you can see that the new business value component of the return represents the return on the goodwill, or franchise value, not the return on the in force book portfolio, or embedded value.

To track, or to target, new business values as a percentage of embedded value, as is implied by the calculation of an ROE ratio, does not seem logical. Why should the two be related? This is, however, what is implicitly implied by imposing a uniform return on embedded value across a global life group. So, how should returns on embedded value be targeted? A simple return approach is to segment the total return into different components and to set targets appropriate to each. The in force component of return rewards significant analysis of the market and

insurance risk to which the embedded value is exposed. The nature of those risks, and their market prices, give a good guide as to the appropriate target returns.

The setting of new business value targets is more difficult. While financial risk and return analysis might suggest how you avoid destroying value, the question is more often about how much value creation is required. If new business targets are to be meaningful to management in a geographically diverse group, they need to reflect many factors, essentially those encapsulated in the franchise value of the company or the business unit. These include the growth and profitability of prospects for each market, the historic performance of the business unit, and the scope for operational improvements, often as represented by competitors' results. So, benchmarking of performance against peers often has an important part to play.

The inevitable conclusion of many of these arguments concerning return on franchise value is that increasing value for shareholders is primarily a matter of managing and exceeding their expectations.

Mr C. D. O'Brien, F.I.A.: This paper provides another step forward to a better understanding of the links between accounting, valuations and risk.

The authors refer to economic capital and to RORAC, and then throw in some doubts, referring, in ¶3.7.9, to extreme percentiles not being the main determinant of expected market returns. Economic capital is the amount required such that a firm has a specified probability of meeting a defined liability over a defined time horizon. However, not only is that specified probability an extreme percentile, it is only one particular percentile. It is a Value at Risk (VaR) measure at 99%, or whatever is being used.

VaR is fine as a measure of risk if the 99th percentile is what we are interested in, and not anything else, but VaR gives no weight whatsoever to other risk levels, and, in particular, within the tail, whether the losses are relatively high or low. Indeed, VaR is not what we would now call a coherent risk measure. This suggests that VaR is not a good measure to use in economic capital, although it is still the main risk measure shown in the reports and accounts of banks. We could take the argument further, and question whether individual capital assessments for insurers, and perhaps Solvency II, are right in focusing on the capital requirements which arise just at one percentile. There may be better ways of measuring risk and determining the required capital.

Who bears the risk? In ¶5.10 the authors refer to why firms should spend money on managing risk. If a firm has diversified shareholders, this is not entirely a straightforward matter. However, there are valid reasons for spending money on managing risk; for example to reduce the cost of financial distress or to increase franchise value by enabling the firm to be more attractive than otherwise to customers and suppliers, who may then be more willing to enter into long-term relationships with the firm.

However, risk is also borne by managers. They, typically, have non-diversified wealth, and may prefer a more secure firm in order to protect their jobs and their earnings. There are a number of empirical studies of why firms manage risk; for example the extent to which they buy derivatives or buy insurance, and also evidence on why firms undertake acquisitions which are expected to reduce the volatility of their earnings. This evidence has found that firms differ in the extent to which they reduce risk, depending on, among other things, how the remuneration system for managers responds to risk; for example using salary, shares or share options, which are all different in this respect. This means that managers' interests can be relevant in risk management decisions. However, it also implies that shareholders should incentivise managers if they feel that managers are being unduly risk averse.

It is also customers who can bear some of these risks. A feature of many financial services firms is that they issue long-term contracts, or at least contracts which have some long-term dimension, such as a bank account, which is likely to last several years. It may be that customers are bearing some of these business risks. This is clear in the case of with-profits life insurance policies, but may also apply elsewhere as well. If a bank or an investment firm suffers some adverse event, it may be able to increase the charges which it makes to its customers. In a competitive market the customer would move to a firm which has not suffered these adverse

events. However, markets for financial services products are not always competitive, partly as the products are often complex and the customers ill-educated. So, consumers can end up bearing some of those risks.

Indeed, if we were to think about mutual financial services firms, which include building societies as well as mutual insurers, then, yes, the customers there are bearing risks. However, customers are unlikely to have diversified portfolios, and therefore have a different perspective from shareholders. Capital market theory has certainly enhanced our understanding of how shareholders can manage risk to increase value, and we are grateful to the authors for their work on this. We may also need to consider how mutual financial services firms use their capital and manage their risks.

Mr D. W. Dullaway, F.I.A. (in a written contribution that was read to the meeting): I shall restrict myself to expanding on the authors' discussion of RORAC. This is a performance measurement approach used in many large financial conglomerates, often to compare the performance of such different businesses as life and general insurance, banking and fund management. It is usually defined as some form of annual profit measure, possibly embedded value earnings for life insurance companies, divided by some measure of the economic capital required to support the business.

The authors bring out a number of reasons why RORAC is a poor method of performance measurement for financial firms, including biases in the determination of profits and in the determination of economic capital. I also mention another problem with RORAC which is not addressed in the paper; that RORAC does not correctly reflect investors' views of the return required for taking different types of risk. It is generally accepted that well-diversified investors require greater returns for taking on market risk (which they cannot diversify) than for taking on similar amounts of non-market risks (which they can largely diversify). Any performance measurement system should reflect this fact. However, RORAC implicitly assumes that investors require the same return on all types of risk.

We can see this if we consider two life insurance products, defined so that they both generate the same level of profit and require the same amount of economic capital, but have different risk exposures. For example, consider a term assurance product (mainly exposed to non-market mortality risk) and a with-profits bond (mainly exposed to equity market risk). As both have the same economic capital for a given level of profit, both will generate the same RORAC, so that, under this performance measurement approach, both seem equally attractive. Yet most investors would prefer to be writing the term insurance product than the with-profits bond, as the term insurance contains less market risk. Therefore, without adjustment RORAC will lead companies to take too little diversifiable risk and too much market risk, from an investor's perspective.

It may be possible to adjust RORAC, either by using separate hurdle rates for each type of business or by risk adjusting the profit to strip out the cost of risk, as in the so called RARORAC (risk adjusted return on risk adjusted capital) approach, but this is moving a very long way away from the idea that economic capital can be used to allow for risk. After all, either of these approaches requires an entirely separate performance measurement system for valuing risk — which is the problem which RORAC was originally meant to solve!

RORAC effectively assumes that investors cannot diversify away non-market risk; that is that it considers valuation from a company rather than from an investor perspective. This might be appropriate for a partnership, a venture capitalist, or, indeed, if we are concerned for the benefit of senior management (whose fortunes are largely tied to the firm which they run); it is not appropriate for a publicly traded company.

Mr A. H. Silverman, F.I.A.: In ¶2.3.2 it states that the normal required shareholder return calculations put life insurers at a disadvantage, and suggests that they are unfair to life companies, at least in down markets, a point related to that made by the previous speaker. The phenomenon described is just part of a long-standing and well rehearsed argument as to why fund management companies, and, perhaps more recently, life insurers, deserve a beta of more

than one, and there is a very strong argument that any business where the underlying earning capacity is directly impacted by markets should have a beta greater than one.

In ¶2.6.2 the authors suggest that most WACC calculations use the accounting values of equity and debt as weights rather than the market values of these items, and that this is unsatisfactory. However, the reasons for using accounting numbers in other industries are not necessarily those stated in ¶2.6.3, when one is looking at the accounts of other industries. They are not because market values of subsidiaries are not available or to avoid volatility. The primary reason is that management, or analysts, are usually assessing the success of decisions around the deployment of organic cash flow. The question being asked is: “Is the company making good use of capital?” The market capital is not available as organic cash flow. It is simply not what management has invested. A large part of the franchise value is genuine goodwill. The accounting equity is used to a degree as a proxy for what management and investors have put in. In other words, these calculations are a substitute for the cash flow return on investment (CFROI), which is alluded to in ¶3.1.8, and which is not really accessible to direct calculation, certainly by an external analyst. Of course, the CFROI is close to the internal rate of return with which many of us will be very familiar.

However, the problem is that in life insurers, accounting equity, whether it is GAAP or MCEV, is not close enough to the amount of cash invested to produce a meaningful answer to that calculation. Indeed, if the accounting ROE has too distant a relationship with the CFROI, which is where we are in the life insurance industry, then that is indeed a serious drawback for the fundamental relevance of the accounts to users.

Section 3 and, to a degree, the rest of the paper, try to put things right by constructing an ROE which equals the total shareholder return, but, again, this does not serve the same purpose as the ROE which is calculated for other industries. A large part of the answer will remain the IRR.

Mr P. D. Needleman, F.I.A.: I shall add a few comments on the subject of managing risk, capital and value, primarily from the perspective of a life insurer.

First, why is this subject so important and currently attracting so much attention? Over the last five years the life industry has focused on risk and capital — developing better tools to help measure risk, and taking action to minimise losses and preserve its depleted capital base, largely defensive measures. Over the next five years, it is to be hoped that the industry will be re-focusing on creating value, and it will need effective risk management and value measurement tools to do this. Thus, the topics covered by the paper are very relevant.

Why, therefore, does it appear to be so difficult to come up with appropriate metrics and measurement tools to assess performance and returns for life companies? It is clear that some measures derived from other industries are, at best, inadequate on their own, and, at worst, may drive inappropriate or value destroying actions.

The answer is, in part, because the risks to which the financial services industry is most exposed, such as asset/liability risk, credit risk and insurance risks, are different from the risks which are most important for other industries. We also have to recognise that within financial services there are two views of everything — the policyholder and regulatory perspective and the shareholder perspective — and we easily confuse these two views.

So, for example, when thinking about risk, the policyholders or regulators are concerned with the ‘risk of ruin’, whereas, from a shareholder perspective, the concept of ‘risk and return’ is more appropriate; that is: “What is the required shareholder return for a particular type of risk (the ‘price of risk’)?” Similarly, there are at least two differing concepts of capital: the regulator or policyholders’ view, being the capital available to secure the benefits; and the shareholders’ view, being the total market value of the company. So, there seems little point in using regulatory capital — or, for that matter, economic capital — to measure shareholder returns. This would only be appropriate if these were good proxies for the total market value, which is often not the case, especially for a company with any significant franchise value.

So, we find that management is faced with a bewildering array of metrics which often causes great confusion. An effective financial management framework needs to ensure that the metrics

used are consistently calculated and appropriately used. They should be based on a clear and consistent underlying methodology which links risk and capital and value together.

Any such framework needs to be theoretically robust, but also practical to implement, and it should drive consistent 'value creating' decisions across all areas of a company, such as life, general insurance and asset management, and at all levels of the business. I believe that a market consistent, value-based approach can meet these requirements.

These changes in financial management are already influencing the ways in which some companies are running their businesses. The most obvious examples are in the area of asset/liability management, where, increasingly, sophisticated dynamic hedging and other techniques are being used. Most companies are also changing the way in which they price for risk in their new products, and last, but not least, in many merger and acquisition transactions, market consistent embedded values are an important part of the array of information influencing transaction prices. I expect that we will hear a lot more about these topics over the coming months and years.

Professor R. Macve, Hon. F.I.A.: I am a chartered accountant, and have been teaching my students about the current proposals from the IASB and the FASB on the future of accounting for business combinations, most of which appear to represent 'fantasy accounting'. However, they are obviously very relevant to this paper, because they are all about accounting for 'goodwill' — and, as I understand it now, goodwill is the franchise value which the authors are discussing. So, we do need to get the right treatment for it.

My comments are primarily directed at increasing the understanding between the accounting, the actuarial and the financial economics literatures. In Section 3 the authors convincingly argue the dangers of looking at returns on equity based just on the accounting numbers, and not building in the fact that shareholders are expecting a return on the stock market value of their investment, i.e. including the franchise value. However, there must be some feedback between these things, because the stock market value of the investment on which shareholders are expecting a return has been set, primarily, by expectations that the company is going to be able to deploy its internal assets and activities sufficiently well to earn a return which justifies the franchise value, and thereby the market value at which the shares currently trade. So, the question is: "How far do the numbers within the accounts provide information which enables the stock market to set a realistic value for the firm, on which it can then deliver an appropriate risk adjusted return?" Obviously, if the firm fails to do this, those market expectations will change, and so will the franchise value and the stock market value.

Still in Section 3, there is a reference to the paper by Solomon & Laya (1967) as being one of the earliest demonstrations of accounting errors. I think, given notice, that I could probably find some from much earlier writers on accounting, so, maybe we need to go back quite a bit further.

This brings me to my last comment, which is on Section 4.2, and, in particular, ¶4.2.12. Professor Klumpes alluded to this. This concerns a well-known theorem about accounting, which corresponds to Modigliani & Miller's famous theorem (which is, of course, not that dividends are irrelevant, but that — in the absence of taxes and other frictions — the *timing* of dividends is irrelevant to the firm's stock market value). The corresponding accounting theorem is not that accounting is irrelevant, but that the *timing* of accounting profits is irrelevant. In other words, it does not really matter whether you recognise the profits now or later, at the end of the day you can only recognise the same amount; and the authors have captured this theorem very well. Unfortunately, it is not new. There has been much work in the past ten years or so, particularly by James Ohlson, who is now at Arizona State University (starting with Ohlson, 1995), which has spawned a whole literature on this 'accounting irrelevancy' theorem, which it would be worthwhile for the authors to explore further.

My overall impression of the paper is that it is very helpful in bringing out some important accounting issues. There is a growing relationship between the actuarial literature and the financial economics literature which is having an impact on practice. Similarly, there is a growing relationship between the accounting literature and the financial economics literature which is having an impact on practice. Clearly, what we need is a growing relationship between

the actuarial literature and the accounting literature, and then we shall have completed the triangle.

REFERENCE

OHLSON, J.A. (1995). Earnings, book values and dividends in security valuation. *Contemporary Accounting Research*, **11**, Spring, 661-687.

Mr T. W. Hewitson, F.F.A.: I shall focus on the ongoing discussion within Europe on what is called Solvency II. For this Solvency II project, the basic principle at which many of us are looking is that the provisions for insurance liabilities within the financial returns should be essentially the amount which will be needed to allow a transfer of the business, on commercial terms, to a well diversified third party with a good credit rating. Therefore, these provisions would be assessed on market consistent principles for hedgeable risks, but with appropriate margins to cover the possible misestimation of the means of the parameters for unhedgeable risks, as viewed over the outstanding duration of the policies.

There are, then, at least three possible conceptual approaches still being considered within Europe which could achieve this purpose. The first option would be a percentile approach to the assessment of risk margins for longevity, general insurance claims and other risk factors, which might, for example, be set at a level sufficient to provide 75% confidence that claims can be met, this being the current benchmark which has been suggested by the European Commission; or, indeed, it could be set at some higher level of confidence, as we understand that some European supervisors would prefer.

The second option, which would be of interest in the context of this paper, is what is called a cost of capital approach, which has been advocated by some of the larger European players within the industry, and this is based on the projected cost of servicing either the economic capital or some regulatory capital, as held by firms, each year. However, this approach would still need much more developmental work, and this is an area where the Profession could make a useful contribution.

The third option, which comes, in part, from some proposals from the German insurance industry, is really a U.S. GAAP style approach, principally applied for life insurers, whereby some fairly arbitrary margins, say 15%, are added to the best estimate rates for mortality, persistency and expenses.

In addition to the assessment of these provisions and financial returns, there will then be added a regulatory capital requirement designed to ensure, with a high level of confidence, such as perhaps 99.5%, that there would be sufficient assets to cover the insurance provisions in 12 months' time. In other words, for the purpose of setting the capital requirement, there would be some sort of stress test, or scenario, or standard formula designed to ensure that, if this very adverse event occurs in the next 12 months, there will still be enough assets to cover the market consistent value of liabilities assessed in 12 months' time.

One way forward on all of this, which might be seen by many people here, would be the industry developing a sound, but pragmatic, approach to determining an appropriate level of economic capital to support its business, and, indeed, for individual business lines, which might then be taken into account, both in the pricing of policies and in the overall management of business. However, such an ideal might still be some considerable distance away.

Meanwhile, the percentile approach to assessing provisions does not appear to be very popular with many people in the U.K. industry, and the alternative cost of capital approach, which has been advocated by some of the larger players, is still in need of considerable further development. Accordingly, we might end up with some standard approach within the E.U. for the setting of provisions which, at least for life firms, might comprise, as a proxy for the percentile approach, some U.S. style pads attached to the various assumptions for the unhedgeable risks such as mortality, persistency and expenses, but, at the same time, with an option for firms to convince the regulators that they can apply some internal modelling approach instead, perhaps based on this cost of capital.

For the internal modelling approach, regulators would be looking to see that such models were fully integrated into the running of the business.

There is, of course, still time for the Profession, both through the Groupe Consultatif in Europe, and through the IAA, which is assisting some parallel work in the IAIS, to inform this ongoing international debate. However, the cost of capital approach to the assessment of provisions will now need to be developed fairly rapidly if it is to be taken forward as a viable option within Europe.

Mr R. J. Baddon, F.F.A.: This paper sheds light on the trickiest issues for the value management of financial services. Much of my work is in the merger and acquisition context, so that I deal, on a daily basis, with the challenges of producing values and comparing them to market values. I particularly welcome the efforts which the authors have made to reconcile their performance calculations to market value.

If we step back and look at embedded value; whilst this looks like a simple idea, there are tricky issues in practice, including: the allocation of tax and expenses, including irrecoverable tax and changes of tax bases as part of a transaction; the allowance for the firm's own possibility of failure, which people tend not to look at; and the quantification of the costs implicit in the three tiers of capital.

It is tempting to push all the difficult issues into the goodwill box. If we think that agency costs or possible corporate failures are too hard, we may convince ourselves that these effects are only future business effects, and therefore stop thinking about them in the context of embedded value or in pricing computations.

Worse still, these difficult items may turn out to be balancing items. I am not the first to notice, or to comment on, how similar in numeric terms an embedded value on a European basis is to a traditional embedded value. This sometimes seems to happen because an increase in the stated guarantee cost is associated with the removal of that agency cost or operational risk cost which is embedded in the risk discount rate.

In a transaction context, parking hard questions in the future new business box is not an answer at all, particularly in the context of recent transactions where people are looking at closed blocks of business. Similar to that in the paper, we are interested in the firm's whole value, or enterprise value. It is pointless to persuade oneself that agency costs are zero, just because they are easier to deal with in that way or because the assumption is offset elsewhere in the computation. This paper offers an exciting way forward. What matters is the firm's total value, which includes anticipated asset liability margins net of tax, less the tricky items of agency costs and credit risk, the credit risk of default to the shareholder. Nevertheless, you can estimate these on an aggregate level by calibrating back to market value. The value in force is then an attribution of that franchise value, considering the marginal impact of the business. So, in summary, this framework gives a unified measure of value which allows us, probably for the first time, to understand fully the impact on shareholder value. That is not to say that we should not consider economic capital and all the other important regulatory measures, but when we are looking at an enterprise, we need to know the value to the investors.

This framework gives us a robust, market consistent technique, which moves us significantly on from embedded value to a broader enterprise value approach.

Mr T. J. Sheldon, F.I.A.: Banks and insurance companies already employ a wide range of capital measures. We have accounting equity, market capitalisation, risk-based capital, regulatory capital, economic capital, and now, for U.K. insurers, we have individual capital assessment (ICA) and its offspring individual capital guidance (ICG). Do companies need yet another measure — optimal capital, or optimal capital structure — as proposed in the paper, where the authors make a strong case for the adoption of this new capital measure?

There is much similarity between some of the currently used measures of capital. For U.K. insurers, in particular, risk-based capital (RBC), economic capital (EC), and ICA/ICG are really variants of one measure, based on VaR; that is the capital required, at some chosen confidence level and time period, to meet the risks accepted by the business. It seems likely that, over the

next few years, regulatory capital will move to an RBC approach. The purpose is to demonstrate to customers and shareholders that the company holds sufficient assets to meet its liabilities to the former group with an acceptable degree of confidence.

Does it make sense to use RBC or economic capital for other purposes? The paper gives an example of the use of economic capital in a RORAC framework, as used by many companies in published reports and accounts. Here the aim of demonstrating attractive rates of return on capital is consistent with the desire to show that net assets are a multiple of EC. Reducing EC improves both the apparent rate of return on that capital and its cover.

While the flaws in this approach, including accounting bias, are exposed by the authors (and in the relevant references in the paper), it is interesting to note that, in the example given in Figure 4 and Table 5, the total shareholder return (TSR) is very close to the ROE for all industries over the period. Does that suggest that, on average, ROE gets it right, or is this just coincidence? Whatever the answer, the use of RORAC to take business decisions leads to the development of less capital intensive business lines and strategies aimed at reducing capital requirements, neither of which will necessarily result in maximising shareholder value.

From a shareholder perspective, a different framework is required. Demonstrating an adequate level of RBC or EC is a necessary requirement in order to continue to trade, but it is clearly not the whole story. First, shareholders would like a return on market capitalisation, the share price, and not just on part of it. Secondly, shareholders should be interested in maximising the value of their firm, relative to the capital contributed. This, as the authors demonstrate, is equivalent to maximising the franchise value.

Most life insurance companies use embedded value (EV) techniques to monitor and to report progress in enhancing firm value. EV is good at measuring how an uncertain franchise value is being converted to an in force value through the writing of new business, but it needs extending, at least for internal management purposes, to embrace the management of franchise value.

The search for a model of total firm value and the optimal capital structure requires a calibration of the various components of firm value (comprising net assets, value of in force, future new business value and frictional costs) to the observed market value, and a model of how the value depends on the capital structure, or the amount of capital and gearing. This has presented a formidable challenge, as it involves two unobservable items — future new business value and frictional costs — both of which depend on several parameters.

The hedge model developed by the authors provides a way of estimating some of the frictional costs, in particular the main component of financial distress costs, the loss of franchise value, from other market data, thus reducing the number of parameters which need to be inferred to give the desired fit to the market value. This should make models for firm value, frictional costs and capital structure more robust and easier to apply in practice. I hope that this excellent paper will encourage more companies to investigate this approach to capital structure, and to develop its applications in pricing and in the other key financial decisions which our clients make.

Professor A. D. Wilkie, C.B.E., F.F.A., F.I.A.: I have a comment on one quite small aspect of this paper, which is not an integral part of the authors' argument at all, but, nevertheless, it appears without much comment. Figure 1 shows mean equity returns in excess of T-bills from 1900 to 2000, quoted from the very large survey by Dimson *et al.* (2002). The geometric and the arithmetic returns are shown. In most countries, in the bottom three-quarters of the table, the difference between these is about 2%. Now, the mean of a log normal distribution with parameters μ and σ^2 is $e^{\mu + (\sigma^2/2)}$, the median is e^μ , the long-term return is e^μ , and the extra $\frac{1}{2}\sigma^2$ term is the difference between the geometric and the arithmetic means. A typical standard deviation for annual share returns is 20%, or 0.2: if you square that you get 0.04; halve it and you get 0.02; and that explains the 2% difference.

The authors do not quite say that one should look at this table in order to find out what shareholders will require in future, but it is actually very misleading. First of all, the final date for the period at which Dimson looked is 2000, and there was a very significant asset bubble in the last few years of the twentieth century, so that share prices at the end of 1999 and for some

time during 2000 were at very high P/E ratios or price/dividend ratios (i.e. very low dividend yields). I do not know what the figures were for 1900, but I doubt very much whether the rating of shares in 1900 was as high as it was in 2000. So, part of that return is to do with the asset bubble, and if you believe the table, you have to believe that the asset bubble will continue in future until P/E ratios go to even higher levels and dividend yields to even lower levels.

A further point is that the countries at the top of the list, with the highest equity returns, Italy and Germany, maybe Japan, and France had very high inflation during this period. Germany had extraordinarily high inflation after the First World War and hyperinflation in 1923. Italy and France had pretty high inflation in the 1930s and during the Second World War.

It is interesting to note that the old Italian lira, the old French franc (which was 100th of the newer French franc), the Swiss franc and the Belgian franc all had the same value in 1900 at an exchange rate of 20 to the pound. If we ignore the introduction of the euro, the exchange rates with the pound would now be about 3,000 for Italy, 1,000 for France, 50 or 60 for Belgium and two for Switzerland. That shows the difference in the value of money in those different countries.

It is a nice concept that Treasury Bills keep up with the rate of inflation and give a small real return plus a compensation for inflation, but, in practice, they do not, certainly not when inflation is exceptionally high. Inflation in the U.K. in the mid 1970s was 25%, but interest rates did not get up to 28%, Bank rate being more like 15% or 17%. Long-term interest rates were a little higher than that, but not over 20%. On the other hand, share prices frequently go down when there is an initial burst of inflation, but keep pace with inflation in the long run, because dividends and earnings do go up in the long run in line with other prices, even in countries where there has been very high inflation.

So, what I think of as fairly high figures are really the result, in many countries, of one or more bouts of inflation. Inflation does not come along randomly year by year, but in attacks. You may have an acute attack or an extremely acute attack, like Germany; a fairly acute attack, like Italy and France; or chronic attacks, as most of these countries had during the 1970s and the 1980s. Therefore, if one is wanting to use these figures as a projection for the future, or as the required rate of return for the future, you are having to build a bout of inflation which *you* expect, but the rest of the market does *not* expect, into your estimates. I do not see how you can be consistent with market values on that basis. The market is not sensible if it is assuming, first of all in relation to gilts, that inflation would be no more than 2.5% or so, but in relation to shares that it would be of the order of somewhere between 5% and 10%. One can derive that from these figures.

I wonder whether, given that companies of all sorts, financial and others, tend to use high hurdle rates for future prospects, they are not finding those capital investments which would be worth doing, with some companies returning capital to shareholders, and all because they are using too high desired rates of return. If you are wanting to work out whether an investment is profitable or not, you need to put in some numerical hurdle rate, one way or another. Top managers may well like junior managers to use a high hurdle rate, because junior managers are very optimistic about their particular ideas. This is just a way of counterbalancing optimism. One has to be careful about Figure 1 and similar figures and the data quoted by Dimson *et al.* (2002) and by many other people.

Mr D. J. McLean, F.I.A. (closing the discussion): Much of the discussion has been about the right return on capital, what is the right part of capital, and about how we get round this. There was one near dissenting voice as to whether we needed yet another measure. There did not seem to be much support for RARORAC. There was a clear conclusion that we should support an emphasis on the return to shareholders. However, no one picked my favourite bit, which is in ¶3.1.7, the reference to the return on iron. This is a very valuable ratio which we should be looking at in the financial sector!

The paper is clear as a theoretical description, but I would love to see a worked example. This is a challenge for us, to go back and to look at a particular business, and to see how close we can get to a real worked example.

One technical point, which was queried by a couple of speakers, occurs in Section 4.9. The authors refer to the zero recovery credit default swap. It is right that what we are looking at there is not a conventional credit default swap, but there is still a subtlety in terms of what franchise value is still preserved, even at the point of the default trigger. So, although it is not a conventional credit default swap, it may be that we need to define the credit default swap there in a subtly different way. At the very least, as at least one speaker mentioned, there is the possibility of the tax losses being valuable to somebody else, even if the franchise value in one or other of the brands is not.

There is a challenge arising from Figure 15. If we are looking at the interaction between tax and capital, what it seems to be saying is that the current tax regime works to encourage people to have less capital. That would seem to be a somewhat perverse situation for those setting taxes and those setting capital ratios. I am not sure how I would resolve that circularity, because the authorities do not want something which lowers the tax take. It is hard for us in the financial sector to demand that and to be taken too seriously. So, we need something which is probably neutral from a tax take perspective, but at the same time does not produce a one-off financial sector tax arbitrage. This is something which we could take further. The challenge to try to resolve is to give greater security for any given level of tax take, without the incentive pointed out by the authors.

Mr A. D. Smith (replying): I shall pick up a few points which have been raised in the discussion, starting with Professor Wilkie's comments about estimating long-term returns. Of course, he is absolutely right that estimating long-term returns over the next 100 years is even more difficult than estimating what they were over the last 100 years, and that is hard enough, and it is not surprising that it is contentious, because there are many different theories. The good news, from our point of view, is that we do not advocate the cost of capital methodology which requires you to do that. Instead, in Section 2.4 we suggest that, for a financial firm, if you mark to market the performance of that firm, then the right thing with which to compare it is the actual market return over the previous year on a replicating portfolio. So, if, for example, you were a financial firm with a beta of one and no exposure to interest rates or any other market figures, then you might want to compare your return to the actual return on the stock market in the previous year. You might say: "Would that not introduce all sorts of volatility?" It does, both to the firm's return and also to the market return, but these two cancel out. So, you are then freed of this need to estimate long-term returns. I am not saying that long-term returns are useless for all purposes, but, for this particular purpose, we could get round some of the difficulties which Professor Wilkie mentioned by going about it in a different way.

I now consider embedded values and market consistent embedded values. We deliberately wrote our paper to apply to financial firms generally, but much of the emphasis in this discussion has focused on life assurance and the embedded value issue. We must not lose sight of how odd it looks to the outside world when we value liabilities prudently, and then try to nudge them back to reality using value in force adjustments.

Imagine that you are the finance director of a car manufacturer, that your warehouse contains some blocks of steel, and that you know the historic cost of those blocks of steel. These are somehow turned into components and assembled into cars, distributed to dealers, and then cash finds its way back from the dealers through the system, and some of it gets back to you. You then use some of it for dividends and some of it to buy more steel, pay workers and for other things. All of these you forecast meticulously month by month in a detailed business planning model.

A life actuary now knocks on your door, and tells you to extract from your planning forecast the passage of a particular cohort of steel through the manufacturing, assembly, distribution and cash settlement process. You have to discount the cash at the end and compare it to the book value of the steel, and the difference is the value in force of your steel inventory. Your actuary suggests that it would be a good idea to build a management framework recognising the value of steel in force as an additional asset. However sensible that seems to a life actuary, most industrial firms would find the suggestion more entertaining than practical. I think that actuaries

need to wake up and 'smell the coffee' — the outside world does not, and probably never will, understand embedded values.

The invention of market consistent embedded values claims continuity with previous, slightly less mysterious, versions of embedded value. Continuity is great from the point of view of actuaries, but from an outside perspective it is harder to see why it is such an advantage. On the other hand, our method provides an impressive arsenal of market consistent valuation tools. My preference is to continue to extend these tools, particularly in relation to credit risk. When we talk about market consistent embedded value, what that means, in practice, is calibrations to interest rates, equity markets and derivatives on each of these, not usually involving the calibration to credit risk instruments. Our hedging construction gives some idea of how we can extend it to credit risk and also to the consideration of frictional costs.

Having got that enhanced toolkit which is consistent with credit markets as well as with the more conventional markets, my preference would be to apply that toolkit directly to the liabilities. In that sense I agree with the IASB, that this method is the way forward.

Professor Klumpes and others have mentioned the limited liability put option, the option which shareholders have to default on liabilities. He is right that the literature places quite a heavy emphasis on this. We did a great deal of numerical work in building this paper, and we found it pretty difficult to construct an example where that default option was material relative to the contingent loss of franchise value in financial failure. That means, other things being equal, that managers who maximise shareholder value should be risk averse rather than risk seeking, which is the opposite conclusion to that given in the finance literature, which looks only at the default option, and not at the loss of franchise value.

The numerical magnitude of these two is important for its implications of how you run a firm. However, even in theory this limited liability default option only affects decisions of very high levels of leverage, where the capital is well below regulatory requirements. In practice, the decision in that situation is simple. You have to raise more capital or you will not be allowed to stay in business. So, while we accept that it would be nice to have an explicit limited liability put option number, it is a considerable effort to calculate it, and we question whether that will be repaid in terms of improved decision making.

I now respond to some comments of both Professor Klumpes and Professor Macve about accounting irrelevance, and I am grateful to Professor Macve for drawing our attention to Ohlson (1995). They talked about accounting irrelevance, that it does not make any difference how you account, it is just a matter of timing. In some ways that is true, but it is not quite true when it comes to the default. The reason is that the measurement of whether you are allowed to continue trading or not is essentially an accounting measure. That is an important point. I have lost count of the number of times when people have said: "You should not focus on accounting principles; you should focus on economic principles." The underlying idea is that you can somehow abstract from the messy details of accounting standards and go to a beautiful, intellectually perfect world of economic values. The reason why that does not work is when you get a call from the FSA telling you: "We are going to close your business." At that point you have to close for new business. The FSA's basis for doing that has to be on some sort of accounting measure.

If you are considering how much capital you have to hold, the sensible starting point is to look at the accounting measure which would be used to judge whether or not you are allowed to continue trading. That might not be the accounting measure that it seems, because there are a few stops which you can pull out in terms of securitising and turning inadmissible assets into admissible ones, and so on; but, in principle, you have to have some accounting measure there if you are looking at the risk of financial failure. Otherwise, the shareholder equity is always going to be a positive number. You can go bust by having negative shareholder equity if you include the franchise value in that.

I now pick up on some remarks from the opener and from Mr Hewitson, about the regulatory implications of cost of capital and introducing cost of capital as a possible liability. It has already happened at the end of 2005 in Switzerland. Swiss firms had to calculate and to project one year ahead and to work out, not only the discounted value of their liabilities, but also

the cost of the capital which would be required to run off those liabilities. In that sense, it does actually make some sense. The regulator does not want you to have just enough of the present value; it is argued that, in one year's time, after your stress scenario, if you need to get more capital you do not need to have the capital on that day, but you need to be able to get to the capital in order to continue running off the business in an orderly fashion. So, logically, it does make sense. However, the numbers which I have seen do not make much sense. There is a submission from the CRO forum which shows that the cost of capital adjustment seems to increase the liabilities by some small fraction of 1%, which is rather less than the uncertainty in assessing the liabilities in the first place. It is hard to construct examples where it is material or makes a lot of sense.

What I suspect is that, whichever of Mr Hewitson's alternatives (a percentile approach, a cost of capital approach or an arbitrary risk margin approach) is used, when a report lands on a regulator's desk claiming that you have enough capital, there is a relatively small number of options which the regulator has, and it does need to be able to compare companies. I suspect that that comparison will end up being relative to a list of rules of thumb. I do not think that the regulator has any access to scientific research which the rest of us do not have, so, to the extent that we are using arbitrary rules of thumb in our own businesses, they probably would not be any less arbitrary than the regulator's rules of thumb.

So, although you have three approaches, which look rather different in theory, one wonders whether, in practice, you will end up having the same numbers from all three approaches, and the only effect of the chosen approach is in what rhetoric you have to produce in the 20 kilo report which justifies the answer which you have achieved.

The President (Mr M. A. Pomery, F.I.A.): It only remains for me to express my own thanks, and I am sure the thanks of us all, to the authors for the paper. We have had a very good discussion. I noted a number of challenges, going forward, for the Profession in a number of different areas. I ask you to join me in thanking the authors, the opener and closer and all of those who participated in this discussion.